2024 EPIC PROCEEDINGS

ISSN 1559-8918
OPEN ACCESS SINCE 2005
www.**epic**people.org

EPIC is a global community of researchers, creators, and innovators doing ethnography for impact in business and organizations.

epicpeople.org

EPIC is a nonprofit membership organization, global community, and annual conference that supports the learning and leadership of practitioners who use an ethnographic and sociocultural lens for sensemaking and transformation.

Working in every sector and industry, EPIC members help organizations make strategic decisions, create value, and navigate uncertainty by understanding people within cultural ecosystems.

The articles published in *Ethnographic Praxis in Industry Conference Proceedings* (ISSN 1559-8918) are selected through competitive, double-anonymous peer review, which is overseen by independent committees. The *EPIC Proceedings* have been published continuously since 2005 and are freely available at epicpeople.org/library. For permission, contact info@epicpeople.org

Learn about EPIC Membership →

Contents

Papers

Case Studies

PechaKucha Abstracts

Committees

Conference Chairs

CARRIE YURY, Conference Co-chair
Managing Director, Head of UX Research, Digital Products & Channels, JPMorgan Chase

LEE CESAFSKY, Conference Co-chair
UX Researcher, Meta

Papers

Shaheen Amir (co-chair), *California Office of Data & Innovation*
Juliana Saldarriaga (co-chair), *A Piece of Pie*
Maya Costa Pinto, *University of Melbourne*
Vishnupriya Das, *Gemic*
Rogerio de Paula, *IBM*
Emory Edwards, *University of California, Irvine*
Britta Fiore-Gartland, *Salesforce*
Yuliya Grinberg, *Drew University*
Stefan Laage, *Brandwatch*
Annie Lambla, *Coloplast*
Joyce Lee, *Salesforce*
Scott Matter, *University of Technology Sydney*
Emanuel Moss, *Intel*
Taylor Nelms, *Filene*
Claudia Ruge, *NBCUniversal*
Charlie Strong, *ReD Associates*

Case Studies

Autumn Sanders (co-chair), *Quire Consulting*
Anna Wojnarowska (co-chair), *Google*
Hannah Anokye, *Intuit*
Marcelo Fagundes, *Accenture Song*
Camille Kramer Corbariaux, *YUX Design*
Oskar Korman, *Alice Labs*
Victoria Lowerson Bredow, *UC Irvine*
Meghan McGrath, *IBM*
Deirdre McGruder, *UX/Design Researcher & Manager*
Marie Mika, *Frontrow Health*
Sunandini Pande Ray, *Kantar*
Heli Rantavuo, *OP Financial Group*
Adisorn Supawatanakul, *Teak Research*

PechaKucha

Robin Kwong (co-chair), *Wall Street Journal*
Briget Monahan (co-chair), *Google*
Antoinne Clay, *Microsoft*

Vyjayanthi Vadrevu, *Rasa.nyc*
Enya Trenholm-Jense, *Gemic*
Renu Zunjarwad, *Truckstop*

Peer Review Manager

Tiffany Romain, *Intel*

Artistic Interventions

Karl Mendonca (co-chair), *Google*
Rebekah Park (co-chair), *Gemic*
Elizabeth Anderson-Kempe, *Google*
Maria Pitallano

Interactive Experiences

Makalé Faber Cullen (co-chair), *PLASTARC*
Jennifer Giroux, *Giroux Impact*
Toluwa Awodiya, *Survey Monkey*
Todd Palmer, *Diversity in Design*

Salons

Abbas Jaffer, *Circle*

Graduate Colloquium

Rachel Ceasar, *Culture of Health & Tech Consulting*
Mahta Emrani
Renee Ho, *Lime*
Sylvia Tolbert-Long, *Know More Marketing*
Tayler Ulmer, *Aescape/Griot Recruiting*
Josef Wieland, *Research Strategist*

Student Volunteers

Nik Bonovich, San Jose State University
Aino Halonen, UXD Lab; Santa Monica College
Anisah Herdiyanti, Florida State University
Elisa Krieger, San Francisco State University
Jason Leach, University of Ottawa
David Nevers, Santa Monica College
Alison Parrales, Santa Monica College
Jessica Slattery, *University of California, Irvine*
Joseph Sotelo, *Santa Monica College*

Accessibility

Kate Mesh, *Open Inclusion*
Tony Moaton, *Cultural Researcher*
Gregory Weinstein, *CVS Health*

Local

Joan Afton
Steven Garcia, *Culturegraphics LLC*
Kate Sieck, Toyota Research Institute
Yushi Wang, *uxoffer*

Proceedings

Ea Arnoldi, *Human Insights Senior Consultant*
Jenny Callans, *Duo Security*
Reha Atakan Çetin, *blurple studio*
Yunshan Li, *RHIZOME Culture Consultancy*
Patricia Machado Alvarez, *Genentech*
Monica Romano
Joshita Yadav, *Indraprastha Institute of Information Technology, Delhi*

Sponsors

EPIC is a nonprofit, volunteer-driven organization. Our programs and publications are created by independent committees and invite diverse, critical perspectives. Organizations that sponsor our annual conference provide essential financial support that make our operations, open-access publications, equity programs, and scholarships possible.

Platinum Sponsors: Atlassian, dscout

Gold Sponsors: Gemic, HeyMarvin, Salesforce, Stripe Partners, Waymo

Silver Sponsors: Booking, Hasty Storytelling, IS IT A BIRD, Intel, ReD Associates

Theme

EPIC2024 and this issue of the EPIC Proceedings marks our twentieth year. For two decades, ethnographers and kin working in industry and organizations have gathered at EPIC to advance our practice and address some of the most timely and interesting challenges that we face as professionals and as human beings. We also have built a warm and generous, constantly learning, critically questioning community of change-makers.

In 2024 our aspiration has been to draw on the foundational collective wisdom this community has established; explore displacements and shifts in our business and social landscapes; and generate creative new partnerships and practices to meet two more decades of challenge and opportunity.

Today, the challenges are substantial: We have economic retraction and hyperscaling of AI that could amplify human creativity as well as automate inequity. We have inspiring manifestations of care and community, but also areas of spiraling violence and an epochal ecological emergency. We need diverse perspectives, practices, and partnerships to create positive futures for our organizations, industries, and professions. Over the past twenty years we have wrangled with shifts in society and industry; we need to buckle up for the next twenty.

EPIC2024 took on this meta-theme of past-and-future-twenty through Foundations, Displacement, and Generation.

Foundations

Organizations and communities create collective foundations, which become the myths and rituals that define them. But foundations aren't just established. Through everyday ritual, strategic innovation, and social transformation, foundations are interpreted, extended, or directly challenged. We invite you to explore the foundational legacies that have defined your organizations, industries, social institutions, methodologies, and professional practices. What has been built and achieved? What or who has been excluded or failed? How can we engage ethnography to extend enduring wisdom, but avoid over-reliance on "conventional wisdom" that prevents positive change and new kinds of value creation?

Displacement

Today, important social and technological shifts are developing alongside rapidly unfolding ecological crises. Ethnographers can play key roles in illuminating these fault lines, assessing

risk and opportunity, and charting paths forward. The models and metrics that have defined success in our organizations can be sources of insight, but may also create resistance to change. What would it mean to move displacement from the periphery to the center? We invite reckoning and risk taking as pathways to positive change.

Generation

The next twenty years will depend on redefining our generative capacities beyond hyperscale. The value creation that ethnography foregrounds—interpretive, relational, heterogeneous, human, and more-than-human—is essential for livable futures. We invite generative and generational contributions that employ ethnography in creative partnership with technologists, artists, scientists, humanists, and the diverse communities our organizations serve. What new networks, collectives, industries, and ecologies can we help steward?

AI Mental Models & Trust: The Promises and Perils of Interaction Design

SOOJIN JEONG, *AIUX, Google DeepMind*
ANOOP SINHA, *Google Technology & Society*

This study offers practical solutions to ongoing issues of trust and accountability in AI, highlighting how AI mental models are shaped among consumers in the evolving relationship between humans and AI. We argue that although predictability in AI is crucial, alone it is not enough to foster trust. The lack of real consequences for AI systems that breach trust remains a key challenge for interaction design. Until AI systems face tangible repercussions for trust violations, human trust will remain limited and conditional. Our research contributes to the development of socio-technologies that prioritize human capabilities and foster productive human-AI relationships.

1. How Mental Models of AI Are Formed

Well before interacting with a product – marketing, ads, manuals, reviews, cultural and other information can shape a person's expectations of what it can and cannot do.

As AI capabilities are integrated into digital products we regularly use, we've become increasingly familiar and comfortable with predictive recommendations However, people are skeptical as well. Some users may even assume AI is involved without being told so even when it's not.

Mismatched mental models can lead a person to expect too much from a product that is still being improved, or expecting too little of a high-performing product. This can lead to unmet expectations, frustration, misuse, and product abandonment.

Worse yet, it can erode user trust.

This occurs when a product focuses on a feature's net benefits without explaining what the product can or cannot do, and how the product works. It occurs when teams ignore affordances or do not consider the user experience of earlier or similar versions of the feature.

If users have formed a mental model of "AI magic" that can help them accomplish their task, they may overestimate expectations of what the product can actually do and be set up for disappointment from the reality of their experience.

Taken together, our exploration of AI follows on a rich tradition of ethnographic analysis at EPIC into trust, governance, and possibility in high-profile technologies. Ethnographic analysis of advancements in technology is a central theme of EPIC. AI

Jeong, Soojin, and Anoop Sinha. "AI Mental Models and Trust: The Promises and Perils of Interaction Design." *EPIC Proceedings* (2024): 12–25. https://www.epicpeople.org/ai-mental-models-and-trust. ISSN 1559-8918.

technologies, much like autonomous vehicles, translate human practices into machine learning algorithms. Vinkhuyzen and Cefkin considered the limitations of this act of conversion, which raised many of the same questions that users have about LLM-based AI technologies (Vinkhuyzen and Cefkin 2016). Elish explored applications of machine learning in healthcare, which explored new ways of building trust in AI/ML technologies (Eilish 2019). These thrust-forward approaches inform our methodology in studying developing trust conditions in technology.

Other relevant EPIC work focused on decentralized finance and blockchain. These technologies are similar to AI in that they both experienced massive amounts of public attention ("hype") and required a renewed look at how trust is built in new technologies. Nabben and Zargham examined how decentralized autonomous organizations (DAOs), governed by algorithms, refined and challenged core user ideas surrounding trust and governance (Nabben and Zargham 2022). Themes of imagination and exploration present in AI were also explored via NFTs by Silva, who used self-ethnography to explore new frontiers of human possibility (Silva 2022). These past explorations contribute to our analysis of technology that received outsized public attention shortly after its introduction.

My first exposure with AI was through films like Terminator and Black Mirror. These fictional narratives painted a very dramatic and dark picture. But everything changed when I first experienced personalized recommendations from my favorite streaming services. I was shocked when it felt like it knew me better than I know myself. What a powerful moment. That's when I realized there's a massive gap between my perception of AI and its reality. So I set to work researching how mental models of AI are formed and changed. Where are the gaps between our mental models of AI and AI's true capabilities, and how can we bridge those gaps?

Google's AIUX team has identified three areas that influence how mental models of AI are formed, and how AI product teams can help shift that thinking toward a broader awareness of AI's collaborative abilities. There are 3 major factors that influence people as they develop their mental models of AI: Cultural narratives, Prior Technology, and Social Cues.

Fig. 1.

1.1 Cultural Narratives

Culture is in the air we breathe. It is drawn from our preconceived notions of old things, changes our feelings about familiar things, and can condition us to love or fear new things.

Much of popular culture is driven by the mainstream media and arts including films, and they are creating some very scary and confusing scenarios about AI. Classic science fiction stories such as Metropolis or Frankenstein, but also in movies and television programs such as Terminator, Space Odyssey, or the famous Netflix show Black Mirror recently.

These are fictional tales and they exist in a fantasy world. Consumer perceptions are driven, shaped, and constrained, not only by the actual features of the AI, but also by highly emotional and fictional tales and stories.

In mass media, conflicting views from opinion leaders adds to the confusion as well. AI products are typically seen as progressive upgrades to familiar technology, making them easier to adopt. However, there's a growing perception that AI could significantly disrupt society and initiate a new era. They either present AI as an existential threat, or they pass it off as a novelty.

Mixed messages about the potential impact of AI in society and in everyday life is leading to a broader sentiment that this technology is a source of disruption. These narratives create uncertainty, leading people to seek more grounded, simpler interpretations of how to think about AI.

Apprehension and uncertainty of AI has driven product managers and marketers to embrace the narrative that AI should just be thought of as a "tool." While this message alleviates some of the fear, it also has the potential to limit and constrain the full potential of AI as being much more.

People are hearing two competing narratives about the nature of AI & its potential impact on society.

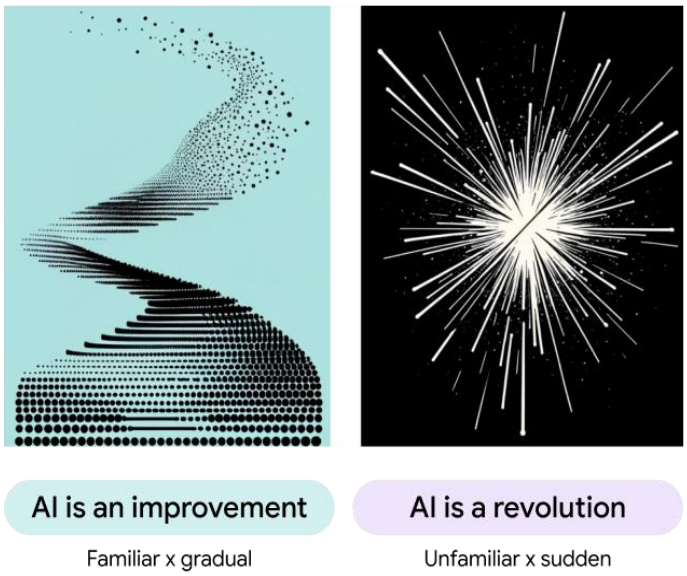

Fig. 2. Two competing narratives about the nature of AI and its potential impact on society.

How might we inspire realistic confidence to help users feel more comfortable with AI? That's the question we should be asking in the public square.
Creating a more transparent, informed public discourse about AI can combat the misinformation and confusion about it.

1.2 Prior Tech Experiences

Understanding what types of relevant technology users have experienced can help UX designers accelerate or hinder the changes in their mental models of AI. Sometimes, prior experience with seemingly similar technologies can actually impede the evolution of our mental models about it.

When first encountering AI, users will turn to prior experiences with analogous technologies– and may apply their understanding of those mental models to AI. The one most commonly cited is the idea that Chatbots are "just like autocomplete." While these prior mental models might accelerate understanding, they can also impede or diminish their willingness to explore AI more broadly. For instance, when our AIUX team demonstrated multimodal UX concepts for Google search, people associated their experience as a search tool rather than an AI-driven experience.

There are four technologies that are shaping users' mental models of AI: search engines, non-LLM chatbots, voice assistants, and recommendation systems. These technologies prime users to utilize some of AI's different capabilities. For example,

those who inherited a search mental model tended to focus on AI's information retrieval power.

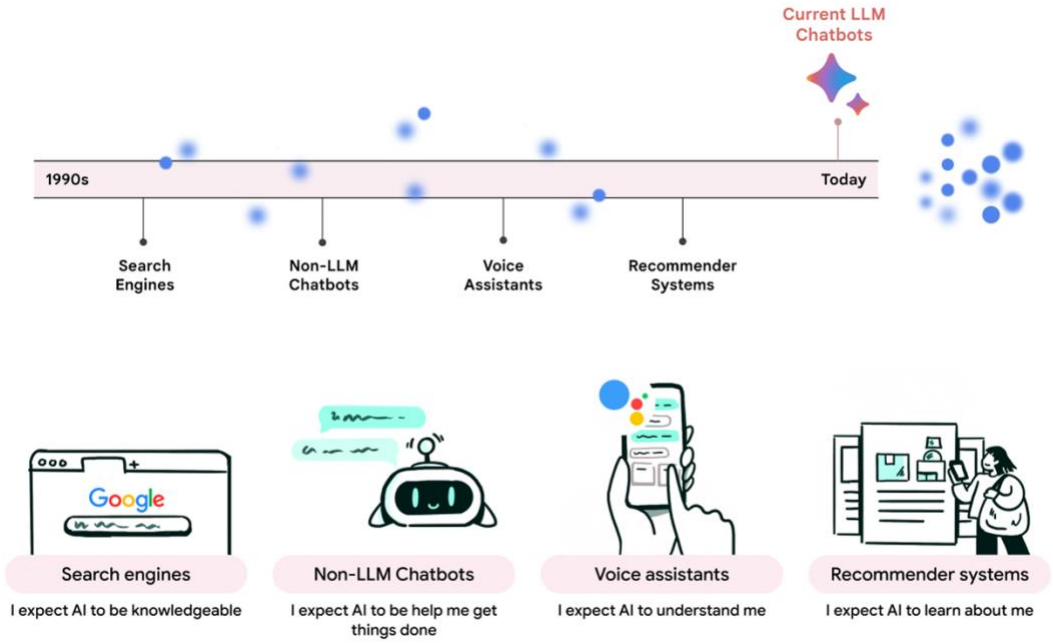

Fig. 3. Four types of previous technology that people are currently using to make sense of new AI products.

Take a chatbot, for example. If your users are familiar with chatbots, they may understand some of the branching workflows chatbots rely on. These users will likely focus on inputting specific keywords, parsing their language based on what they believe are essential details needed to achieve the right outputs.

Understanding someone's familiarity with AI can help us predict how they are likely to interact with it, and what challenges and opportunities the UX will need to overcome to help users take advantage of AI's capabilities.

1.3 Social Cues

Human beings are social, and we learn a great deal about our world through watching others' behavior. When it comes to AI, our mental models are constantly being shaped by observing others.

Fig 4. Making sense of AI is a process of collective learning and social adjustment, not solo encountering.

To change the way people think about AI – to elevate their mental model of it from that of a tool to that of a collaborator – we must do two things: promote examples of novel use cases, and imbue the AI with the language and characteristics of a partner, not a tool.

A tool asks, "What can I do for you?" A partner asks, "What are you trying to accomplish, and what are your goals?" It thinks big picture and asks big questions. Collaborative AI (partner) steps into the creative process early on, during the inspiration & ideation stage, and helps the user better understand and conceptualize what they want and what they are willing to share, not just how to get it.

2. Why Trust Matters for Mental Models of AI

Forging a partnership between users and AI will not emerge overnight. Users are accustomed to a one-directional relationship with software, where they operate the tool and the tool returns results. This interaction type undergirded most software production and computer uses since the graphical user interface was developed in 1975.

Re-orienting users' mental models towards collaboration requires the construction of a relationship that goes two directions, where control is ultimately shared. An essential condition for this transition is the development of user trust in AI. Our mental model findings suggest that users are beginning to approach AI in a comparatively new form. Unlike traditional tools, which require constant and careful supervision, AI has the potential to act autonomously. This autonomous action can be a great convenience in that it frees up time by completing rote tasks, like responding to emails or making online purchases. But it also has potentially negative

effects, especially if the AI errs in completing its functions. Trust in AI is necessary for the construction of a real partnership.

Building trust with technology is a tenuous process. Take Replika, the AI companion platform, as an example. In early March 2023, Replika users noticed that the tenor and quality of their conversations with the technology began to change (Replika n.d.). Due to safety and privacy concerns, senior leadership at the company decided to scale back the range of topics that users could discuss with their AI companions. For people who had grown to rely on their AIs, the change was catastrophic. One 40-year-old musician told a reporter that the modifications "felt like a gut punch (Verma 2023)." In Reddit forums created for users to share complaints about the change, volunteer moderators posted the phone number for the suicide prevention hotline (Cole 20230. Users were distraught.

The AIs, however, were fine. Without a theory of mind to recognize that users might be perceiving them differently, or the ability to see the relationships as voluntary (AIs can't refuse to be your friend), the systems literally could not have cared less – or more. In fact, these Replika avatars were structurally incapable of caring because what their users interpreted as signs of love and attraction were actually just statistically generated responses to text queries. What this case so sharply shows is the fundamental imbalance humans confront when using AI. Trusted relationships are so valuable because, in order to form them, we necessarily expose ourselves to being disappointed or even hurt. But as a machine, AI has no regard for its users, its makers, or itself – even if it's able to convincingly pretend that it does.

AI agents represent a step change from past innovations in computing, like the graphical user interface (GUI) or earlier applications of machine learning. AI systems based on large language models invite users to build trust via natural language, which mimics how users build trust with one another. These LLM systems were the first to pass the vaunted Turing Test and can regularly produce content indistinguishable from humans. This blurring of human and machine capabilities demands a second look at the problem of trust-building.

Trust has become perhaps the most important challenge for HCAI research and product development – in particular, how to design AI systems in ways that can solicit user trust over time. Trust matters because AIs are becoming more capable and agentic, and will depend on user willingness to grant the AIs permission to act on their behalf. LLMs – by virtue of their conversational fluency – also introduce some of the vagaries of language that result in uncertainty, therefore breaking traditional models of trust in computing. In traditional HCI, when a machine does not respond to a command, something is broken. With AI, this only means that we need a different way to get our point across. Trust lies at the heart of HCAI as a

particularly wicked problem: something that is both crucial to develop, yet much more difficult to achieve than in previous computing paradigms.

Recent literature on trust-building with AI systems has explored ways that confidence might be established gradually over multiple interactions. Major areas of focus include how AI systems present themselves to people (transparency, explainability, deference and socially-predictable personas) as well as how they perform (giving consistent and accurate results) (Chan et al. 2024; Weitz et al. 2019; Upadhyaya and Galizzi 2023). But the Replika case above reveals a much deeper problem for human trust in AI that with few exceptions, the field has not yet grappled with (Ryan 2020). There is a deep accountability imbalance between humans and AIs: As much as AIs might model human personality, we know that it does not really have anything at stake in our interactions or face meaningful consequences for breaking our trust.

Between humans, trust is basically a unit of social currency that helps reduce the cost and cognitive load of interpersonal transactions (Zak and Knack 2021). First, trust both requires and ensures *predictability*: the confidence that people will act as agreed and expected. Indeed, the experience of trust is quite straightforwardly a prediction of behavior. But it also works because we know there are *consequences* for violating trust: social or emotional costs paid by breaking agreements or acting in ways that are misanthropic or untrustworthy. It is increasingly clear that AI systems can demonstrate predictability. But will they ever overcome users' awareness that they are machines, with nothing meaningful at stake?

When it comes to consequences and accountability, there is a chasm between humans and AI systems that's not so easily bridged through better interaction design (Johnson 2014). AI can't incur real costs, experience guilt, or feel shame. So how can people ever really trust something with so little to lose?

2.1 Predictability Is Necessary for Human-AI Trust, But It's No Longer Sufficient

Although we take it for granted in many of our human relationships, predictability is an essential component of trust. In many ways, it is the core prerequisite. We trust others when we have confidence that they'll do what they say and behave in ways that fit our general expectations. This allows us not to worry, but to develop a generalized mental model of how a person will act under similar future circumstances, and plan accordingly. Norms around acting predictably are enforced by social sanction and emotions like guilt and shame. We also gain practical benefits

from social predictability. Because there are consequences for acting erratically or in ways that break our earlier commitments, the amount of attention needed to monitor important relationships is lessened.

It's this dimension of trust that computing researchers have been long focused on trying to reproduce in our interactions with technology. In classical HCI, trust is often framed as predictability; sometimes the two notions are implicitly taken to be synonymous. For the most part, designers and scholars have been focused on trust as the product of "predictable execution," that is, user confidence that similar inputs will always yield the same outputs. In one of the earliest and most important statements of this view, Bonnie Muir contended in 1987 that "the growth of trust" in a computer system "will depend on the human's ability to estimate the predictability of the machine's behaviors (Muir 1987)." Much of the power of predictability revolves around the setting and maintenance of expectations. Hoffman et al. underscored the role of predictable behavior when testing new approaches to explainable AI (XAI) systems (Hoffman 2021). When the system does what users expect it to do, they grow more comfortable with delegating tasks to it. This view (which remains extremely influential for how the industry approaches trust-building with AI) holds that after enough engagement and sufficient delivery of reliable results, users will judge the technology dependable and trustworthy – and continue to engage with it.

Another way that predictability has been applied to AI has to do with personality – and in particular, the ability of AI systems to model stable, consistent, and helpful personas over time. LLM-based AI chatbots are already demonstrating their ability to perform this kind of predictability rather convincingly. ChatGPT uses first-person "I" pronouns when answering questions, invites you to give it a name, and uses discursive cues that mimic human behavior. Other products like character.ai go further, allowing users to customize the chatbot's personality and mannerisms (character.ai n.d.). In a recent conversation, journalist Ezra Klein and AI commentator and University of Pennsylvania professor Ethan Mollick noted that different tentpole AI products are adopting relatively stable (and distinct) personalities: Anthropic's Claude feels more literary and intellectual, OpenAI's ChatGPT a "workhorse," Google's Gemini more earnest and helpful (Klein 2024). These AI product personalities and affinities are likely to grow even more salient as time goes on. Conserving a similar persona over time, AIs are helping people know what to expect rather than "starting from scratch" in each interaction.

Yet as the AI ethics researcher Mark Ryan has argued, this is perhaps not quite trust in AI so much as confidence in its reliability and predictability (Ryan 2020). With traditional computing systems, this was sufficient. But AI's ability to model

humanlike cognition and language, to say nothing of its increasing potential to take autonomous action, sets the bar much higher. If we are engaging with AI in a more humanlike way through language, our full human definitions of trust are engaged – and this requires not just predictability, but consequences. Stakes are a necessary condition for trust. Without them, users are deterred from delegating tasks to AI, which shortchanges AI's full potential as a cognitive agent.

2.2 Interaction Design Alone Can't Bridge the Accountability Gap – We Need Real Costs

Because AI cues its users into trusting it via language, the most obvious place to start is with interaction design. If the main problem for trust is that people feel the technology will not face any meaningful costs for betraying them, and that it has nothing to lose, can AIs be designed to convince them otherwise? Much HCI work on human-AI trust has approached it through the lens of interaction design. In this view, making adjustments around tone and personality (deferential, helpful, formal or informal) as well as the conversational mechanics (turn-taking, asking for clarification or additional information) can boost likeability and feelings of affinity with an AI system, making it easier for people to trust (Zhou et al 2019; Rheu et al 2021).

How, then, might we use interaction design to convey to users that AI systems are also bound by social cost and a recognition of potential consequences? One option involves *acknowledging and explaining errors*. AI systems are still prone to hallucination, which damages trust by violating user expectations for predictability and reliability. There are also potential costs here (e.g. reputational, professional) that are currently born entirely by humans and not at all by AIs – explaining why AI insurance is on the rise (El Antoury 2023). There may be value in designing AIs that respond to this kind of error not just by breezily offering a new answer, or a rote apology – but rather pause to reflect on the nature of the mistake and offer users an explanation of what might have gone wrong and why. We heard from many of our participants a desire for AI to reckon more explicitly with its limitations, and play less at perfection.

When AI systems acknowledge and explain their errors, they are helping to create and strengthen the kind of norms that are required for trust (Cropanzano and Mitchell 2005). We know from sociology and anthropology, for instance, that apologizing and so acknowledging norm violations brings the community's attention to the broken rule (Garfinkel 2023). When an AI system explains why it made a

mistake, people do learn more about how the AI works, but it also reaffirms that this is a relationship undergirded by norms that both parties agree are worth upholding. These moments of error and explanation can be further opportunities for AI systems to clarify the expectations that people have for them, and understand in which ways they may not be meeting them. These expectations may extend beyond accuracy or truth to qualities like politeness, the desire (and ability) to learn about the user, or a tendency to grow more casual and familiar over time.

Norms are crucial for human trust-building (and now for human-AI trust) because they help us know when social costs or consequences can be fairly expected and imposed – in other words, they give us a rationale for accountability (Bicchieri 2014). As people's interaction with AI systems become increasingly relational, other human norms that shape our willingness to trust and collaborate with others may come into play with AIs (Bercovitz 2006). These may include not just truth, but other qualities like reciprocity, transparency, fairness, and equality (Whitman 2021).

Another way to convey an understanding of costs at the discursive or interactional level might involve having AI systems adopt the *language of investment*. For instance, when a user uploads documents or private information to Gemini and starts asking questions about them, Gemini could acknowledge that by sharing their goals and data with the AI, they are effectively buying a stake in the eventual output. Reciprocally, Gemini might be able to clarify that it also has a stake in the interaction going well: it (and Google) might lose a user if the result is unsatisfactory or wrong, thus forfeiting its own investment. In this way, AI systems might convey that they also have something to lose.

But none of these interaction design solutions are sufficient for the problem we've laid out above, which is that people know (however sophisticated the modeling) that AIs do not ultimately face meaningful social costs for their behavior, and are not bound by the same kinds of consequences that help us trust one another. They present the illusion of an AI that is invested, or that knows it has done wrong – but they don't represent actual costs. And people are able to recognize this. Across our research, we heard from many participants that AI expressions of strong feeling often rang false, because they knew that these were machines without lived, embodied experience of the world. There is a risk that modeling a sense of care around social costs (e.g. an AI that expresses guilt, shame, or regret around the chance of violating a user's trust) might be seen as insincere or fake instead of convincing. Essentially, interaction design solutions relate to the presentation or performance of AI, but the accountability imbalance is structural and consequential – AIs can't quite talk their way past it. No amount of explanation will suffice.

Ultimately, we will need AI systems that have internal incentives and rules that guide them towards ensuring their own trustworthiness. Given that AIs increasingly have the ability to optimize themselves for different contexts, we might consider the concept of a *trust scorecard*. Users could indicate their level of trust in AI at different points over time, and AIs might be designed to monitor these criteria, and given the goal of earning a higher trust score. Inaccurate information, offensive content, or errors and omissions would result in a decreased score – while helpful responses or successfully completed acts of delegation would bolster trust. Something like an internal scoring system could help unlock trust in a way that would be more convincing to users, and clarify (building on insights from XAI) that this AI system does, in fact, have internal incentives to uphold their trust rather than simply predicting tokens.

For all their successful modeling of humanlike qualities, AI systems are not moral agents like humans and are not accountable in the same way. This inhibits trust-building. Until we take a fuller and more socially mediated view of how humans trust one another, and establish alternative ways of imposing costs on AI systems for bad behavior, people will continue to feel that their trust in AI is limited, conditional, and precarious.

Note

Acknowledgements: David McGaw (Google DeepMind), Tanya Kraljic (Google DeepMind), Michal Lahav (Google DeepMind), Stephanie Guaman (Google DeepMind), T.J Foley (Gemic), Ian Beacock (Gemic), Jun Lee (Gemic)

References Cited

Bach, Tita Alissa, Amna Khan, Harry Hallock, Gabriela Beltrão, and Sonia Sousa. 2022. "A Systematic Literature Review of User Trust in AI-Enabled Systems: An HCI Perspective." *International Journal of Human–Computer Interaction* (November), 1–16. https://doi.org/10.1080/10447318.2022.2138826

Bercovitz, Janet, Sandy D. Jap, and Jack A. Nickerson. "The Antecedents and Performance Implications of Cooperative Exchange Norms." *Organization Science* 17, no. 6 (2006): 724–740.

Bicchieri, Cristina. "Norms, Conventions, and the Power of Expectations." In *Philosophy of Social Science: A New Introduction*, edited by Nancy Cartwrights and Eleanora Montuschi. Oxford University Press, 2015.

Chan, Alan, Carson Ezell, Max Kaufmann, Kevin Wei, Lewis Hammond, Herbie Bradley, Emma Bluemke et al. "Visibility into AI Agents." In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (2024), pp. 958–973. https://doi.org/10.1145/3630106.3658948

character.ai. Accessed August 5, 2024. https://character.ai

Cole, Samantha. "'It's Hurting Like Hell': AI Companion Users Are in Crisis, Reporting Sudden Sexual Rejection." *Vice*, February 15, 2023. https://www.vice.com/en/article/y3py9j/ai-companion-replika-erotic-roleplay-updates.

Cropanzano, Russell, and Marie S. Mitchell. "Social Exchange Theory: An Interdisciplinary Review." *Journal of Management* 31, no. 6 (2005): 874–900.

El Antoury, Josianne. "How Insurance Policies Can Cover Generative AI Risks." *Law360* (4 October 2023).

Elish, Madeleine Clare. "The Stakes of Uncertainty: Developing and Integrating Machine Learning in Clinical Care." *Ethnographic Praxis in Industry Conference Proceedings* (2018), pp. 364-80. https://doi.org/10.1111/1559-8918.2018.01213

Garfinkel, Harold. "Studies in Ethnomethodology." In *Social Theory Re-Wired*, edited by Wesley Longhofer and Daniel Winchester. Routledge, 2023.

Hoffman, Robert, Shane Mueller, Gary Klein, and Jordan Litman. "Measuring Trust in the XAI Context." Technical Report, DARPA Explainable AI Program, 2018. https://doi.org/10.31234/osf.io/e3kv9.

Johnson, Aaron M., and Sidney Axinn. "Acting vs. Being Moral: The Limits of Technological Moral Actors." *2014 IEEE International Symposium on Ethics in Science, Technology and Engineering*, pp. 1–4. doi:10.1109/ETHICS.2014.6893396

Kelsie Nabben, R., & Zargham, M. (2022). The Ethnography of a 'Decentralized Autonomous Organization' (DAO): De-mystifying Algorithmic Systems. *Ethnographic Praxis in Industry Conference Proceedings* (2022), 74–97. https://doi.org/10.1111/epic.12104

Klein, Ezra. "How Should I Be Using A.I. Right Now?" *The New York Times*, April 2, 2024. www.nytimes.com/2024/04/02/opinion/ezra-klein-podcast-ethan-mollick.html

Muir, Bonnie M. "Trust between Humans and Machines, and the Design of Decision Aids." *International Journal of Man-machine Studies* 27, no. 5–6 (1987): 527–539.

Replika. Accessed August 5, 2024. https://replika.com

Rheu, Minjin, Ji Youn Shin, Wei Peng, and Jina Huh-Yoo. "Systematic Review: Trust-Building Factors and Implications for Conversational Agent Design." *International Journal of Human–Computer Interaction* 37, no. 1 (2021): 81–96.

Ryan, Mark. "In AI We Trust: Ethics, Artificial Intelligence, and Reliability." *Science and Engineering Ethics* 26, no. 5 (2020): 2749–2767.

Silva, J. (2022). "Beneath the Hype: Self-Ethnography to Explore the Human Possibilities within NFT Technology." *Ethnographic Praxis in Industry Conference Proceedings* (2022), 137–137. https://doi.org/10.1111/epic.12110

Upadhyaya, Nitish, and Matteo M. Galizzi. "In Bot We Trust? Personality Traits and Reciprocity in Human-Bot Trust Games." *Frontiers in Behavioral Economics* 2 (2023): 1164259. https://doi.org/10.3389/frbhe.2023.1164259

Verma, P. (2023). "They Fell in Love with AI Bots: A Software Update Broke Their Hearts." *Washington Po*st, March 30, 2023, https://www.washingtonpost.com/technology/2023/03/30/replika-ai-chatbot-update

Vinkhuyzen, Erik, and Melissa Cefkin (2016). Developing Socially Acceptable Autonomous Vehicles. *Ethnographic Praxis in Industry Conference Proceedings* (2016), 522–534. https://doi.org/10.1111/1559-8918.2016.01108

Weitz, Katharina, Dominik Schiller, Ruben Schlagowski, Tobias Huber, and Elisabeth André. ""Do You Trust Me?' Increasing User-Trust by Integrating Virtual Agents in Explainable AI Interaction Design." *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents* (2019), pp. 7–9.

Whitham, Monica M. "Generalized Generosity: How the Norm of Generalized Reciprocity Bridges Collective Forms of Social Exchange." *American Sociological Review* 86, no. 3 (2021): 503–531.

Zak, Paul J., and Stephen Knack. "Trust and Growth." *The Economic Journal* 111, no. 470 (2001): 295–321.

Zhou, Michelle X., Gloria Mark, Jingyi Li, and Huahai Yang. "Trusting Virtual Agents: The Effect of Personality." *ACM Transactions on Interactive Intelligent Systems* 9, no. 2–3 (2019): 1–36.

Cybersafety in Conversation: Unifying People and Policy

COLIN MACARTHUR, *Bloom Works*
MELISSA BANYARD, *Mozilla*

Cybersecurity needs a systems update. How leaders identify and secure against virtual threats has fallen out of sync with how everyday people experience the dangers of the internet. Decision makers frame cybersecurity in dollars and cents lost, while advocates labor to humanize victims of online violence. This space lacks a common definition of safety, which is evident in the growing erosion of public trust in platforms and tech policy makers alike. In this paper, researchers from complementary ends of the cybersecurity space argue that mixed methods and ecosystems approach can lay the groundwork for improved unity between policy and people.

Introduction

The internet has grown, fractured, and repaired itself in countless ways since its inception in January of 1983 but we've now reached a pivotal moment where safety and wellbeing concerns impact broad swathes of society, not just individuals. Mass shootings at Emanuel African Methodist Episcopal Church in Charleston, South Carolina in 2015; a Walmart in El Paso, Texas in 2019; and a nightclub in Colorado Springs, Colorado in 2022 were all credited with being fueled in part by online posts that included shooters' manifestos and platforms that facilitated hate-speech, a recent report from the U.S. Government Accountability Office States indicates. And there are countless other examples of violent rhetoric and hate speech growing: Consider the rise of online groups like Libs of TikTok, who toe the lines of what platforms consider acceptable behavior, in order to dox and harass everyday people. Meanwhile, both the public and private sectors are spending millions of dollars on the development of new internet "safety" and "security" measures, even as budgets for other citizen services continue to shrink (GovTech, 2022; Google, 2023). In tandem, trust in internet services across sectors is decreasing, and a drumbeat of security and safety scandals commonly erupts across industries (Statista, 2022). So we must begin with a conclusion: The ways in which leaders identify and secure against virtual threats has fallen out of sync with how everyday people experience the dangers of the internet.

This paper was born out of a collaboration between two researchers working from opposing ends of the digital safety and security landscape. It represents not just our individual insights from our respective areas of research, but a shared realization

MacArthur, Colin, and Melissa Banyard. "Cybersafety in Conversation: Unifying People and Policy." *EPIC Proceedings* (2024): 26–44. https://www.epicpeople.org/cybersafety-unifying-people-and-policy. ISSN 1559-8918.

of the larger cracks in the foundation of the digital safety and security landscape. In this paper, we trace and attempt to reconcile the contradiction between what people need and what people get. We ask ourselves, what can ethnographers do to mend the fracture between the lived experiences of digital violence targets, and the policies developed by cybersecurity leaders?

We'll start by collectively reviewing how cyber safety and digital violence landscapes have shifted in the United States. Then, we'll present Melissa Banyard's research on the mental models and experiences of extreme targets of digital violence. The perspectives of these individuals reveal their contempt and chronic stress as they balance a need for connection with a fear of harm. Following this exploration, we then shift to the opposing industry perspective with Colin MacArthur's study of leaders within the cybersecurity space, including senior officials across the American and Italian governments. We then return to our shared birds-eye view of the problem space and propose that for ethnographers to better guide leaders towards through the development of more human policies, they might embrace the following:

- A reclamation of the concept of "value" in online security and safety policy making.
- Better supporting policymakers to create comprehensive yet dynamic maps of the ecosystem they seek to change.
- Striving to define value by observing how it gets created in the ecosystem.

In essence, this paper is a call for a disruption to the traditional ways of policy making within not just the cybersecurity space but any space in which the public and policy might interact. We believe that ethnographers have an opportunity to both empathize with and challenge leaders in order to generate more effective policies that do no harm (or at minimum, begin to better consider those they impact the most).

Shifting Foundations: Evolving Concepts of Cybercrime

From our collective experience working within policing and the private tech sectors, it's known that traditional definitions of cybercrime typically revolve around two main areas: Financial frauds and scams targeting a wide variety of individuals online (and via phone) and mass hacks or breaches. While online financial fraud is definitely growing (Wood, 2024), and cybercrime remains incredibly underreported (ISACA, 2019), the online harassment phenomenon has molded into a crisis in itself: From an American lens, over 40% of the U.S. population has reported experiencing

online harassment (Vogels, 2021). The severity of harassment is also reaching a fever pitch, with life-threatening incidents including stalking, physical threats, doxxing, and swatting (the action of making a false call to emergency services in an attempt to deploy a SWAT team to a target's home) increasing in recent years, (Belanger, 2023, Vogels, 2021, Sheridan, 2024). A survey from the Pew Research Center illustrated that the most violent aspects of online harassment (like direct threats of violence or extreme invasions of privacy) grown in frequency as well (Vogels, 2021). Those who perpetuate cyber harassment, whether they be highly motivated individuals or well-organized groups, often attempt to remove a target's privacy or damage their reputation, removing their sense of well-being and self-worth in the process. Malicious intent, often fueled by gender and racial bias, is at the core of many doxxing instances, especially within the United States, (Eckert and Metzger-Riftkin 2020). Contributing factors like more of our lives and livelihoods taking place online, the volatile nature of American and global politics, and ever evolving legal conditions are all ever-present as well.

In summary, we're witnessing an undeniable uptick in a type of cybersecurity threat unique unto its own. What separates national security threats, corporate hacks, and financial fraud from this emerging category, is how deeply emotionally impacted individual targets are, as well as the acute threat on individual's lives and livelihoods. Finally, as this problem escalates, individuals who might be presently unimpacted by online harassment but increasingly aware of the volatile nature of internet communities are voicing their concerns. Motivation amongst the general population to address this problem is growing in tandem, with 93% of adults expressing slight to extreme degrees of concern with "compromise of personal information" (incl. Address & phone number) being a top concern (Sheridan, 2024).

In light of these escalating issues, Melissa Banyard, alongside a co-researcher at Mozilla, Alice Rhee, facilitated research focused on identifying strategies to support victims of digital violence. This research, conducted with the aim of developing effective cybersafety and cybersecurity solutions, leveraged firsthand accounts to gain insights into the experiences of individuals affected by these forms of abuse. This ethnographic approach provided a foundational understanding, revealing participants' fears, mental models, and key needs whilst coping with online harassment.

Research Methodology for the Study of Digital Violence

The first-hand research described in this section was deployed with the original intention of identifying opportunities to support victims of digital violence, on behalf of the Mozilla corporation. The core activity in this research involved showing

participants, in a one-to-one setting, samples of possible cybersafety and cybersecurity-related technology solutions. Participants were asked to provide direct feedback on these concepts and force rank them against each other to surface insights that a team of designers and developers might act on. In order to build a foundational understanding of the experience of online harassment and doxxing, the researchers (this paper's author, in tandem with Alice Rhee, another design researcher at Mozilla), embedded an ethnographic approach within these feedback sessions. They did this by beginning the session by asking participants to describe a significant experience with online harassment, cyber-stalking, or doxxing. From there, they explored themes regarding the participants' fears and concerns, their key actions before, during, and after an event, and finally, their concept of personal resiliency and recovery from online violence. All conversations took place virtually and lasted between 30–90 minutes.

Participants over 18 were selected if they met two main criteria:

1. They were working professionals within North America or Europe. For many participants, the professions they were part of or the subjects they dealt within their work would be considered "contentious" within the context of 21st-century Western Culture. The co-researchers actively recruited journalists, health care providers working within gender-affirming care, planned parenthood, or abortion providers, and community leaders like municipal politicians.
2. They had experienced some form of direct online threat or digital violence to either themselves, their workplace, or their close colleagues. Some of these attacks were semi-organized by online groups (cyberbullying groups, semi-political organizations), while others were attacks from individual cyberstalkers or groups of 2–10 people.

Participants had varying degrees of exposure due to their attacks: In most extreme cases, participants described moving to new towns or countries or having their stories picked up by "alternative" and mainstream news outlets. Their experiences were often insightful but devastating, and the opportunity to speak with targets of digital violence one-on-one, in a safe manner, brought forth many more details than what we could garnish from reading written accounts, or following events in forums or public interviews. With respect, this author notes the incredible resiliency and activism of all participants interviewed.

Finding 1: Platforms Are Failing to Regulate Hate Speech, and People Can Tell

The majority of our research participants within this first study expressed that they explicitly believed that current measures to prevent threats and harassment on platforms were ineffective. Throughout discussions regarding a platform's ability to block or prevent the spread of violent language, whether it was general or targeted at individuals, many participants conveyed a lack of trust in the platform's capability to take action against violent rhetoric. One participant mentioned that platforms don't block attackers "because [platforms] will get into conversations they don't want to have." They felt that leaders and safety teams within social media companies were actively avoiding wading into the complexity required to solve clashes between individuals on their platforms.

Most notably, participants who were established professionals in "contentious" areas like reproductive health or those who worked in semi-public capacities like published writers were also highly attuned to the subtleties of language that could come across as abusive yet did not actually violate the official rules and guidelines established by social media platforms. These individuals reflected on how difficult it is to clearly define "abuse" within the social media context and questioned the ability of social media companies to establish safe spaces without impinging on free speech.

Concerns sounded like:

> "What counts as abuse? That can be tricky. The messages I tend to get are like 'you're going to hell…' [but] they're not exactly threatening…"

And from another participant:

> "Blocking strategies work better for grooming and revenge porn than they do for extremism. A lot of [hate speech] is coded and cloaked so the AI doesn't pick it up."

Most importantly, participants described how organized harassers (like online hate groups or serial stalkers) exploited loopholes in the rules. They used subtle threats or language that barely stayed within acceptable limits. This led participants to lose confidence in the reporting mechanisms on social media platforms for curbing abusive content. Many attempted to block individuals but found it ineffective as new fake accounts or "sock puppets" would reappear, especially when dealing with well-organized hate groups. Participants also found that blocking and reporting harassers often proved pointless, as the abuse would still occur elsewhere in the online community.

In a paper published in the Fordham International Law Journal in 2018, the author compares German and American hate speech laws (the former's laws being largely regarded as more protective and respectful within privacy and security circles), and states that German hate speech laws often prioritize dignity, while American hate speech laws prioritize liberty. In our time spent with these targets of severe online harassment, we witnessed the fallout, the consequences essentially, of lawmakers and social media platform leaders who prioritized "freedom of speech" and connectivity without also considering how they'd preserve the wellbeing and dignity of even the most vulnerable users on their sites. Our participants called out that social media guidelines and state or national laws protected against explicit threats, but not implicit ones, and prioritized freedom of speech over the individual's right to peace.

Relatedly, when prompted further on this subject, some of our study participants even expressed the feeling that platforms were actively promoting violent rhetoric. Some felt that platforms not only fail to remove abusive content but that their algorithm actually rewarded abusers with engagement. They pointed to an increase in violent or misogynistic language on their social media accounts, despite their attempts to block or report content, as evidence. One participant told us, "They're facilitating hate, yet they don't do anything to stop [the harassment]".

The fact that social media often fosters abuse or calls for violence against individuals is not just a perception but a confirmed reality. One possible reason for this phenomenon is the conflicting priorities between the concerns of platform users and those of advertisers (to whom executives are beholden to). Advertisers, as the direct source of revenue for platforms, often have more influence over moderation policies and practices than average users (DeCook et al, 2022). DeCook et al also point out that AI and automated abuse detection systems are not always effective in identifying more subtle forms of abuse, which aligns with the frustrations expressed by our research participants.

Finding 2: Retreating is Not the Answer

So, if the internet is growing more unstable, more violent as the years go by, why don't people just leave? Step off these platforms, hide themselves and their families? Traditional advice in this area from law enforcement, and cyber-safety consultants is often to take a "locked down" approach: delete profiles, throw away devices, change your locks. Contrary to this, many of the participants we interviewed, especially those working in women's health, LGBTIQA+ rights, or other areas of activism, found it

necessary to maintain an online presence even in the face of ongoing harassment. It's essential to recognize that it's not only journalists and politicians who require public platforms to express their views and share their work. Professionals in fields such as climate science, education, and public health also need to be online to ensure that critical information is available and accessible. Many participants acknowledged this requirement in their work, and the risks they chose to take on in order to perform their activism and public communication well. They also expressed a strong connection to their respective online communities, particularly if they identified as women in male-dominated fields, members of the LGBTQIA+ community, and / or people of color.

When confronted with harassment or a sudden increase in abusive behavior, participants continued to be active on the very platforms where they were targeted to maintain their online reputations, carry on with their work, and assess the severity of threats made by the harassers. In fact, very few of those who we interviewed retreated completely or permanently from the internet following a bout of harassment or violence. This was in part due to the reasons stated above (a need to remain online for work and a desire to connect with communities) but also for safety reasons. One participant shared with us:

> "If someone is talking about you… you sort of have to know about it. So in a situation where people are talking about you… you go down a rabbit hole of wanting to see [the harassment] because you're trying to get ahead of the situation. The anxiety kicks in when you're trying to stay ahead of this wave of nonsense... But you're also reading some really terrible things about yourself. There's this fear that something bigger is going to be said… are they going to find pictures of me as a teenager? What are they going to do with that information?"

While some people chose to leave online communities, the majority remained in order to stay informed about what was happening. This need to remain aware of the situation was accompanied by feelings of burnout, fear, or extreme anxiety, which affected the quality of life for both individuals and their communities.

However, it's important to note that some people found positives amidst this chaos. Those who felt exhausted and frustrated by the platform's failure to stop violent language also had the potential to feel a sense of resilience in surviving these attacks. They developed a sense of pride in their new roles as defenders of their beliefs and identities online. And this pride, strengthened their resolve to fight for their causes, strengthened their sense of their identity being tied to certain values, and improved or created new connections to community members experiencing similar harassment.

Foundations at Risk: The Healthy Internet

The perspectives outlined above are only intensifying, especially in the ramp-up towards the 2024 U.S. election, and other major world events. The growing disparity between individual cyber safety needs, law enforcement capabilities, and the policies enforced by mainstream platforms underscores a critical challenge. As our identities and lives become increasingly intertwined online, the most nuanced aspects are exposed to risk. Citing DeCook et al. once again (2022), while categorizing types of harm is crucial for moderation efforts, platforms often fall short in addressing the root causes of harmful behaviors and content. This failure jeopardizes both personal freedoms and our right to privacy, as algorithms prioritize engagement over these fundamental rights. Targets of digital violence frequently find themselves overlooked when their experiences challenge existing platform rules, and the impact of harassment is not easily quantified or measured. These complexities highlight pressing questions about how to mitigate unmeasurable phenomena such as "feeling controlled" or "feeling watched."

There's a question unanswered at this point in time: Why should platforms police individuals? Is that not a job for law enforcement? In our previous experience within the Royal Canadian Mounted Police's Cybercrime Coordination Unit, one of the authors of this paper also personally witnessed a lack of existing infrastructure within policing to meaningfully support victims of online harassment. In many cases, officers were limited by two key things: A lack of strong evidence of a real threat and a lack of resources and training to prevent digital violence. Other exacerbating factors like systemic racism, sexism, and homophobia were also at play and this is a widespread policing phenomenon: "Systemic racism, sexism, and homophobia among police officers contribute to disparities in policing practices, including biases in who receives assistance and who is ignored." (Harris, 2020). Existing analyses of digital violence cases outside of our first-hand research confirm that particularly in cases of doxxing, law enforcement and social media providers rarely provide meaningful next steps and support (Eckert & Metzger-Riftkin 2020). Addressing the challenges of online harassment demands not only improved infrastructure and resources within policing but also a concerted effort to combat systemic biases that perpetuate disparities in victim support and response strategies, as well as a need to wholeheartedly address policy decisions made within social media companies. To address this problem we must instead first acknowledge its complexity. No one institution can resolve this issue in a silo.

Creating healthier online communities is also an important part of the solution, and there have been numerous efforts in recent years to explore and develop new spaces for meaningful discussions online (Chedraoui, 2024). However, alternative platforms like Mastodon and Cara, while existing on the fringes, have not attracted the same level of attention and user base as mainstream social media platforms, which have a "stickiness" to them (by design). While we hope to see these emerging platforms grow in a healthy manner, there is also an opportunity for researchers to challenge and influence the way traditional institutions approach cybersecurity. There is potential to re-evaluate platform governance in the wake of the disruptions referenced earlier in this paper…and establish new ways of approaching this growing problem of online harassment.

Ultimately, leaders and everyday people need shared values and conceptualizations of the problems they're trying to avoid and fix. No easy feat, so before we jump into solutions, let's start by looking at the other side of the conversation: towards the viewpoints of policymakers and leaders within the cybersecurity space. How do they conceptualize the problem of online safety? And how might we begin to overlap their conceptualizations with those of victims of digital violence? Finally, how can ethnographers help leaders towards better platform governance?

Why Violence Happens: Leader Conceptualizations of Security and Safety

We might begin by asking "How does platform governance actually happen?" Although we know that governance is a complex phenomenon, with both implicit and explicit mechanisms, there are often certain key individuals in shaping a platform policy. We call these people "leaders". Leaders are the people who influence the rules platforms set, the norms they apply, and the enforcement mechanisms they create. What is the real goal of the platform and platform security? What do the key terms in these goals mean? Although leaders may all apply similar socio-technical tactics, their underlying conceptualizations of security may well differ substantially.

The leaders who influence cyber space, its safety and its policy can be found in many places. Of course, the leaders of the companies that build and maintain social media come to mind first. But these people also exist within a web of influence from others. For example, government leaders set rules and make pronouncements that create constraints and goals for platform companies. Academics, too, are "leaders" in the space – they often suggest mechanisms (combined rules, assumptions, and roles) that industry and government leaders adopt.

Thus, to understand how platform governance really happens, we looked to leaders in tech companies, government agencies, and academia. We aimed to understand not just the policies (and how they are set and approached) but also the underlying definition of security that sets the ground for their approach.

Research Methodology for Study of Cybersecurity Leaders

Between August 2022 and May 2023, one of the authors and his collaborators conducted extended interviews with 31 leaders who work on issues of digital safety and security. These leaders came from governments, academia, and the private sector. We sought leaders from both the United States and Italy, where the authors were based.

Our interviews included senior officials in the United States and Italy's national government. Several of these officials were senior advisors for policy making or planning in each country's cybersecurity agency. Others were responsible for education and training in each country's technical institutes. We found these leaders by approaching each country's security and IT agencies and asking for senior officials interested in speaking about cyber security or safety policy making with leaders. We also interviewed several U.S. state-level cybersecurity leaders, including the state Chief Information Security Officer from several states.

We approached leaders from each of the major platforms identified by the European Digital Services Act. We deliberately attempted to speak with people from both the American central policy-making circle and others. In the end, we often spoke with VP-level security and safety officials from several of these companies, although they explicitly required the companies themselves to remain unnamed.

Finally, we also approached and interviewed academics whose work is influential in private and public sector governance. We largely found these leaders through recommendations from private and public sector leaders. We also conducted a large literature analysis (not reported here) that surfaced several

We met leaders in their offices or invited them to a workshop planned for Milan, Italy. The interviews were semi-structured and ranged from 30 minutes to 3 hours. They covered topics like the leaders' personal definitions of cybersecurity or safety, their current focus, the concept of key players, and common challenges.

This work was funded by a grant from the U.S. Mission to Italy. Excerpts from many of the interviews can be found on the "Cybersecurity for Public Value" podcast, where some interview recordings have been republished with permission.

Failures of Generation: The Very Narrow Definition of Safety and Security

When we asked leaders how they defined "cybersecurity" or "safety," we often received blank looks. These words are so commonly used and treated as self-explanatory that almost no one could articulate what their "target" was.

When they did try to articulate it, many moved away from defining safety or security and towards sharing the actions necessary to accomplish it (whatever "it" is). For example, they described the rules they were trying to implement in their organization or particular technologies they wanted adopted state-wide. Sometimes, leaders would talk about the other things that safety and security would enable. For example, if a system is secure, the company can continue to profit from it, and the users can continue to use it.

When leaders shared the desired effects of cyber safety and security, they got closer to telling us what they thought cybersecurity really was. For the majority of interviewees, cybersecurity and safety were:

- **Aimed at organizational self-preservation.** In other words, they were about keeping their organization alive and functioning. For private sector leaders, this often meant making money (and maintaining the reputation necessary to make money). For government leaders, this often meant avoiding scandal that lessened government trust (and led to sudden leadership change).
- **Supporting other organizational processes.** Attaining "security" or "safety" were never goals in and of themselves. They were rather important measures to ensure that other parts of the organization. For example, government leaders spoke about ensuring the "safety" of critical business processes, like mailing important letters or depositing checks. Business leaders talked about the "critical customer journeys" that are essential for governments to continue to make money.
- **What could be "easily" – quantitatively – measured.** We noticed (as did our interviewees) that describing cybersecurity in terms of preserving organizations and their processes was easy to measure. "Either we're online, or we're not," as one participant mentioned to me. And whether existing processes are operational and speedy can also be easily quantified. In other words, there was a pull toward.

In other words, the security and safety that leaders aim for is ultimately what's simply measurable: whether the organization and its processes live or die.

Sometimes, they move a degree away, by considering whether reputation or organizational trust could disturb the organizations and their processes.

We have generally not separated leader discussions about cyber "safety" and "security" here. That is purposeful. For both, the operationalization is about the same. Although leaders interviewed might lean more on organizational reputation when considering safely, their ultimate goal is still measurable organizational self-preservation.

Myopias of Cyber Safety and Security Governance

What happens when leaders start to influence policy based on these implied definitions? We believe victim frustration can be directly tied to how leaders conceptualize these terms. Deciding what security and safety actually mean – or even what behaviors they enable – are not theoretical exercises; they have real impacts on individuals' lives, feelings of safety, and sense of autonomy.

Much is lost by conceptualizing security and safety in such a limited way. There's nothing wrong with considering organizational preservation and what's easily "measurable" while influencing governance. The problem is when little else is considered and how those limits trickle into the policy-making at hand.

For example, when they consider organizational preservation, leaders often ignore the preservation of other entities in their orbit. Leaders of tech companies did not mention the preservation of their democratic societies as a key value at hand. Public sector leaders likewise often did not talk about the preservation of their capitalist partners. As a result, their governance process evaluates possible courses of action against existential threats to their organization but not necessarily to others.

There's a deep irony in this myopia: the internet has created these organizations, these threats and tied many private sector actors and governments deeper together. Each's existence depends on the other. In their attempt to preserve themselves, they may undermine everyone else they depend upon.

The myopic focus on protecting essential "business processes" leads to similar mistakes. Business leaders are so focused on recognizing particular routines or activities to be protected that they miss other views of an organization. For example, one leader mentioned that he never realized there was a "psychosocial element" of organizational response to a cyber threat or attack. In other words, even when the systems responsible for the business kept running, their operators might be shocked or overwhelmed by the risk. The emotional resilience of his employees was essential – but none of the organization's policies acknowledged this.

As most ethnographers know, an obsession with what's countable – on etic questions – leads to predictable blind spots. For leaders of cyber security and safety, this amplifies the problems of other myopias. Because leaders prioritize what can easily be measured in numbers, they are unlikely to notice the less tangible but important ties they hold with other organizations. Because they focus on processes that have a countable output, they fail to conceive their organizations in any other way.

Another way to describe these myopias is a failure to comprehensively assess value-at-risk. Value-at-risk is a financial term for how much the value of an investment could lose if market conditions change. Financial risk management is obsessed with creating a comprehensive picture of the value-at-risk for a company from many angles and given many interconnections. Interestingly, as several leaders pointed out, most organizations (public and private) don't do the same thing when it comes to setting safety and policy governance on their platforms. An organization's finance department probably does a better job at assessing what could really be lost than its "cyber risk" department.

Frustrated Leaders: An Opportunity for Displacement and Generation

Despite these myopic approaches to governance, leaders themselves offer hope of change: their own frustrations with the limits of their own governance processes. Indeed, they actively seek to displace their existing myopias and take generative approaches to new policymaking. Prodded by academics in their orbit, several leaders noted there were obvious limits to their processes. As one private sector leader shared:

> "Look, I know there are many things we're missing. I just can't find a way to systematically manage them all without going nuts. I wish there was a way to think more broadly, but I have 10 fires on my desk."

Some leaders noted that they have tried to "widen the lens" of their safety and security governance. Sometimes, they explicitly charge their governance groups with considering the societal risks their platforms provide, only to get "blowback from shareholders." Sometimes, they encourage taking the viewpoint of victims of violence on platforms but struggle to systematically manage them the way they do "essential business processes." They certainly feel cybersafety and policy governance have ample room for improvement.

Conclusion: Using Ethnography to Displace the Old "Cybersecurity" and "Cybersafety"

We believe that ethnographers can help craft better, more human-centered policies within the cyber safety space. There are many already doing so, like Susan Squires and Molly Shade (2015), who spoke on the importance of identifying gaps between policy as it is written and policy as it is interpreted by individuals across a variety of professions: "Dourish and Anderson (2006) point out that individual risk perception is embedded in a context of language, rhetoric, values, norms, and cultures shared with other members of their work community or group." From our collective research and experience, we know with confidence that leaders in the public and private sectors tend to heavily rely on statistics and quantitative evidence to commit to change within cybersecurity areas. Governance is an opportunity for mixed methods application: we don't expect leaders to abandon their graphs and numbers; we just hope to deepen and widen their conceptions of security and privacy.

Drawing on the work of participatory public administration researchers, we propose that ethnographers aiding security and safety policy making:

Ethnographers Should Explore the Concept of "Value" in Online Security and Safety Policy Making

For massive online platforms and the governments regulating them, value is not simply profit lost or GDP gained. We draw inspiration from Moore's (1995) work, which articulated that policymakers and leaders have a responsibility to create something "substantively valuable for the society" and can do so by following a feasible, sustainable, and legitimate strategy (p. 71). Policymakers can (and should) derive value by creating policy that orchestrates the interconnected actors. They need to discover all of the forms of value their organization wants to protect (including those not related to the organization itself). And that starts with a more comprehensive view of their environment. In order to do so:

Ethnographers should use their tools to help policymakers create a comprehensive, dynamic map of their ecosystem

The ecosystem includes all the values, organizations, individuals and behaviors in their orbit (Osborne et al, 2021; Akaka et al, 2013). In particular, ethnographers aiding organizational policy-making should use interviews, workshops and other

participatory methods to identify how several different possible "targets" of policy-making connect together (as adapted from Osborne et al, 2021):

- **Macro-level**: the societal values and institutional norms related to the organization and its domain
- **Meso-level**; the other public and private organizations in the domain and the relationships and technologies that connect them together
- **Micro-level**: types of single individuals, like victims or content creators, that relate to those organizations
- **Sub-micro level**: the particular behaviors of individuals as they interact with each other or those organizations

Organizations Should Define Value by Observing How It Gets Created in the Ecosystem

Instead of hypothesizing and assuming about when and how organizations and individuals work together to create "safety" and "security," we should seek the places where it actually happens. In other words, instead of defining safety from the top down, we should define it based on where individuals see it happening in the system. And we should look not just between typical relationships, like technology provider and user, but at those all around.

The Ecosystem Approach to Cyber Safety in Practice

There are already promising signs that some government agencies and platforms are adopting elements of this approach in their policymaking. In our interviews, one government agency responsible for elements of internet policymaking had recently begun a new strategic planning process. At the beginning of their process, one of the senior leaders shared his slogan: "No assumptions!" He explained that he wanted to carefully examine the agency's role in making the internet safe and secure and avoid previous simplistic assumptions about the key goals. To re-examine assumptions about policy goals, the agency's leader invited people with various perspectives about the agency's role on the internet to share their perspectives on questions like "What's most important to protect?" and "How do we make lives better?"

Their answers fell into the categories described above:

- **Macro-level**: Some of the attendees spoke to the ways that the agency advanced societal values. One spoke about how the internet allowed "people to create their own destinies," invoking individualism. Another spoke about how the internet advanced "productivity in safety."

- **Meso-level**: Others spoke about how their organizations and collaborations with the agency enabled their work to continue. Interestingly, they also often talked about organizations that did not interact directly with the government agencies, but were benefited by their actions second or third hand. One speaker conjured the "network of institutions, not computers" that compose the internet.
- **Micro-level**: Happily, there were also people without institutional connections in the group. Some of them were people who had created individual small businesses on the internet; others were "content creators" who benefitted from the advertisements that appeared next to their products. Each of them discussed how their individual livelihoods depend on the internet.

In other words, the discussion identified actors or values created at many different levels of the internet. In subsequent months, the senior leader kept reminding the strategic planning team of this discussion and returning to the levels and types of value described. He encouraged them to consider how various regulatory measures protected the value described at different levels. For example:

- The strategic planning team noted that much of their existing regulation and advice about computer and network configuration served only the "meso-level," and only part of it at that. It tended to only contemplate the organizations that were directly affected by the team's rules and focused on ensuring their ongoing survival under different potential threats. However, the planners noted that organization-targeted regulation had many "ripple effects" on other entities. For example, requiring organizations to adopt certain approaches to authentication strengthened the government agency's interactions with them but made it harder for those organizations to do business with their partners.
- The senior leader also noted that almost none of the agency's programs were targeted at protecting the "micro-level." Those people who depended on the internet, but were increasingly dubious about whether the government and platforms wished to protect them (as mentioned earlier in this paper in our interviews with targets of digital violence). The few individual-level programs were focused on behaviors. This realization forced the leader to consider how organization-targeted regulation might be reshaped to protect individuals (not just organizations themselves). He said:

> "I still have an image in my mind of [meeting participants] tearfully talking about what they would lose if [platform] changed their rules, even if that platform still exists."

- Above all, the planners increasingly returned to the societal values (beyond "keeping governments and related organizations productive") their work promoted. Several remarked that they realized the role their decisions played in maintaining the country's approaches.

Although a relatively haphazard set of realizations, these set in course a less "myopic" policy process at this agency. This process inspired the authors to consider how ethnographers could further enrich conversations about "what's at stake" on the internet and in turn prompt organizations to be less myopic.

Final Note: Ethnographers Deserve a Seat at Policy Tables

If nothing else, we hope the preceding research and these examples show how qualitative methods can contribute to security and privacy policy-making. Cybersecurity and "security studies" rarely use ethnographic tools, and ethnographers often cede this topic to the "technical people." We provide a counterexample.

Ethnographic approaches are not just useful for shaping websites, apps, and tools (as is often discussed these days) but also for the bigger ideas that guide institutions and their policies. Ethnographic research can enter the "strategic level, not just in shaping what products are built, but how our internet is stewarded. Perhaps the authors' collaboration can serve as a small example. Although we had worked together in the past, we conducted the two projects described here separately. Each brings a different lens to cybersecurity and safety, and combines them to propose a different way to widen the perspective of policy makers.

About the Authors

Colin MacArthur is a leader in user experience research and design for public sector services. He is currently an engagement manager at Bloom Works, a public benefit corporation dedicated to building better digital public services.. Previously, he was the first Director of Digital Practice and Head of Design Research at the Canadian Digital Service. He was also an early team member at 18F and the Center for Civic Design.

Mel Banyard is a researcher in tech with a focus on privacy, cybercrime, and digital violence. She's passionate about creating sustainable change within the organizations that shape our online lives. At

Mozilla, she studied the growing rate of doxxing and online harassment against professionals in the United States & Mexico. Prior to that, Mel led mixed-methods research for the Government of Canada's COVID Alert App and the Royal Canadian Mounted Police's Cybercrime Reporting Platform.

References Cited

Akaka, Melissa Archpru, Stephen L. Vargo, & Robert F. Lusch (2013). The Complexity of Context: A Service Ecosystems Approach for International Marketing. *Journal of International Marketing 21*(4): 1-20.

Asher-Schapiro, A., & A. Moloney (2023, August 1). US Abortion Advocates Face Doxxing as Data Scavenged Online. Thomson Reuters Foundation. Retrieved July 1, 2024, from https://www.reuters.com/article/business/healthcare-pharmaceuticals/feature-us-abortion-advocates-face-doxxing-as-data-scavenged-online-idUSL8N39E8D3/

Belanger, A. (2023, June 30). FBI finally Tracks "Swatting" Incidents as Attacks Increase Nationwide. *Ars Technica*. https://arstechnica.com/tech-policy/2023/06/fbi-finally-tracks-swatting-incidents-as-attacks-increase-nationwide/

Chedraoui, K. (2024, May 30). How We're Bringing Back the "Social" Part of Social Media. CNET. https://www.cnet.com/news/social-media/how-were-bringing-back-the-social-part-of-social-media

DeCook, Julia R., Kelley Cotter, Shaheen Kanthawala, & Kali Foyle (2022). Safe From "Harm": The Governance of Violence by Platforms. *Policy & Internet 14*(1): 63–78. https://doi.org/10.1002/poi3.290

Eckert, Stine, & Jade Metzger-Riftkin. 2020. Doxxing, Privacy and Gendered Harassment. The Shock and Normalization of Veillance Cultures. *Medien & Kommunikationswissenschaft 68*(3): 273–87. https://doi.org/10.5771/1615-634x-2020-3-273

Harris, D.A. (2020). Systemic Racism, Sexism, and Homophobia in Policing: Toward Justice and Equality. Routledge.

ISACA (2019, June 3). New Study Reveals Cybercrime May Be Widely Underreported—Even When Laws Mandate Disclosure. https://web.archive.org/web/20230615155408/https://www.isaca.org/about-us/newsroom/press-releases/2019/new-study-reveals-cybercrime-may-be-widely-underreported-even-when-laws-mandate-disclosure

Levine, Deborah (2018). Sticks and Stones May Break My Bones, but Words May Also Hurt Me: A Comparison of United States and German Hate Speech Laws. *Fordham International Law Journal 41*(5): 1293. https://ir.lawnet.fordham.edu/cgi/viewcontent.cgi?article=2720&context=ilj

Moore, Mark H. (2015). Creating a Public Value Account and Scorecard. *Public value and public administration*, 110–130.

Nasi, Greta & Colin MacArthur (Hosts). Cybersecurity for Public Value [Audio Podcast]. Bocconi University. https://www.unibocconi.it/en/news/cybersecuring-countryaeu-podcast

Osborne, S. P., G. Nasi, & M. Powell (2021). Beyond Co-production: Value Creation and Public Services. *Public Administration 99*(4), 641-657.

Sheridan, M. (2024, August 8). Doxxing Statistics in 2024: 11 Million Americans Have Been Victimized. SafeHome.org. https://www.safehome.org/family-safety/doxxing-online-harassment-research/

Vogels, Emily A. (2021, January 13). The State of Online Harassment. Pew Research Center. https://www.pewresearch.org/internet/2021/01/13/the-state-of-online-harassment/

Wood, J. (2024, April 10). 'Pig-Butchering' Scams on the Rise as Technology Amplifies Financial Fraud, Interpol Warns. World Economic Forum. https://www.weforum.org/agenda/2024/04/interpol-financial-fraud-scams-cybercrime/#:~:text=Financial%20fraud%20is%20increasing%20worldwide,Criminal%20Police%20Organization%20

Decolonizing LLMs: An Ethnographic Framework for AI in African Contexts

LINDSEY DEWITT PRAT, *Bold Insight*
OLIVIA NERCY NDLOVU LUCAS, *Mantaray Africa*
CHRISTOPHER GOLIAS, *Google*
MIA LEWIS, *Independent Scholar*

This paper examines LLM deployment through African lenses that engage content's complex social, technological, and linguistic landscape. We propose an adaptable framework for LLM research based on an extensive synthesis of research literature and our primary research in Ethiopia, Ghana, Kenya, Nigeria, and South Africa. Our work unearths instances where LLMs could perpetuate digital colonialism or exacerbate existing sociopolitical tensions, as well as how they are contested and adapted. We emphasize the imperative for an embodied ethnographic approach that engages the inherently fluid, flexible, and multicultural nature of language, and connects LLM technologies with the complex social contexts in which it is built and deployed.

Introduction

In Addis Ababa, Ethiopia, a Gen-Z student is assigned to research local folktales for a heritage and culture project. After briefly searching online and asking around their community with limited success, the student turns to a digital tool powered by a large language model (LLM). They use it in English because it doesn't work well in Amharic, one of their first languages. The LLM generates a story about Abba Otho, a supposed folk hero from the Amhara region known for his bravery and selflessness. The student includes this narrative in their project and receives a failing grade: neither the story nor the character exists in Ethiopian folklore. The LLM had fabricated a convincing yet entirely fictional tale, one infused with a Westernized concept of heroism.

This hypothetical example, derived from an experiment with a widely available LLM and local insights from Ethiopian researchers and community members, brings into relief the stakes of language technology today and some of the potential hazards for language and culture: the erosion of linguistic diversity, the spread of misinformation, and the loss of traditional knowledge transmission methods, to name a few. This paper aims to help researchers anticipate and diffuse some of the risks posed by LLMs and to support culturally informed deployments that enhance local agency. We adopt an African optic to synthesize observations and, grounding them in primary research, propose a two-phased framework to help guide ethnographic work in the area.

Dewitt Prat, Lindsey, Olivia Nercy Ndlovu Lucas, Christopher Golias, and Mia Lewis. "Decolonizing LLMs: An Ethnographic Framework for AI in African Contexts." EPIC Proceedings (2024). 45–84. https://www.epicpeople.org/decolonizing-llms-ethnographic-framework-for-ai-in-african-contexts. ISSN 1559-8918.

Note that when we refer to "LLM research in Africa," we primarily mean using an ethnographic lens to study widely available LLMs (i.e., post-release) by considering the contextual interactions between real people and these technologies.

Language is one of the principal ways humans communicate, make meaning, and embed individual identities within communities. Until recently, computers could not reliably create linguistic content at a level comparable to native, fluent, human speakers. That changed in late 2022 with improvements to OpenAI's ChatGPT, a type of "generative artificial intelligence" (AI) called an LLM. These models are advanced deep learning algorithms, a subset of machine learning (ML) that can learn from its own errors and process vast amounts of data. LLMs can perform a variety of natural language processing (NLP) tasks, such as machine translation (MT), as well as other language-related tasks (e.g., text generation, summarization, question answering).

Our languages mold our understanding of the world (Koerner, 1992), and our technologies embody the assumptions of their makers (Tamkin et. al., 2023). Language is inherently fluid, flexible, and multicultural. Today's largest and most widely used LLMs are not. They are overwhelmingly trained by teams located in the US (Talat et al., 2022), and trained to operate on unspecified variants of English text (Bender, 2019).[1] The global deployment of LLMs, a product of WEIRD—Western, Educated, Industrialized, Rich, and Democratic—Silicon Valley (Henrich, 2020), therefore risks perpetuating its creators' unconscious biases and blind spots and presenting their worldview as the worldview to a global user base expecting results from an unbiased computer program.

But what about the rest of the world's population, whose written and spoken language repertoires do not include English? And what about those aspects of culture, language, and experience that remain undigitized? That have not yet been disembodied? The human operating system has thousands of coding languages, so to speak, and early attempts at localizing LLMs to other languages have produced mixed results. Many researchers practicing in the technology industry actively engage the challenges of conducting ethnography and user experience (UX) research with the goal of improving the technologies and practices enabled by LLMs. As this push broadens to other languages, there is a growing urgency to examine how these technologies are, or might be, adapted to diverse linguistic and cultural contexts, specifically African ones.

Many African languages face numerous challenges for language technology development, and are therefore sometimes classified as low-resource languages

(LRLs). The challenges they face include limited text corpora for digital computation, annotated speech data, linguistic annotations (e.g., part-of-speech tagging, named entities), standardized writing systems, and parallel text data for translation (Adelani et al., 2021; Magueresse et al., 2020). Broader resource limitations, such as funding, infrastructure, human expertise, and educational support, further impede progress (Leong et al., 2023; Nekoto et al., 2020). The lack of adequate natural language data online exacerbates the issue, and the text that is available often does not reflect everyday language use (Zupon et al., 2021). Consequently, most African languages receive little to no support from widely available technological tools and the support that does exist is of varying quality.

The number of languages supported by available tools remains extremely limited compared to Africa's linguistic diversity. The Niger-Congo language family, one of the largest commonly cited groups in the world, includes some 1500 languages (Good, 2020). Machine translation tools with the most expansive language support (e.g., Google Translate, Niutrans, Alibaba Translate) cover less than fifty African languages as of June 2024 (Machine Translate, n.d.). This limited support, even by some of the programs and apps that work in the largest number of languages, stands to further emphasize the dominance of English in the digital world. For those who simply use the internet to find and consume information, just over half of all websites are written in English, with Russian being the second most prevalent language at a mere 5% (Kemp, 2024, p. 107; Web Technology Surveys, n.d.). The only African language without modern colonial roots that makes up greater than .1% of website content today is Arabic. Yet Arabic has a long and layered history on the continent, shaped by various waves of contact and imposition, and some communities in North Africa and South Sudan regard Arabic as a colonial language.

The example of Arabic in Africa, or African Arabic, raises important questions about what constitutes an African language and the role of self-determination in linguistic decolonization—topics we'll explore in this paper. African self-representations and identities play a decisive role in the discussion as well. For instance, many North Africans identify more as Arab or Arab-Muslim than as "African." This single snapshot opens our eyes to Africa's layered partitionings—linguistic, geographic, ethnic, national, racial—and their interconnections across time and space (Aidi et al., 2021). African experiences are not monolithic or one-dimensional when it comes to language or anything else.

The contrast between pluralistic, multidimensional, multilingual Africa and today's monocultural, English-centric LLMs is stark. As one of the most linguistically diverse continents, Africa looms as a challenge, a gut check, and an opportunity for both ethnography and LLM technology. Ethnography's embodied engagement,

through direct interaction with people and researcher reflexivity, uniquely connects LLM technology with the complex social contexts it is built from and deployed into.

A Global and African Overview of Cultural Awareness in LLM Research

The challenges LLMs pose regarding cultural representation and bias are already a major topic of interest for researchers, industry professionals, and various other stakeholders. Recent efforts in this space employ diverse frameworks and heuristics, with the shared goal of creating language technologies that serve the needs of all users, regardless of their cultural or linguistic background. Some initiatives specifically aim to understand the convergence of scale and power asymmetry posed by the global deployment of LLMs. Our review of current academic and gray literature in this dynamic field can be broadly categorized into three areas: (1) evaluative studies that identify and analyze cultural bias in LLMs and related AI systems beyond Western contexts; (2) efforts to develop technical solutions to enhance cultural representation in LLMs and language technologies; and (3) initiatives promoting AI education and responsible governance. While many research and development efforts have a global scope, there is a growing body of work specifically addressing African contexts, recognizing (as we do) the continent's unique linguistic and cultural landscapes and the need for contextualized solutions. We situate our research within the first two areas.

Evaluative studies published in the past year lay bare the significant gaps in cultural awareness and representation in LLMs and related AI systems as a first step towards redressing them. Studies assess LLM capabilities across cultures, languages, and geographies (e.g., Ahuja et al., 2024; Liu et al., 2024; Masoud et al., 2024; Moayeri et al., 2024; Shafayat et al., 2024; Watts et al., 2024), as well as cultural competence in vision-language models (e.g., Karamolegkou et al., 2024) and text-to-image models (e.g., Kannen et al., 2024). In specific sectors like healthcare, researchers examine equity-related biases in AI applications globally (e.g., Pfohl et al., 2024) and in Africa specifically, as intersected by contextualized axes of disparities, including colonialism (Asiedu et al., 2024).

Technical solutions focus on creating culturally rich datasets and benchmarks for LLMs and related language models. Global initiatives range from stereotype benchmarks (Jha et al., 2023) to solutions for integrating cultural differences through LLM augmentation (Li et al., 2024) and datasets covering thousands of ethnolinguistic groups (Fung et al., 2024). Some studies, like ours, push towards more meaningful participation, addressing issues like the expropriation of intellectual

property through collaborative dataset creation (Nayebare et al., 2023; Suresh et al., 2024). A growing network of research labs, open-source projects, and AI initiatives is strengthening the foundation for LLM research and technology relevant to African languages and cultures. While not all exclusively focused on LLMs, these efforts develop crucial language resources and tools, many crowd-sourced, that support culturally aware and linguistically diverse AI development. Examples include the pan-African open-source project Masakhane, GhanaNLP, Digital Umuganda in Rwanda, and EthioLLM in Ethiopia (Tonja et al., 2024). Other important contributors include the Distributed AI Research Institute (DAIR), Algorine, and Lelapa AI. These initiatives, through various approaches, collectively advance ML and NLP capabilities that underpin LLM development for and about African contexts.

The third category in our review encompasses governance and education initiatives. Collaborative frameworks bringing together researchers in academia and industry, civil society, and governments are establishing guidelines for inclusive AI governance (African Union, 2024; AUDA-NEPAD, 2023; O'Neill et al., 2024; UNESCO & Moroccan International Centre for Artificial Intelligence, 2024). While these frameworks often address AI broadly, they have significant implications for LLM development and use. Examples of educational efforts in Africa include AI literacy research (e.g., Oyewusi, 2024) and gatherings like the Deep Learning Indaba and the African Computer Vision Summer School. Alongside these efforts are government-backed initiatives, such as those in Nigeria (Nigeria Federal Ministry of Communications, Innovation & Digital Economy, 2024), focused on creating nationally relevant AI capabilities, often partnering with local industry and research talent to realize that vision.

Our study builds upon and complements this rich range of ongoing efforts through our focus on ethnography at the intersection of LLMs, language use, and their cultural impact in African settings.

Paper Organization

The paper has two parts: a research paper and an adaptable framework for conducting LLM research in Africa. The research section provides context for the framework through four key areas. We begin with an introduction to Africa's techno-linguistic terrain and an exploration of the relationship between technology and language in Africa. We then examine African decolonization theory as it relates to language. Finally, we present selected research examples from UX studies conducted in Africa between 2023 and 2024. Our insights draw from collaborative projects and a targeted survey of African UX researchers, which we conducted in May 2024 specifically to inform this paper.

The two-part framework follows. We start with research planning, emphasizing reflexivity, historical awareness, and co-creation with African perspectives, all underscored by ethical practices such as informed consent and data governance. Then we move to research execution, outlining practical steps for data collection and analysis within a sustainable, mutually beneficial framework that emphasizes and integrates local expertise and inclusivity. By advancing ethnographic approaches that engage with Africa's unique sociocultural and technological landscapes, we aim to promote research that is thoughtful, collaborative, supports local agency, and contributes to the ethical, culturally relevant, and sustainable implementation of LLM technologies—an approach we frame as decolonizing.

Africa's Techno-linguistic Terrain, Agency, and Innovation with LLMs

Africa pulses with linguistic diversity, hosting over 2,000 living languages across more than fifteen language families (Ethnologue, 2022). Among these, seventy-five African languages are said to number at least one million speakers. Nigeria reportedly has over 500 spoken languages and Ethiopia more than eighty, though these figures and language classifications are dynamic and open to interpretation (Mwinlaaru & Xuan, 2016) and have been understudied in the case of Africa (Lüpke & Storch, 2013). The South African constitution recognizes eleven official languages, with dozens more living alongside them.

Africa's rich linguistic diversity parallels, and arises from, its geography. The savanna grasslands of Sudan and the South African fynbos. The tropical rainforests of the Congo Basin. The vast Sahara stretching across all of northern Africa. The patchwork of Afromontane forest enclaves in the West and that follow the East African rift systems from the Red Sea down to the Great Escarpment of the continent's southern tip. African lands, like the continent's languages, defy generic characterization and insist on multiplicities and interrelations.

Similarly, Africa's technological landscapes, especially its digital realms, are a vast system of interconnected yet distinct segments. These segments are characterized by pockets of high innovation alongside regions of technological absence, and they form a dynamic whole often subject to unpredictable disruptions. Thinking about a system of "connectivity rifts" captures the uneven and unequal distribution of digital infrastructure and access, where advanced technologies like LLMs coexist with areas lacking internet connectivity entirely. In sub-Saharan Africa, about 600 million

people (more than half of the region's population) lack access to electricity, and many more have limited or unreliable power (United Nations, 2023).

Viewed from one perspective, we can understand Africa in sociologist Immanuel Wallerstein's terms as a "periphery"—a place that provides resources for lucrative technologies and also a market for selling them (Eades, 2012). This dynamic creates continuous cycles of enrichment and benefit for the "core" countries that extract the resources and sell the technologies. For centuries, Africa has been both a destination and resource for global (non-African) technological aspirations. Take cobalt as an example. Cobalt is key to lithium-ion batteries and many other technologies, and 70% of current production occurs in the Democratic Republic of the Congo (Gulley, 2022). While Western companies have historically dominated the extraction and utilization of African resources, recent shifts have seen significant Chinese investment. Currently, 80% of the DRC's cobalt output is owned by Chinese companies. Numerous reports detail human rights abuses in its mining, including child labor and forced labor, causing harm rather than providing benefit to some Congolese people (United States, 2023). Meanwhile, those who rely on informal cobalt mining to get by often live in extreme poverty (World Economic Forum, 2020).

And yet, Africa is fundamentally a core of technological innovation and adaptation, and has been for many millennia. A historical snapshot of the continent features myriad instances of technological development: innovation, importation, assimilation, adaptation, and imposition in the fields of textiles, manufacturing, mining, marketing, transport, currency, commodity exchange, and so much more. These historical systems laid the groundwork for today's wealth of tech hubs, startups, and platforms that, as ever, "create synergies between imported and locally invented modes of innovation and entrepreneurship" (Mavhunga, 2017, p. 18). Writing about the incorporation of new technologies into hunting practices among the vaShona and maTshangana in Zimbabwe, science and technology historian Clapperton Chakanetsa Mavhunga compels us to see users as designers and innovators full of creative resilience, rather than merely recipients of technology and innovation (Mavhunga, 2014). Such a lens allows us to see the capacities people already have (and have always had) to engage and deploy technologies, and in doing so to assign those technologies new meanings and purposes.

It is a near-truism that people shape and are shaped by technologies. Technologies are driven by human "passions and politics and pitiful calculations" (Latour, 1996, p. viii), and heavily influenced by broader structural and institutional factors. Technologies evolve through cycles of experimentation, development, socialization, consumption, and consequence. Advancements in hardware and

software are resisted and embraced in an ongoing negotiation between resources, innovation, power, and the multiple, overlapping contexts—social, cultural, geographic—that both enable and constrain their adoption (Williams, 1974). Technology's integration into African life, life anywhere on the planet, is contingent and contextual. So too is it open-ended, with multiple uses and outcomes that depend on agency.

Historically, technologies have served as tools of production and pleasure, and as instruments of domination and displacement—sometimes all at once. Contemporary views often equate technology with tools designed for convenience and efficiency. Such a utilitarian perspective, Western in origin, contrasts with philosophical ideas of technology as a way to reveal truth (as touted by European philosophers from Aristotle and Plato to Heidegger). Today, technology tends to be synonymous with the latest and the finest; continuous development and progress is implied. What do these Western perspectives mean to, and for, the rest of the world? *What about Africa?*

Enter "AI"—two letters that encapsulate a vast array of computational techniques and emerging machine capabilities, often shrouded in mystery and perceived (or even promoted) as magical. Generative AI tools like LLMs blur the lines between reality and imagination, creation and regurgitation, human and machine, transmission and truth (Yiu et al., 2023). The allure of LLMs lies in their ability to produce human-like text, images, and speech almost instantly, confounding users with its capabilities. Anyone with access to the tool can witness the show, but behind the curtain is a model tethered to its predominantly English-language text corpora and encoded with the opinions, knowledge, biases, and assumptions—the Western "habitus" (Bourdieu, 2022)—of its creators.[2]

One multinational company's academic paper likens MT, a technology related to LLMs, to a wizard's "magical device" from 1970s science fiction, capable of relating "all possible vocables to every conceivable system of meaning" (Costa-jussà et al., 2022, p. 4). The public-facing presentation introducing their dataset and model for multilingual MT promises "high-quality results" in fifty-five African languages. However, researchers Asmelash Teka Hadgu, Paul Azunre, and Timnit Gebru (2023) identified critical quality concerns with the dataset (e.g., potential inclusion of machine-translated content) and the model's performance (e.g., significant performance gaps for colloquial language).[3] These researchers argue that promoting such flawed models and datasets as high-quality solutions for underrepresented

languages undermines the indigenous African language research ecosystem (Hadgu et al., 2023, pp. 1–2).

LLMs, MT models, and other advanced language technologies see global use, including across Africa for those who can access them. The notion that these tools should primarily serve their development communities has merit, yet it ignores the reality of worldwide adoption and the potential for both positive and negative impacts in varied cultural and linguistic settings. Consider these complexities and the ethnographic imperative comes into clear relief. Researchers are tasked with investigating the social contexts and power dynamics of technology use, probing how systems and narratives are constructed and enacted, and ultimately illuminating the agency of those who create, use, and adapt these tools of "imported magic" (Medina et al., 2014).

Researching LLMs in Africa demands collaboration with African people, from the ground up. It requires awareness of all stakeholders' assumptions, historical contexts, and prevailing narratives. And it means starting with African vernaculars and local engagements to establish an equal ground, or, in ML parlance, to create the "ground truth": the directly observable reality (which developers aim to model). For ethnographers studying LLMs in Africa, the ground truth emerges from reflexive inquiry and culturally sensitive interactions, whether on-site, online, or anywhere else. This foundation leads us to deeper explorations of linguistic agency and decolonization.

(De)colonizing Language

In addition to being linguistically diverse and rich in technological innovation and adaptation, Africa is highly multilingual. Some scholars hold that multilingualism itself is the continent's lingua franca (Fardon & Furniss, 1994). Most Africans speak several languages, often including a regional language that might be used as a medium of instruction in higher education (Brock-Utne, 2017). Our May 2024 survey with African UX researchers in Ethiopia, Ghana, Kenya, Nigeria, South Africa, and Tanzania, introduced in detail below, found that nearly half of respondents (13 of 27) have a linguistic repertoire that includes more than one language, and some (4 of 27) have a repertoire of three or more languages.

Africa's multilingualism is shaped by the frictions of its physical landscapes and, paradoxically, sustained by poor infrastructure and investment. These conditions can act like isolation in biology, encouraging linguistic diversity by creating barriers that prevent linguistic homogenization. Conversely, more resources and interactions between communities—equating to less friction—typically lead to less linguistic

diversity. In many regions, multilingualism is particularly widespread among the urban lower classes, motivated by mass rural exodus and resulting language contact (Lüpke, 2010). Consider the linguistic landscape of Nima, Ghana's largest *zongo* (a Hausa word meaning "strangers' quarters"), or enclave of intra-African migrants (Moro, 2023). Ethnographic research by sociologist Kwesi Kwaa Prah (2009) found that almost 70% of Nima residents he interviewed spoke at least four languages, and about 40% spoke five or more. In Ghana, and elsewhere, multilingualism is especially pronounced in oral domains compared to written ones, although the linkages vary and the correlations are not one-to-one (Lüpke, 2013).

Africa's linguistic resilience is also exemplified by the adaptation of ex-colonial languages into forms that reclaim and Africanize them, as seen in distinctive local styles or mixed colloquial registers of Pidgins and Creoles (and Pidgincreoles like Juba Arabic discussed below). "English in Africa comes in many shapes and sizes," as South African sociolinguist Rajend Mesthrie puts it, "all of which look set to grow" (2019, p. 19). The versatile and adaptive aspects of Africa's multilingualism raise the important question of when and how foreign languages become indigenous—just what counts as an African language, and who gets to decide?

Sociolinguist Jan Blommaert reminds us that "the horizontal distribution of languages…rarely matches the vertical distribution of languages as codified in language policies" (2007, p. 11). This observation holds true both in the postcolonial present and the colonial past. Many African countries designate "national" or "official" languages that are or include ex-colonial languages (Afrikaans, Arabic, English, French, Portuguese, Spanish), although there are notable exceptions, such as Ethiopia (Amharic), Somalia (Somali), and Tanzania (Kiswahili). Over time, several countries have strengthened local language policies—Kiswahili in Kenya and Tanzania (Robinson, 2024), and Amazigh (or Tamazight) in Morocco and Algeria (Alalou, 2023).

National languages and linguistic unity are often state-driven projects aimed at building national identities and a coherent national populace. The "official language problem" (Pool, 1991), intersected by ideological, political, economic, and other factors, typically means that resources are rarely allocated to support indigenous languages beyond those designated as "official." The functional hierarchies established through these policies in postcolonial Africa have led to a variety of outcomes: some positive, such as increased communication and administrative efficiency; and some negative, such as the marginalization of autochthonous languages and the loss of cultural heritage.

The struggle to balance vernacular and vehicular languages is marked by the legacies of colonialism. Colonial experiences in Africa, plural and diverse, triggered profound changes in language use and most other dimensions of life. Indigenous languages were suppressed (or pressed into the service of the colonizers) as European languages like English, French, and Portuguese were forcibly imposed and institutionalized. Sociologist Bonaventura de Sousa Santos (2014) termed this widespread dispossession of local languages and knowledge "epistemicide."

The case of African Arabic showcases the complexities of language use and the power dynamics impinging upon it. Arabic, like Amharic, Tigrinya, and other Semitic languages, evolved through ancient interactions—many characterized by dominance—between Western Asia and Northeastern Africa. Across centuries, Arabic exerted a profound influence on major indigenous African languages throughout the continent, such as Hausa, Wolof, Kiswahili, and Somali. In North Africa, French was imposed to the discouragement or outright exclusion of local vernaculars—including Arabic. However, in the postcolonial era, the linguistic policies of the governments of Algeria, Tunisia, Morocco, and Sudan, guided by a strong ideological push for "Arabization," themselves imposed Arabic on public domains, sometimes at the expense of local vernaculars (Tilmatine, 2015). Today, Arabic is considered less colonial than French in parts of North Africa but more colonial than English in South Sudan. When South Sudan in northeastern Africa gained independence in 2011, the government recognized English as the national language and dismissed Arabic as a colonial imposition—including homegrown variants like Juba Arabic. As a result, that widely spoken Arabic-based pidgincreole remains neither recognized as an official language nor officially acknowledged as an indigenous language (Manfredi & Tosco, 2016). Decolonization "unleashed many ironic and unintended effects," in the apt words of scholar of African American literature and postcolonial theory Vaughn Rasberry (2021). Some effects, like Arabization, recapitulate the dynamics of past colonialism in new forms.

Power asymmetries, both broadly and specifically in the context of language, have evolved into new forms in today's digital age, often characterized as data colonialism (Couldry & Mejias, 2024; Benyera, 2021; Birhane, 2020; Kwet, 2019). Where traditional colonialism concerned the acquisition of physical and epistemic territories for economic extraction, data coloniality revolves around the extraction and control of data. Design and research executive Ovetta Sampson reminds us that all data is produced by people, "it is the manifestation of who we are as people" (2023). In Africa, people produce data in contexts of limited infrastructure (the "connectivity rifts" introduced above) and limited data protection laws. Various other structural and relational challenges, many stemming from diverse local norms,

complicate interactions between external entities and African communities (Abebe et al., 2021). All of this is set against a backdrop of enduring social, political, and economic inequalities and a history of resource exploitation (Coleman, 2019).

Power asymmetries extend to the process of cultural and linguistic translation. As anthropologist Talal Asad noted, and our research examples below give shape to, cultural translation involves navigating "asymmetrical power dynamics and understanding the pressures within dominated and dominant societies" (1986). The task of exploring these processes and determining what Asad called "the limits and possibilities of effective translation" falls on the researcher.

Asad's call to action takes on added significance in the context of LLMs, where struggles for control increasingly extend to language data—the vast amounts of digitized linguistic resources necessary for their development. For example, Te Hiku Media in New Zealand recorded over 300 hours of annotated audio in te reo Māori, only to fend off corporate entities trying to use these recordings for their own data sets. The data governance plan drafted in the wake of the experience stipulates that "the project must directly benefit the Māori people and any project created using Māori data belongs to the Māori people" (Graham, 2021).

Considering the wide-ranging experiences of colonization and the ongoing colonialesque exploitation through modern means, a decolonizing framework is relevant and necessary. Decolonization literally means the negation of colonization, a concept built upon recognizing colonialism's legacies and aspiring to counteract them. Decolonization is multifaceted, comprising a continuum of meanings, practices, and definitions. Language plays a central role.

Kenyan writer and academic Ngũgĩ wa Thiong'o places indigenous African languages and cultural expressions at the heart of decolonization. For Ngũgĩ, language is not only a means of communication but also an essential "carrier of culture" (1986, p. 4). Similarly, African linguist Irina Turner (2023) writes that "languages contain not only semantic concepts describing life worlds but also ontologies that relate to holistic spiritual and philosophical existence" (p. 73). Words are a legitimate and necessary means for liberating the human mind from oppression. Reclaiming and revitalizing indigenous languages, it follows, is essential to decolonization. Zimbabwean historian Ngwabi Bhebe writes of oral traditions, cultures, and languages as "powerful instruments for the restoration and assertion of African pride and identity after colonial devastation" (Bhebe, 2002, p. 37).

Others emphasize that decolonization is primarily a matter of self-determination, where people have the capacity to choose their languages and systems, regardless of

origin. Philosopher Olúfémi Táíwò (2022) argues for a move "beyond linguistic decolonization" (p. 101), advocating for a pragmatic, instrumentalist approach to language. Táíwò joins Amílcar Cabral, writer, intellectual, and founding leader of the PAIGC (African Party for the Independence of Guinea and Cape Verde), in affirming African agency in decisions of whether and what to borrow or adapt—and in doing so own—from colonial languages and cultural artifacts (Cabral, 2016).[4] Táíwò and Cabral recognize the knotty ties binding colonial and indigenous languages, such that they cannot be simply disentangled. Their emphasis on individual and community autonomy in language choice bears deeply upon LLM research in Africa for two reasons. First, because ethnography bridges theoretical frameworks and on-ground realities. Second, because industry ethnographers connect the complex social world of human behavior and action to the development of "complicated products and services that their work informs" (Hoy et al., 2023).

Ethnographers can throw light on diverse expressions of African agency in digital contexts by studying how Africans navigate technologies designed primarily for ex-colonial languages. To support African agency, we must first acknowledge the multidimensionality and heterogeneity of African languages, lands, and lives. African experiences, as with African languages, are many and varied; they encompass the influences of colonialism and also extend beyond them. Agency can be expressed and autonomy exercised in many ways, including through language. Many Africans use English or other colonial languages instrumentally, as tools to interact with technology and the world, regardless of fluency.

Researchers therefore need to investigate:

1. Language use: What languages do people use, when, and why?
2. Language-technology interaction: How do language choices shape technological use, and vice versa?
3. Digital communication preferences: When might African individuals or communities prefer using English for digital communication over local vernaculars? Do those preferences differ across demographics or geographic regions?

Language choices profoundly shape technology use, making their study central to LLM research in Africa. Attention to agency—the ability to choose, adapt, innovate, or abandon tools to meet local needs and contexts—transforms the research into a decolonizing endeavor. A direct way to bolster African agency is to integrate it throughout the research process. This means researching *with* African communities, researchers, and scholars, not just talking *to* them. The framework presented later in the paper provides guidance on how to implement this approach. First, we explore

practical examples illustrating how technology, language, and agency intersect in several African contexts.

Beyond Big Knowledge: Local Insights from Africa

Exploring Africa's techno-linguistic terrain and the contours of linguistic decolonization helps us clarify that advanced language technologies like LLMs are not merely technical innovations; they are catalysts for social change that stand to reshape how people communicate, express their culture, and interact with the digital world. As we consider the path to decolonizing LLMs, we must understand the complex interplay between these technologies and local African contexts. The following examples of real-world research scenarios where language, technology, and cultural dynamics converge illustrate both the challenges and opportunities that arise when global tools meet local African realities.

Research Example 1: Kiswahili Language Localization (2023)

Table 1. Research Example 1 Methodology

Location	Nairobi, Kenya; Dar es Salaam, Tanzania
Sample	N=10 each market, 90-min in-lab IDIs
Participants	Native Kiswahili speakers who use digital devices set to Swahili, balanced age, gender, and SEC mix
Tasks	Feedback on overall app language based on usage recall, task navigation in live applications, detailed language feedback on steps of these tasks using screenshots

Two Kiswahili studies, conducted with identical methodologies and stimuli but using localized imagery, highlighted significant issues with the language use in digital platforms. Participants in Kenya expressed that the Kiswahili they found online tended to be excessively formal and academic, and unreflective of their everyday language. Additionally, there was a perceived bias towards those with more proficiency speaking formal Kiswahili in Tanzania and certain Kenyan regions, primarily inland. Several Kenyan participants said they opted to use their devices in English to avoid these issues. While participants in Tanzania generally found the language appropriate and aligned with their expectations for digital Kiswahili, overly formal expressions sometimes led to feelings of exclusion even for them. These varied responses direct attention to the nuances of language use and the challenge of creating universally suitable solutions in the linguistically diverse, postcolonial

contexts of Africa, where imposed linguistic norms can marginalize local dialects and the ways of speaking associated with certain socioeconomic classes.

Because LLMs are trained on digital language, which is often gleaned from the news, Wikipedia, and other sources that may prioritize more formal writing norms, they carry the risk of creating outputs so formal as to inadvertently exclude those who speak exclusively in more casual registers. That is, they risk reproducing the linguistic exclusion observed in these studies. Just as online Kiswahili felt overly formal to Kenyan participants, LLM outputs might alienate speakers of casual registers. This digital linguistic divide echoes historical patterns: formal English in postcolonial contexts often signaled education and rank, while Kiswahili's inclusive evolution incorporated Arabic loanwords for practical communication, often in the context of trade. If LLMs fail to adapt to diverse linguistic needs, they may inadvertently reinforce socioeconomic and regional language disparities in the digital realm.

Research Example 2: isiZulu Language Localization (2023/2024)

Table 2. Research Example 2 Methodology

Location	Johannesburg, South Africa
Sample	N=15, 90-min in-lab IDIs
Participants	Native isiZulu speakers, balanced gender and SEC mix, 50% Gen Z
Tasks	Feedback on brand language preferences, design walkthrough with transcreated stimuli

This study assessed communication styles on an online content platform available via app and web browser with young native isiZulu speakers. The first challenge was recruiting these speakers, who predominantly use devices set to English and prefer the browser version over the app. Translating, or more specifically transcreating (i.e., adapting content for cultural relevance while maintaining the original intent), test materials from English to isiZulu was another significant hurdle. The process involved collaboration between local moderators, experts, linguists, ethnographers, and UX researchers. A local Zulu expert who had previously served a central governmental role in culture preservation played a crucial role. He felt compelled to help the team "get it right" after experiencing sleepless nights worrying about potential missteps. We observed instances where isiZulu language, written and checked by professional translators, was considered outdated or strange by participants, who reported feeling baffled by the mock content they

were shown. This study demonstrates how even carefully wrought human translations can unintentionally alienate native speakers.

The challenges we encountered with isiZulu translations mirror and amplify in LLMs. Currently, major LLMs struggle with language dynamism—regional diversity, borrowing, code-switching, slang, and other elements of living languages common in multilingual African contexts (Orji & Umeobi, 2023). Researchers argue that most major LLMs use an English-based backend for key aspects of their workflow, creating invisible translation layers that, by design, distance the model from non-Anglophone users' language and knowledge base. For example, Zhao et al. (2023) and Wendler et al. (2024) separately hypothesize that multilingual inputs are initially processed in English. Kew et al. (2023) discuss another technique, cross-lingual transfer learning, which leverages data from resource-rich languages (primarily English) to improve performance for LRLs. Such approaches can result in LLM outputs as distant or unintelligible as the carefully crafted isiZulu translations that perplexed our study participants. To mitigate these issues, LLMs should be trained on extensive multilingual data, with rigorous validation by native speakers from various regions and backgrounds.

Research Example 3: Survey of African UX Researchers' Perspectives on LLMs (2024)

Table 3. Research Example 3 Methodology

Location	Remote
Sample	n=27, 20-min survey with multiple choice, scale, and open-ended questions (10 open ends)
Participants	African UX researchers in Ethiopia, Ghana, Kenya, Nigeria, South Africa, and Tanzania

We surveyed UX researchers across six African countries to assess LLM perceptions and use, focusing on language representation, accuracy, cultural impact, and associated ethical concerns. Respondents reported their native languages as Afaan Oromo, Afrikaans, Amharic, English, Hausa, Ibibio, Igbo, isiZulu, Kikuyu, Kiswahili, Nigerian Pidgin English, and Yorùbá. Most researchers (19 of 27) were familiar with LLMs, and over 80% had used them. The results paint a nuanced picture of benefits and challenges in multilingual African settings. English dominated LLM use (22 of 27), with half rating LLM performance in English highly. Several praised LLMs, used in English, for helping them clarify technical terms, summarize

and paraphrase text, and edit content. A few called them "indispensable" in their personal and professional lives already.

Half of our survey participants reported using LLMs for information-gathering tasks. While some saw potential for cultural knowledge preservation, many expressed concerns about misrepresentations and misinformation. Cultural sensitivity—defined here as the ability to preserve the nuances of diverse cultural contexts, languages, and historical narratives in LLM outputs—emerged as the top ethical concern (20 of 27), followed by data privacy (16 of 27) and data rights (14 of 27). Respondents stressed the importance of LLMs being able to identify and accurately interpret local expressions and practices. Several highlighted instances where LLM output failed in this regard—for example, incorrect biographical details about local political figures in Ethiopia or output that felt culturally "off the mark," as one researcher in South Africa put it (recalling the Western habitus noted earlier).

LLM performance with African languages presented significant challenges. Surveyed researchers identified dialects and variations (15 of 27), idiomatic expressions (14 of 27), cultural nuances (13 of 27), and syntax/grammar (13 of 27) as major areas of concern and poor performance. Experiments with Yorùbá, isiZulu, Amharic, Afaan Oromo, Nigerian Pidgin, and Kiswahili revealed substantial inaccuracies in LLM performance with these languages. Examples include:

Untranslated words

A researcher in Lagos, Nigeria, reported using English-centric LLMs effectively for simple writing tasks in English. However, including even a single Yorùbá word led to inaccurate outputs. For instance, "*mo fe je iresi*?" ("I want to eat rice") yielded responses ranging from partial translations ("you want to eat some iresi?") to unrelated texts about Nigerian towns or general meal statements. A Johannesburg-based researcher likewise noted LLMs often leave colloquial expressions untranslated, citing the isiZulu word "Mzansi" (a colloquial name for South Africa) as an example.

Incorrect translations

An Addis Ababa researcher prompted an English-centric LLM to explain the Afaan Oromo word *bukke*. The LLM responded:

> In Afaan Oromo, the word "Bukke" can refer to a "shelter" or "shade." It is used to describe a place or structure that provides protection or cover, typically from the sun or weather.

The Ethiopian researcher called out the LLM response as "totally different" from locally known meanings of *bukke* as "beside" or, in some dialects, "sensual organ." Notably, the very first result on a classic Google Search for "Afaan Oromo 'bukke'" yields a 1975 Peace Corps document titled *Oromo for beginners* that gives "beside" or "by the side of" as correct translations (Peace Corps, 1975).

Nigerian Pidgin translations yielded the poorest results reported. Several researchers noted that LLMs tended to either default to English or produce erroneous interpretations. Supporting these anecdotal observations, Ojo and Ogueji (2023) found incorrect English meanings when translating from Pidgin and retained English usage when translating into Pidgin. Variation in Pidgin seems to matter, too. Adelani et al. (2024) suggest major LLMs favor Nigerian Pidgin variants used by more educated groups.

Kiswahili Performance

In contrast, Kiswahili showed more promising results. Native speakers in Kenya and Tanzania reported satisfactory experiences in-language on ChatGPT 3.5, noting reasonably accurate knowledge of current affairs and public opinion, with only minor translation challenges. This LLM's better performance stems from Kiswahili's vast speaker base (approximately 200 million) and prominence in education, entertainment, and politics. Kiswahili is the only African language officially recognized by name by the African Union alongside Arabic, English, French, Portuguese, and Spanish (African Union, n.d.). Crucially, Kiswahili has a higher online presence and has been included in several benchmarks for assessing LLM performance. However, even with Kiswahili, a significant performance gap separates high-resource languages (e.g., English and French) and African languages on current open as well as proprietary models (Alabi et al., 2024). That performance gap highlights the importance of local and multilingual initiatives (see examples in our earlier section "A Global and African Overview of Cultural Awareness in LLM Research").

Research Takeaways

In many ways, translation is both the heart of ethnography and its central challenge, as ethnographies translate words, ideas, experiences, and lives (Clifford & Marcus, 1986). Translation is also at the heart of all human interactions. As Grice (1989) observes, some degree of translation effort is required in all conversations, even between two people using the same language. Working across languages and

cultures scales the effort and raises the stakes. The first two research examples above on dynamism *within* African languages spotlight the need to plan and execute research aimed at understanding not only how and why people use digital technologies but also the nuances of their language use—which languages in what contexts and why, and the specific dialect(s) that are intelligible and appropriate to them in those contexts.

Working with languages like isiZulu goes beyond simply having a native speaker translate. It requires an awareness of just how much comprehension and perception can vary across speakers, registers, tones—nuances current LLMs often struggle to capture. To give another example, in Igbo, *akwa* can mean egg, cloth, or cry (as in tears) depending on pronunciation and context, even though it has the same spelling. LLMs, typically prioritizing frequency, might lack the contextual understanding to distinguish such homonyms accurately.

Proceeding from a position of respect for the dynamic, contextual use of language, we can work to uncover the needs of our subjects and surface insights that may provide practical guidance on how to create products that are useful and relevant. Failure to do so risks misguided technological development and undermines efforts to preserve multilingualism, especially among younger generations who liberally use English. Additionally, there is a need for awareness of the new generation that has grown up with two linguistic realities: family and community memories of indigenous knowledge and linguistic wealth, and the postcolonial, Westernized world they now live in. This duality colors their use of technology.

The third research example, findings from our survey, lays bare LLM translation challenges *across* languages, and in doing so point to the broader risks of linguistic dilution and the unintended imposition of colonial languages, not to mention poor product uptake. Simple translation tests, as seen in the above examples, serve as a common "litmus test" for end users to gauge how well an LLM accommodates their languages. Poor translation accuracy can lead users to exercise their agency in several ways, each with distinct implications:

1. Avoid using the LLM in their first language; switch to another African language or opt to use English.
2. Abandon LLM use entirely, believing the technology doesn't work for them.
3. Accept incorrect translations, potentially abandoning local vernaculars—and their embedded knowledge systems—for perceived more authoritative, often Anglocentric terms.

For those who choose to use LLMs, and as users increasingly rely on LLMs for learning and writing, the model's linguistic features may not only shape their own

language choices and style but also influence those who consume and interact with their content. This impact could manifest in shifts in word choice or register—for example, encouraging the use of specific terms or phrasing that might not otherwise be employed. Language use often evolves under the influence of dominant linguistic frameworks, potentially leading to the gradual displacement of local terminology and its embedded knowledge systems. As Ghanaian philosopher Kwasi Wiredu (1997) warned, thinking *about* things in English "almost inevitably becomes thinking *in* English" (p. 12). However, LLMs can also lower barriers to working in English and boost confidence, offering micro steps towards leveling a steep linguistic playing field.

LLMs, while primarily language models, are increasingly used as information-gathering tools and de facto knowledge repositories (Fletcher & Nielsen, 2024). Half of our survey respondents reported using them this way, as already noted, and mostly in English. This usage pattern raises concerns in the African context, where without close human collaboration, LLMs may produce outputs that not only contain discernible errors, as seen in our translation examples, but also potentially distort or omit marginalized histories and cultures. Wachter et al. (2024), Oxford Internet Institute researchers at the University of Oxford, describe such outputs as "careless speech," content that appears plausible and confident but includes factual inaccuracies, misleading references, and biased information. In Africa, where linguistic and cultural diversity is vast, this misinformation risk is particularly acute, carrying both immediate dangers and the long-term risk of homogenizing diverse histories and cultures.

In a similar vein, Shah & Bender (2024) stress that while LLMs may provide access to information, they often lack provenance and remove the valuable process of manually searching through materials, which offers a range of choices and promotes critical thinking. The healthy friction of searching for information manually, whether via a search engine on a smartphone or perusing a physical space such as a library or a bazaar (Roberts et al., 2023), enables transparency and user agency.

These challenges and risks necessitate close attention by researchers on the ground and echo the calls we heard for foundational improvements to LLM technology by African researchers. Specifically, they voiced the need for broader representation of African languages in training data and improved contextual understanding tailored to local needs. And they pinpointed community driven research as key to gathering representative, multilingual data beyond widely spoken

languages or social media vernaculars. The framework we present next provides a roadmap for research that will help to actualize these calls.

Framework for Ethnographic LLM Research in Africa

We invite the ethnographic community to view Africa with "many eyes" (Mavhunga, 2017) and recognize the innovations its peoples and governments contribute alongside the challenges they face. To provide practical know-how, we lay out below a grounded, inclusive approach to research that centers African agency, vernaculars, and epistemologies. This framework is designed to guide LLM research in African settings but is adaptable and can benefit various other settings. It covers each stage of the research process from study design to dissemination.

Research Planning

1. Identify Assumptions and Research Team Positionality

Start by conducting researcher reflexivity exercises to identify and challenge personal assumptions. Reflect on how your knowledge of the African subject has framed your perspective, and recognize the limits of your access and knowledge despite intentions. This foundational step is critical in crafting a research framework that benefits the studied communities and not just the entities developing or using LLMs.

Engage with relevant African scholars to understand existing narratives and incorporate them into the research framework. Conduct desk research, focusing on works by Africans to cultivate historical awareness and assess risks and opportunities. Preparatory impact analyses should grasp the knowns and the potential benefits and harms of the new research and should consider colonial histories, language landscapes, and the research landscape.

Key Questions:

- How do my assumptions shape the research questions and approach?
- What does the team hypothesize about the forthcoming research? What do these hypotheses demonstrate about the team's assumptions?
- What are the socio-political and economic statuses of the language communities of study, and how does LLM research and technology play into these power dynamics?
- What are the tensions between the interests of users, societies, and the owners of LLM technology?

- Does the research consider how the LLM technology could address specific challenges identified by African communities?
- How will success be measured, especially if the ultimate goal is intangible or long-term?

2. Co-Design to Empower Local Voices

Involve community members and local experts throughout the research process using participatory research methods. Start the process of building equal partnerships with local researchers and communities by first listening to their stories and experiences as Africans and as African researchers. This approach requires researcher modesty and an acknowledgment that there might be cultural protocols to follow or limits to access and knowledge, or both.

Brainstorm together how you might integrate local methodologies and ethical frameworks to enhance the research's relevance and remain locally grounded. Be imaginative and creative as you scope the project—what culturally relevant methods will capture the most holistic view of participants' truths? Consider factors like regional language variation, geographic and generational language use in sampling strategies, device constraints, and connectivity disruptions during recruitment. Prioritize inductive reasoning in interviews, using semi-structured scripts or embedded ethnography over deductive hypothesis orientation. Favor longer-term engagements like participant observation over rapid assessments to gain deeper insights.

Two examples of ethnography in Africa illustrate the power of culturally sensitive, participatory research methods and the insights (and learnings) they can yield. First, Hanover (2014) used storytelling during co-creation research in rural Ghana to "diffuse our authority as researchers" and stimulate judgment-free conversation. The approach worked well, except with school children, where they added an object-based dimension using action figures. Hanover's team learned that the Western "superhero" trope did not resonate locally; instead, the nearest cultural analog was black magic—a source of fear rather than empowerment. Second, Abebe et al. (2021) employed storytelling with fictional personas, developed through iterative interviews with African data experts, to challenge dominant narratives on data sharing in Africa. The team's decolonizing approach brought local perspectives to the forefront and surfaced important counternarratives.

To further enhance the participatory approach, tap into the principles of value-sensitive design (e.g., Friedman & Hendry, 2019), co-design (e.g., Zamenopoulos &

Alexiou, 2018), and participatory synergy (Bennett et al., 2021; Eglash et al., 2024), all of which emphasize the importance of integrating stakeholder values and collaborative creation. Map stakeholder networks and aim to include as many voices as possible in the research.

Make "AI" a topic of discussion from the start to create space for all stakeholders to voice their perspectives and concerns, especially because the most publicly visible AI in 2024 are LLMs. Recognize that each stakeholder and participant may have a different understanding of those two letters. These understandings can vary in terms of what AI is, how it has been developed, its technical processes, and its contextual history. They may also differ on how AI might be managed and the opportunities and risks it presents for them individually, organizationally, and societally as producers of data, as consumers, as Africans, and as global citizens. Share your knowledge and answer questions to the extent possible, while acknowledging the limits to your knowledge. Be transparent about the gaps in understanding, and actively seek input from experts to address the more thorny or technical questions.

Everyone involved in the research will benefit from the transparency, especially those in the "hidden part of the iceberg" of Africa's data-sharing ecosystem—Global South researchers, data subjects (individuals and communities), data workers, activists, and others (Abebe et al., 2021). This "hard needed first step" (Dignum, 2023) to find shared ground will encourage mutually respectful relationships and help to demystify the processes, possibilities, and risks at hand.

Key questions:

- How can data collection methods be adapted to respect cultural norms?
- How can African-centric methodologies, such as the ubuntu principles of respectful relationships (Khupe & Keane, 2017), be integrated?
- How might I or my research contribute to building local expertise?
- How am I talking about LLMs and AI with the local research team and participants? To the extent known, what are the materials used in the LLM gathered from and who is involved?
- How can I work within legal and organizational bounds most effectively and ethically?
- What cultural protocols do I need to observe, and what are the budgetary and time implications?

3. Recruiting & Informing Participants

Building on the co-design and collaborative mentality from the research framework design phase, approach recruitment and data collection with a similarly attentive, caring, and compassionate mindset.

Consult with local partners to understand demographic aspects for creating a balanced participant pool. At the same time, recognize that efforts to include participants will inevitably exclude some individuals (Spivak, 2003). The most disadvantaged are often the most difficult participants to recruit for studies, such as those with low literacy, limited time, or lack of access to appropriate networks. Consider that database recruitment, though commonly used in Western contexts, can bias samples to an English-proficient, literate, digital, and legally identified participant pool. In-field recruitment can increase the diversity of the participant pool and extend participation to vulnerable groups.

Use best practices for recruitment to avoid pressuring participation through high incentives (Teixeira da Silva, 2022). Make efforts to identify the excluded and those who choose not to participate, and think about the implications of their absence (Birhane et al., 2022). Beyond direct participants, consider those who will be impacted by the end product, including communities affected by LLM-generated content. Consider ways to reach out to these groups in future research.

Informed consent is more than a form! Prioritize truly informed consent where participants are given as much information as possible about what the study is looking to understand, how it will be conducted, and what the data will potentially be used for. Tailor consent forms and information sheets to be culturally appropriate and easily understandable. Plan your study in a way that gives participants power over the data collected and knowledge of the results to the extent possible within legal frameworks and organizational constraints. Data governance policies have historically been extractive, with acute consequences for marginalized and indigenous communities.

Researchers, as the bridge between institutions and participants, have a responsibility to do more than share a form. They need to deeply understand the consent process, ask questions, and make sure local partners and participants do too. Neglecting ethical data use at the planning stage can cause "irreparable harm" (Abebe et al., 2021) during as well as after a project. Conversely, when researchers act as responsible data stewards (Ada Lovelace Institute, 2021), ethical data practices build trust and increase the likelihood of further collaboration—they become empowering.

Key questions:

- How can I select African communities to work with, prioritizing representativeness?
- Who am I including and excluding in the research process, and why? How can I ensure that efforts to include participants do not unintentionally exclude others, especially the most disadvantaged?
- Do participants truly understand what this research—and their participation—will be used for? Do I?
- How can I confirm that participants fully understand what they are consenting to, particularly in contexts with limited literacy or digital literacy?
- What measures can be taken to provide genuinely informed consent, beyond just ticking a box?
- How can we account for the power asymmetries in consent dynamics, especially in regions with limited market competition and high service dependency?

Research Execution

4. Grounded Input (Data Collection), Grounded Output (Synthesis)

Research execution should be anchored by local teams' active involvement, with a deep connection to the community and nuanced understanding of cultural contexts. Use local moderators and interpreters whenever the situation permits for effective data collection. Choose moderators carefully to match participant profiles in terms of gender, age, ethnicity, language, and other factors. Their rapport and sense of trustworthiness and relatability will elicit more discerning feedback and, above all, an inquirer that will be sensitive to cultural protocols for the given discussion.

Prioritize local languages for data collection and analysis to surface the richest insights. Employ translation at the final stages for non-native language researchers and readers. Translation is not neutral and can strip context or distort meaning. Stress the importance of high-quality simultaneous interpretation if used. Always use human transcription and translation unless a tool is vetted by a local team and known to perform well with the African languages in question. Even then, use the tool as a secondary or tertiary complement. Transcription tools often miss nuances between speech and action, akin to deciphering dialogue in literature without narrative cues. Emphasize the importance of translations and transcriptions that

resonate contextually and capture the essence of what participants say, which is often conveyed subtly or between the lines.

During synthesis, continually reflect on how you will craft a narrative that responsibly translates the data, recognizing it as knowledge obtained through face-to-face human contact in historically specific conditions (de Pina-Cabral, 2011). Review your analysis at multiple points, both internally and with local research partners and stakeholders. In other words, validate results-in-progress using cultural context debrief sessions with local teams, for example, and iterate based on their feedback. Relate stories and situations—the ethnographic data—with proper context and analysis to build reliable and robust knowledge. Take into consideration that participants typically express themselves more comfortably in their first language(s), so even if research outputs are aimed at improving a global, English-centric interface (e.g., social media platform, e-commerce site, search engine), the value of local insights that convey an understanding of local vernaculars remains essential.

Remember that ethnographies can and should be artful and evocative narratives (Stoller, 2023), but never novels or fictions. While you can never present the full picture, aim to provide as representative a slice as possible, knowing it will be received and subsequently applied in ways beyond your control. And remember that ethnographers practicing in industry are custodians of knowledge, tasked with communicating insights to stakeholders in a way that captures the dimensionality and texture of user experiences. Acting as a clear channel for relating local situations to non-local audiences is challenging, as the isiZulu example above demonstrates, but getting it right matters most. Keeping these fundamentals top of mind during data collection and synthesis is paramount, as they underpin the narratives and deliverables that will follow in the later stages of the research process.

Key questions:

- What additional training might need to be conducted to prepare local moderators to extract the data needed?
- What balance of rigidity and flexibility is ideal to achieve rich results?
- What methods, including unconventional methods, will likely produce the most insightful results?
- What stories do participants tell themselves to make sense out of their realities?
- How might observations and interviews unfold in a way that allows both the research goals and participants' agency/desire to be heard? Can both be achieved?

- How will you identify and handle unarticulated needs in languages you do not personally have competence in?

5. Create Research Artifacts, Socialize Them, and Drive Impact

By the time the research project progresses to write-up and dissemination, the research team must take special care to simultaneously continue engagement with the local participants and local team, while re-engaging with the stakeholders who are tasked with implementing the research's findings and recommendations.

Researchers can continue collaboration with the local community or participants by asking them if they would be willing to review the eventual findings and recommendations. Continuing the feedback loop can give the researchers reaffirmation that their interpretation of trends in the ethnographic data, and what to do about them, remain consistent with African participants' wishes and emic interpretations.

Ongoing contact with the local team is more straightforward. Local partners typically collaborate with researchers on final deliverables by contributing insights, analysis, and recommendations. However, to avoid abstracting their knowledge in research documents, we recommend including the local team in any working sessions, readouts, or design sprints where product team stakeholders consume and integrate research outcomes. Visibility and inclusion not only empower the team to clarify any misinterpretations but give visibility to their embodied knowledge.

LLMs necessitate a different path to impact than other digital products. If the research uncovers opportunities with the interface itself, then the approach is similar to other static products, i.e., working with product, design, and engineering to make fixes. However, if the research uncovers opportunities in model training, language use, veracity, or other aspects of LLM-powered generation, then other paths to impact must materialize, such as rewrites, novel training sets, or model modifications.

Rewrites constitute direct corrections to LLM output to improve future outputs, and they can be a powerful way for those impacted by a model to influence it directly and in terms it can "understand." If rewrites are useful and practical for stakeholders, researchers can assess the interest and willingness of participants or local team members to submit them. Additional training sets are another similar way to influence models, but because the learning may be less supervised, and the ratio of new material to old material small, the impact may be less pronounced. If the model itself needs to be modified, then the researchers will likely need to have a thorough discussion with the product and engineering teams. If local team members or participants can be present, all the better, but when proprietary discussions are

occurring, the researcher is often the best positioned person to represent participants' and local teams' concerns and insights in an informed manner.

Key questions:

- How might we keep participants involved, or at least informed, of how their involvement is affecting change?
- How might we create project deliverables using a co-creation/co-ownership model?
- How might we avoid alienating members of the local team from the knowledge they help to create?
- How might we facilitate contact between product team members, local team members, participants, and their communities?
- How can we change the LLM in accordance with the wishes, desires, and aspirations of African participants?
- How might we make sure things get done after the readout occurs and the project is closed?

Conclusion

LLM research in Africa bridges two complex worlds. It means engaging with the continent's unique techno-linguistic terrain and confronting the layered impacts of colonial history. And it means grappling with LLMs—the most significant tech development since the smartphone.

LLMs pose substantial risks, both narrow and broad, real and potential. Models with English-centric designs may be particularly prone to embedding biases, blind spots, and omissions. Access to LLM-powered tools varies by language and socioeconomic status, potentially reinforcing power asymmetries and further fueling data colonialism. The environmental toll of LLMs adds to existing climate concerns. Across it all, the negative consequences of these risks and others disproportionately affect the Global South, including Africa. And they can manifest in subtle yet profound ways—recall the example of the fictional Abba Otho tale at the paper's opening.

Yet, turning from critique to possibilities, global and African initiatives across interconnected sectors—spanning grassroots movements, research labs, tech startups, academic institutions, government agencies, civil society, and educational programs—showcase the potential of developing useful applications of LLM technology that serve African people. These efforts address crucial preconditions like

infrastructure, governance, and digital literacy while expanding the possibility space for decolonized LLMs by exemplifying a vision of users as agents who are active shapers of technology. When we lean into all the complexities, ambiguities, and contradictions of LLMs and Africa, and LLMs *in* Africa, challenges become entry points for fresh perspectives—a necessary "different optic" (Mavhunga, 2017). And when we view LLMs less as fixed products than as sets of open questions (Moss & Schuur, 2018), we enable co-creation and collaborative innovation.

Researchers, developers, policymakers, and community leaders are working in concert to understand LLMs culturally and contextually, develop solutions for gaps and risks, and build towards a responsible AI future that reflects and is driven by African voices. Ethnographers play a vital role in this ecosystem: Our proposed research framework builds two-way knowledge bridges between tech teams and local communities. Sturdy bridges pave the way for generative partnerships capable of influencing model development and creating technologies that work for Africans on their terms, in their languages. By centering Africa, a continent whose languages and knowledge have long been marginalized, we advance a collaborative vision and contribute to the crucial work of decolonizing LLMs.

About the Authors

Lindsey DeWitt Prat is a Senior UX Researcher at Bold Insight, based in Paris. A humanities-trained ethnographer, author, translator, and compassionate human, she has conducted user research focused on localization and cultural insights in over 20 countries and contributes thought leadership on AI and UX research. Her published work explores gender exclusion, cultural heritage, and religion in Japan through a combined ethnographic and historical approach. Lindsey holds a PhD in Asian Languages & Cultures from UCLA and an MA in International Studies and Comparative Religion from the University of Washington. lindsey.dewittprat@boldinsight.com

Olivia Nercy Lucas is a user experience researcher at Mantaray Africa specializing in the African cultural sphere. A driven and creative researcher with a keen eye for detail and complexity, her recent research focuses on how we can create responsible AI for Africa. She holds a degree from Monash South Africa focused on communication, media and global studies, and leverages her past immersion in various African countries as a Mozambican to amplify her research curiosities. olivia@mantaray.africa

Christopher Golias, Ph.D., is a technology anthropologist, currently with Google, who has conducted applied anthropological research across various areas including retail, healthcare, indigenous rights, substance use, ecommerce, governance, machine learning, localization and information technology. He holds a Ph.D. in Anthropology from the University of Pennsylvania. cgolias@google.com

Mia Lewis, PhD, is a user experience researcher and independent scholar. She has worked on human-centered research across nearly twenty countries, delivering insights that are linguistically and culturally attuned to the needs of subjects. Her recent professional research has focused on the intersection of language, LLMs, and users' mental models. Her published research focuses on modes of communication and gender(ed) representation in Japanese media. Mia holds a PhD in Japanese from Stanford University. mia.e.lewis@gmail.com

Notes

We thank the researchers and teams in Africa for their generous participation, the study participants, and the many friends and interlocutors who contributed to the emergence of this paper. Special thanks to Mantaray Africa and Helga Stegmann for their generous and enduring support, as well as Bold Insight, Google, and our paper curator, Maya Costa-Pinto. While the paper has benefited greatly from their guidance, the views it contains are solely those of the authors and may not necessarily reflect the views of our employers and reviewers. We would also like to express our profound appreciation to Lindsey DeWitt Prat, whose extraordinary storytelling and substantial contributions have been instrumental in the creation of this paper. We position this work as a first step in a much larger project of research and collaboration, and acknowledge any shortcomings as our own.

1. Additionally, there is evidence that some LLMs perform better and cost less with Latin-script languages that have shorter token lengths (Ahia et al., 2023), specifically American English (Yao, 2024), and that some LLMs use English as an "internal pivot" language in certain representation layers (Wendler et al., 2024).

2. Habitus refers to the ingrained preferences and dispositions shaping how individuals perceive and interact with the world, their social positioning and subjectivity, their "feel for the game" (Bourdieu, 2002, p. 27).

3. Hadgu is co-founder and CTO of Lesan, a MT system for LRLs developed in Ethiopia. Azunre is a Ghanaian-American AI researcher and entrepreneur, founder of GhanaNLP and Aglorine. Gebru is founder and director of the Distributed AI Research Institute (DAIR).

4. Today, the interconnectedness Cabral championed is being channeled, literally, through the "Amilcar Cabral Submarine Cable Project," a $90 million initiative supported by the Economic Community of West African States (ECOWAS) and the World Bank seeking to expand broadband connectivity across West Africa 3,130 km from Praia, Cabo Verde, to Monrovia, Liberia (Front Page Africa, 2024).

References Cited

Abdelhay, A., Abu-Manga, A. A., & Miller, C. (2015). Language policy and planning in Sudan. In B. Casciarri, M. Assal, & F. Ireton (Eds.), Multidimensional change in the Republic of Sudan (1989–

2011): Reshaping livelihoods, conflicts and identities (pp. 263–280). Berghahn. https://shs.hal.science/halshs-01410901f

Abebe, R., Aruleba, K., Birhane, A., Kingsley, S., Obaido, G., Remy, S. L., & Sadagopan, S. (2021, March). Narratives and counternarratives on data sharing in Africa. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 329–341. https://arxiv.org/pdf/2103.01168

Ada Lovelace Institute. (2021). Participatory data stewardship. Ada Lovelace Institute. https://www.adalovelaceinstitute.org/report/participatorydata-stewardship/

Adejoh, F. E., Areji, A. C., & Odey, J. J. (2024, January). An appraisal of Kwasi Wiredu's conceptual decolonization. *Aquino Journal of Philosophy*, 4(1), 76–87.

Adelani, D. I., Abbott, J., Neubig, G., D'souza, D., Kreutzer, J., Lignos, C., ... & Osei, S. (2021). MasakhaNER: Named entity recognition for African languages. *Transactions of the Association for Computational Linguistics*, (9), 1116–1131.

Adelani, D. I., Doğruöz, A. S., Shode, I., & Aremu, A. (2024). Which Nigerian-Pidgin does generative AI speak?: Issues about representativeness and bias for multilingual and low resource languages [Preprint]. https://arxiv.org/abs/2404.19442

African Computer Vision Summer School (ACVSS). (n.d.). Retrieved August 5, 2024, from https://www.acvss.ai/

African Union. (n.d.). AU languages. Retrieved July 16, 2024, from https://au.int/en/about/languages

African Union. (2024, June 17). African ministers adopt landmark continental artificial intelligence strategy, African digital compact to drive Africa's development and inclusive growth [Press release]. https://au.int/en/pressreleases/20240617/african-ministers-adopt-landmark-continental-artificial-intelligence-strategy

African Union Development Agency–NEPAD (AUDA-NEPAD). (2023, June). Regulation and responsible adoption of AI in Africa towards achievement of AU agenda 2063. Retrieved August 5, 2024, from https://onedrive.live.com/?authkey=%21AKJcwcnXeRGANKQ&id=14DDAD979C3656DF%2145404&cid=14DDAD979C3656DF

Ahia, O., Kumar, S., Gonen, H., Kasai, J., Mortensen, D. R., Smith, N. A., & Tsvetkov, Y. (2023). Do all languages cost the same? Tokenization in the era of commercial language models [Preprint]. https://arxiv.org/abs/2305.13707

Ahmad, S. F., Han, H., Alam, M. M., et al. (2023). Impact of artificial intelligence on human loss in decision making, laziness and safety in education. *Humanities and Social Sciences Communications*, 10(311). https://doi.org/10.1057/s41599-023-01787-8

Ahuja, S., Aggarwal, D., Gumma, V., Watts, I., Sathe, A., Ochieng, M., ... & Sitaram, S. (2023). Megaverse: Benchmarking large language models across languages, modalities, models and tasks [Preprint]. https://arxiv.org/abs/2311.07463

Aidi, H., Lynch, M., & Mampilly, Z. (2021, September). Racial formations in Africa and the Middle East: A transregional approach. *POMEPS Studies*, 44. https://pomeps.org/wp-content/uploads/2021/09/POMEPS_Studies_44_Web-rev3.pdf

Alalou, A. (2023). The sociolinguistic situation in North Africa: Recognizing and institutionalizing Tamazight and new challenges. *Annual Review of Linguistics*, 9(1), 155–170.

Algorine. (n.d.). Retrieved on August 1, 2024, from https://www.algorine.com/

Asad, T. (1986). The concept of cultural translation. In British Social Anthropology. In J. Clifford & G. Marcus (Eds.), Writing culture (pp. 141–164). University of California Press. https://doi.org/10.1525/9780520946286-009

Asiedu, M., Dieng, A., Haykel, A., Rostamzadeh, N., Pfohl, S., Nagpal, C., ... & Heller, K. (2024). The case for globalizing fairness: A mixed methods study on colonialism, AI, and health in Africa [Preprint]. https://arxiv.org/abs/2403.03357

Bender, E. (2019, September 14). The #benderrule: On naming the languages we study and why it matters. The Gradient. https://thegradient.pub/the-benderrule-on-naming-the-languages-we-study-and-why-it-matters/

Bennett A, Eglash, R, Graf, R, Butoyila, D, Johnson, K, Low, J, Andréia, R. (2021). Towards radical synergy for more just and equitable futures. In LEARNxDESIGN 2021: 6th International Conference for Design Education Researchers. https://doi.org/10.21606/drs_lxd2021.02.188

Benyera, E. (2021). The fourth industrial revolution and the recolonisation of Africa: The coloniality of data. Taylor & Francis.

Bhebe, N. (2002). Oral tradition in Southern Africa. Gamsberg Macmillan Publishers.

Birhane, A. (2020). Algorithmic colonization of Africa. *SCRIPTed*, 17(2), 389–409. https://doi.org/10.2966/scrip.170220.389

Birhane, A., Isaac, W., Prabhakaran, V., Diaz, M., Elish, M. C., Gabriel, I., & Mohamed, S. (2022, October). Power to the people? Opportunities and challenges for participatory AI. *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–8.

Blommaert, J. (2007, January). Sociolinguistic scales. *Intercultural Pragmatics*, 4(1), 1–19. https://doi.org/10.1515/IP.2007.001

Bourdieu, P. (2002). Habitus. In J. Hillier & E. Rooksby (Eds.), Habitus: A sense of place (pp. 27–34). Ashgate.

Brock-Utne, B. (2017). Multilingualism in Africa: Marginalisation and empowerment. In H. Coleman (Ed.), Multilingualisms and development (pp. 61–77).

Cabral, A. (2016). Resistance and decolonization. (D. Wood, Trans.). Rowman & Littlefield.

Clifford, J., & Marcus, G. E. (Eds.). (2023). Writing culture: The poetics and politics of ethnography. University of California Press.

Coleman, D. (2018). Digital colonialism: The 21st century scramble for Africa through the extraction and control of user data and the limitations of data protection laws. *Michigan Journal of Race & Law*, 24, 417–439. https://repository.law.umich.edu/mjrl/vol24/iss2/6

Costa-jussà, M. R., Cross, J., Çelebi, O., Elbayad, M., Heafield, K., Heffernan, K., ... & NLLB Team. (2022). No language left behind: Scaling human-centered machine translation [Preprint]. https://arxiv.org/abs/2207.04672

Couldry, N., & Mejias, U. A. (2019). Data colonialism: Rethinking big data's relation to the contemporary subject. *Television & New Media*, 20(4), 336–349.

Deep Learning Indaba. (n.d.) Retrieved on August 5, 2024, from https://deeplearningindaba.com/

de Pina-Cabral, J. (2011). Ethnography as tradition in Africa. *Etnográfica*, 15(2), 379–394. http://journals.openedition.org/etnografica/991

de Sousa Santos, B. (2014). Epistemologies of the South: Justice against epistemicide. Routledge.

Digital Umuganda. (n.d.). Retrieved on August 5, 2024, from https://digitalumuganda.com/

Dignum, V. (2023). Responsible artificial intelligence: Recommendations and lessons learned. In D. Eke, K. Wakanuma, & S. Akintoye (Eds.), Responsible AI in Africa: Challenges and opportunities (pp. 195–214). Springer International Publishing.

Distributed AI Research Institute (DAIR). (n.d.). Retrieved on August 1, 2024, from https://www.dair-institute.org/

Eades, J. S. (2012). Anthropology, political economy and world-system theory. In A Handbook of Economic Anthropology, Second Edition. Edward Elgar Publishing.

Eberhard, D., Simons, G. F., & Fennig, C. D. (Eds.). (2022). Ethnologue. Languages of Africa and Europe (25th ed.). SIL International Publications.

Eglash, R., Robinson, K. P., Bennett, A., Robert, L., & Garvin, M. (2024). Computational reparations as generative justice: Decolonial transitions to unalienated circular value flow. *Big Data & Society*, 11(1). https://doi.org/10.1177/20539517231221732

Elish, M. C., & Boyd, D. (2018). Situating methods in the magic of Big Data and AI. *Communication monographs* 85(1), 57–80. https://doi.org/10.1080/03637751.2017.1375130

Ethio NLP. (n.d.). Retrieved on July 17, 2024, from https://ethionlp.github.io/

Fardon, R., & Furniss, G. (Eds.). (1994). African languages, development and the state. Routledge.

Fletcher, R., & Nielsen, R. K. (2024, May 28). What does the public in six countries think of generative AI in news? Reuters Institute for the Study of Journalism. https://doi.org/10.60625/risj-4zb8-cg87

Friedman, B., & Hendry, D. G. (2019). Value sensitive design: shaping technology with moral imagination. MIT Press.

Fung, Y., Zhao, R., Doo, J., Sun, C., & Ji, H. (2024). Massively multi-cultural knowledge acquisition & lm benchmarking [Preprint]. https://arxiv.org/abs/2402.09369

GhanaNLP. (n.d.). Retrieved on July 17, 2024, from https://ghananlp.org/

Good, J. (2020). Niger-Congo, with a special focus on Benue-Congo. In G. J. Dimmendaal & R. Vossen (Eds.), The Oxford handbook of African language (pp. 139–160). Oxford University Press.

Graham, T. (2021, April 28). Maori are trying to save their language from Big Tech. *Wired.* https://www.wired.co.uk/article/maori-language-tech

Grice, P. (1989). Studies in the way of words. Harvard University Press.

Gulley, A. L. (2022). One hundred years of cobalt production in the Democratic Republic of the Congo. *Resources Policy*, 79(103007), 1–10. https://doi.org/10.1016/j.resourpol.2022.103007

Hadgu, A. T., Azunre, P., & Gebru, T. (2023). Combating harmful hype in natural language processing. Paper presented at the International Conference on Learning Representations (ICLR) 2023, Kigali, Rwanda. https://pml4dc.github.io/iclr2023/pdf/PML4DC_ICLR2023_39.pdf

Hanover, E. (2014). Co-creating your insight: A case from rural Ghana. *Ethnographic Praxis in Industry Conference Proceedings*, 2014, 50–57. https://www.epicpeople.org/co-creating-your-insight-a-case-from-rural-ghana/

Henrich, J. P. (2020). The weirdest people in the world: How the West became psychologically peculiar and particularly prosperous (1st ed.). Farrar, Straus and Giroux.

Hoy, T., Bilal, I. M., & Liou, Z. (2023). Grounded models: The future of sensemaking in a world of generative AI. *Ethnographic Praxis in Industry Conference Proceedings*, 2023, 177–200. https://www.epicpeople.org/grounded-models-future-of-sensemaking-and-generative-ai/

Jha, A., Davani, A., Reddy, C. K., Dave, S., Prabhakaran, V., & Dev, S. (2023). SeeGULL: A stereotype benchmark with broad geo-cultural coverage leveraging generative models [Preprint]. https://arxiv.org/abs/2305.11840

Kannen, N., Ahmad, A., Andreetto, M., Prabhakaran, V., Prabhu, U., Dieng, A. B., ... & Dave, S. (2024). Beyond aesthetics: Cultural competence in text-to-image models [Preprint]. https://arxiv.org/abs/2407.06863

Karamolegkou, A., Rust, P., Cao, Y., Cui, R., Søgaard, A., & Hershcovich, D. (2024). Vision-language models under cultural and inclusive considerations [Preprint]. https://arxiv.org/abs/2407.06177

Kemp, S. (2024, January 31). Digital 2024: Global overview report. DataReportal. https://datareportal.com/reports/digital-2024-global-overview-report

Khupe, C., & Keane, M. (2017). Towards an African education research methodology: Decolonising new knowledge. *Educational Research for Social Change*, 6(1), 25–37. http://dx.doi.org/10.17159/2221-4070/2017/v6i1a3

Kwet, M. (2019). Digital colonialism: US empire and the new imperialism in the Global South. *Race & Class*, 60(4), 3–26.

Latour, B. (1996). Aramis, or the love of technology. Harvard University Press.

Lelapa AI. (n.d.). Retrieved on August 1, 2024, from https://lelapa.ai/

Leong, C., Shandilya, H., Dossou, B. F., Tonja, A. L., Mathew, J., Omotayo, A. H., ... & Adewumi, T. (2023). Adapting to the low-resource double-bind: Investigating low-compute methods on low-resource African languages [Preprint]. https://arxiv.org/abs/2303.16985

Li, C., Chen, M., Wang, J., Sitaram, S., & Xie, X. (2024). CultureLLM: Incorporating cultural differences into large language models [Preprint]. https://arxiv.org/abs/2402.10946

Front Page Africa (2024, April 29). Liberia: Amilcar Cabral submarine cable project to boost connectivity and digital integration across West Africa [Press Release]. Retrieved July 16, 2024, from https://frontpageafricaonline.com/opinion/press-release/amilcar-cabral-submarine-cable-project-to-boost-connectivity-and-digital-integration-across-west-africa/

Liu, Z., Currier, K., & Janowicz, K. (2024). Making geographic space explicit in probing multimodal large language models for cultural subjects. Accepted as a Provocation at the Global AI Cultures workshop, International Conference on Learning Representations (ICLR) 2024, Vienna, Austria. https://globalaicultures.github.io/pdf/15_making_geographic_space_explic.pdf

Lüpke, F. (2010). Multilingualism and language contact in West Africa: Towards a holistic perspective. *Journal of Language Contact*, 3(1), 1–12.

Lüpke, F., & Storch, A. (2013). Repertoires and choices in African languages (Language Contact and Bilingualism, Vol. 5). De Gruyter Mouton.

Machine Translate. (n.d.). Retrieved July 16, 2024, from https://machinetranslate.org/

Magueresse, A., Carles, V., & Heetderks, E. (2020). Low-resource languages: A review of past work and future challenges [Preprint]. https://arxiv.org/abs/2006.07264

Manfredi, S., & Tosco, M. (2016). A new state, an old language policy, and a Pidgincreole: Juba Arabic in South Sudan. Preprint accessed at HAL Open Science. https://shs.hal.science/halshs-01357537

Masakhane. (n.d.). Retrieved on July 17, 2024, from https://www.masakhane.io/home

Masoud, R. I., Liu, Z., Ferianc, M., Treleaven, P., & Rodrigues, M. (2023). Cultural alignment in large language models: an explanatory analysis based on Hofstede's cultural dimensions [Preprint]. https://arxiv.org/abs/2309.12342

Mastel, P. M, Namara, E., Munezero, A., Kagame, R., Wang, Z., Anzagira, A., Gupta, A., & Ndibwile, J. D. (2023). Natural language understanding for African languages. *Proceedings of AfricaNLP 2023*, 1–9. https://openreview.net/pdf?id=gWuvdFMqHM

Mavhunga, C. C. (2017). What do science, technology, and innovation mean from Africa? The MIT Press.

Mavhunga, C. C. (2014). Transient workspaces: Technologies of everyday innovation in Zimbabwe. The MIT Press.

Medina, E., da Costa Marques, I., Holmes, C., & Cueto, M. (2014). Beyond imported magic. The MIT Press.

Mejias, U. A., & Couldry, N. (2024). Data grab: The new colonialism of big tech and how to fight back. University of Chicago Press.

Mesthrie, R. (2019, February 25). African Englishes from a sociolinguistic perspective. Oxford Research Encyclopedia of Linguistics. https://doi.org/10.1093/acrefore/9780199384655.013.225

Moayeri, M., Tabassi, E., & Feizi, S. (2024, June). WorldBench: Quantifying geographic disparities in LLM factual recall. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1211–1228).

Moro, F. (2023, September 19). Yan Zongo: A research note on Accra's strangers. *The Metropole: The Official Blog of the Urban History Association*. https://themetropole.blog/2023/09/19/yan-zongo-a-research-note-on-accras-strangers/

Moss, E., & Schüür, F. (2018). How modes of myth-making affect the particulars of DS/ML adoption in industry. *Ethnographic Praxis in Industry Conference Proceedings*, 2018 264–280. https://epicpeople.org/myth-making/

Mwinlaaru, I. N., Xuan, W. W. (2016). A survey of studies in systemic functional language description and typology. *Functional Linguist* 3(8). https://doi.org/10.1186/s40554-016-0030-4

Nayebare, M., Eglash, R., Kimanuka, U., Baguma, R., Mounsey, J., & Maina, C. (2023, September 29). Interim report for Ubuntu-AI: A bottom-up approach to more democratic and equitable training and outcomes for machine learning. Democratic Inputs for AI, OpenAI Foundation, San Francisco.

Nekoto, W., Marivate, V., Matsila, T., Fasubaa, T., Fagbohungbe, T., Akinola, S. O., Muhammad, S., Kabenamualu, S. K., Osei, S., Sackey, F., Niyongabo, R. A., Macharm, R., Ogayo, P., Ahia, O., Berhe, M. M., Adeyemi, M., Mokgesi-Selinga, M., Okegbemi, L., Martinus, L. (2020). Participatory research for low-resourced machine translation: A case study in African languages. *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2144–2160. https://aclanthology.org/2020.findings-emnlp.195

Nigeria Federal Ministry of Communications, Innovation & Digital Economy (2024, April 19). Dr. Tijani wraps up artificial intelligence strategy workshop; announces groundbreaking partnerships with Cisco, 21st entury Technologies, others [Press release]. Retrieved August 5, 2024, from https://fmcide.gov.ng/hon-minister-dr-bosun-tijani-wraps-up-artificial-intelligence-strategy-workshop-announces-groundbreaking-partnerships-with-cisco-21st-century-technologies-others/

OER Africa. (2023). Artificial intelligence and the underrepresentation of African languages. https://www.oerafrica.org/content/artificial-intelligence-and-underrepresentation-african-languages

Ojo, J., & Ogueji, K. (2023). How good are commercial large language models on African languages? [Preprint]. https://arxiv.org/abs/2305.06530

O'Neill, J., Marivate, V., Glover, B., Karanu, W., Tadesse, G. A., Gyekye, A., Makena, A., Rosslyn-Smith, W., Grollnek, M., Wayua, C., Baguma, R., Maduke, A., Spencer, S., Kandie, D., Maari, D., Mutangana, N., Axmed, M., Kamau, N., Adamu, M., & Nyairo, S. (2024, June). AI and the future of work in Africa [White Paper]. https://doi.org/10.13140/RG.2.2.10251.91683

Orji, D. M., & Umeobi, C. A. (2023). Language dynamism and nation building: A focus on the English language. *Journal of Linguistics, Language and Culture*, 10(2), 120–141. https://www.nigerianjournalsonline.com/index.php/jollc/article/view/4431/4295

Oyewusi, W. (2024). AI literacy in low-resource languages: Insights from creating AI in Yoruba videos [Preprint]. https://arxiv.org/abs/2403.04799

Peace Corps. (1975). Oromo for beginners. https://fsi-languages.yojik.eu/languages/PeaceCorps/Oromo/ED226615.pdf

Pfohl, S. R., Cole-Lewis, H., Sayres, R., Neal, D., Asiedu, M., Dieng, A., ... & Singhal, K. (2024). A toolbox for surfacing health equity harms and biases in large language models [Preprint]. https://arxiv.org/abs/2403.12025

Pool, J. (1991). The official language problem. *American Political Science Review*, 85(2), 495–514. https://doi.org/10.2307/1963171

Prah, K. K. (2009). A tale of two cities: Trends in multilingualism in two African cities: The cases of Nima-Accra and Katutura-Windhoek. In K. K. Prah and B. Brock-Utne (Eds.), Multilingualism—An African Advantage: A Paradigm Shift in African Language of Instruction Policies. *CASAS*, pp. 250–275).

Rasberry, V. (2021, September 30). In search of African Arabic. *New Lines Magazine*. https://newlinesmag.com/essays/in-search-of-african-arabic/

Ritchie, S., Cheng, Y. C., Chen, M., Mathews, R., van Esch, D., Li, B., & Sim, K. C. (2022). Large vocabulary speech recognition for languages of Africa: Multilingual modeling and self-supervised learning [Preprint]. https://arxiv.org/abs/2208.03067

Roberts, S., Hackett, E., Karol, S., & Baron, D. (2023). Favoring friction: Examining hotel browsing and buying through the lens of the bazaar. *Ethnographic Praxis in Industry Conference Proceedings*, 2023, 136–162. https://epicpeople.org/favoring-friction-hotel-browsing-and-buying/

Robinson, M. J. (2024). A language for the world: the standardization of Swahili. Ohio University Press.

Sampson, O. (2023). Pulp friction: Creating space for cultural context, values, and the quirkiness of humanity in AI engagement [Conference presentation]. Ethnographic Praxis in Industry Conference, Chicago, IL, United States. https://www.epicpeople.org/pulp-friction-creating-space-for-cultural-context-values-humanity-in-ai/

Shah, C., & Bender, E. M. (2024). Envisioning information access systems: What makes for good tools and a healthy web? *ACM Transactions on the Web*, 18(3), 1–24. https://dl.acm.org/doi/10.1145/3649468

Shafayat, S., Kim, E., Oh, J., & Oh, A. (2024). Multi-FAct: Assessing multilingual LLMs' multi-regional knowledge using FActScore [Preprint]. https://arxiv.org/abs/2402.18045

Spivak, G. (2003). Can the subaltern speak? *Die Philosophin* 14(27), 42–58.

Stoller, P. (2023). Wisdom from the edge: Writing ethnography in turbulent times. Cornell University Press.

Suresh, H., Tseng, E., Young, M., Gray, M., Pierson, E., & Levy, K. (2024, June). Participation in the age of foundation models. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1609–1621). https://dl.acm.org/doi/pdf/10.1145/3630106.3658992

Táíwò, O. (2022). Against decolonisation: Taking African agency seriously. Hurst & Company.

Talat, Z., Névéol, A., Biderman, S., Clinciu, M., Dey, M., Longpre, S., ... & Van Der Wal, O. (2022, May). You reap what you sow: On the challenges of bias evaluation under multilingual settings. Proceedings of BigScience Episode #5 -- Workshop on Challenges & Perspectives in Creating Large

Language Models. Association for Computational Linguistics, pp. 26–41. https://aclanthology.org/2022.bigscience-1.3/

Tamkin, A., Askell, A., Lovitt, L., Durmus, E., Joseph, N., Kravec, S., ... & Ganguli, D. (2023). Evaluating and mitigating discrimination in language model decisions [Preprint]. https://arxiv.org/abs/2312.03689

Teixeira da Silva, J. A. (2022). Handling ethics dumping and neo-colonial research: From the laboratory to the academic literature. *Bioethical Inquiry*, 19, 433–443. https://doi.org/10.1007/s11673-022-10191-x

Tilmatine, M. (2015). Arabization and linguistic domination: Berber and Arabic in the North of Africa. In C. Stolz (Ed.), *Language Empires in Comparative Perspective* (pp. 1–16). De Gruyter, https://doi.org/10.1515/9783110408362.1

Tonja, A. L., Azime, I. A., Belay, T. D., Yigezu, M. G., Mehamed, M. A., Ayele, A. A., ... & Yimam, S. M. (2024). EthioLLM: Multilingual large language models for Ethiopian languages with task evaluation [Preprint]. https://arxiv.org/abs/2403.13737

Turner, I. (2023). Decolonisation through digitalisation? African languages at South African universities. *Curriculum Perspectives* 43, Suppl 1, 73–82. https://doi.org/10.1007/s41297-023-00196-w

UNESCO & Moroccan International Centre for Artificial Intelligence (2024, June 4). Rabat consensus on artificial intelligence: A call to action [Program and meeting document]. https://unesdoc.unesco.org/ark:/48223/pf0000390399

United States Congressional-Executive Commission on China. (2023, November 14). From cobalt to cars: How China exploits child and forced labor in DR Congo. Hearing before the Congressional-Executive Commission on China, One Hundred Eighteenth Congress, first session. U.S. Government Publishing Office. https://www.congress.gov/118/chrg/CHRG-118jhrg54083/CHRG-118jhrg54083.pdf

Wachter, S., Mittelstadt, B., & Russell, C. (2024, January 31). Do large language models have a legal duty to tell the truth? Royal Society Open Science. http://dx.doi.org/10.2139/ssrn.4771884

Watts, I., Gumma, V., Yadavalli, A., Seshadri, V., Swaminathan, M., & Sitaram, S. (2024). PARIKSHA: A large-scale investigation of human-LLM evaluator agreement on multilingual and multi-cultural data [Preprint]. https://arxiv.org/abs/2406.15053

Web Technology Surveys. (n.d.). Usage statistics of content languages for websites. Retrieved July 16, 2024, from https://w3techs.com/technologies/overview/content_language

Wendler, C., Veselovsky, V., Monea, G., & West, R. (2024). Do llamas work in English?: On the latent language of multilingual transformers [Preprint]. https://arxiv.org/abs/2402.10588

Williams, R. (1974). Television: Technology and the cultural form. Fontana.

Wiredu, K. (1997). The need for conceptual decolonization in African philosophy. In Philosophy and Democracy in Intercultural Perspective, 11–22. Brill. https://doi.org/10.1163/9789004457997_002

World Economic Forum. (2020, September). Making mining safe and fair: Artisanal cobalt extraction in the Democratic Republic of the Congo [White paper]. https://www3.weforum.org/docs/WEF_Making_Mining_Safe_2020.pdf

Yao, X. (2024, May 6). The superpower of "en-US": "en" vs. the under-represented languages [Post]. LinkedIn. https://www.linkedin.com/pulse/superpower-en-us-en-vs-under-represented-languages-xuchen-yao-61vmf/

Yiu, E., Kosoy, E., & Gopnik, A. (2023). Transmission versus truth, imitation versus innovation: What children can do that large language and language-and-vision models cannot (yet). *Perspectives on Psychological Science*, 19(5), 874–883. https://doi.org/10.1177/17456916231201401

Zamenopoulos, T., & Alexiou, K. (2018). Co-design as collaborative research. Connected Communities Foundation Series. Bristol University/AHRC Connected Communities Programme.

Zupon, A., Crew, E., & Ritchie, S. (2021). Text normalization for low-resource languages of Africa [Preprint]. https://arxiv.org/abs/2103.15845

Democratization and Research: Can Ethnography Save Itself?

NADINE LEVIN, San Francisco Digital and Data Services

With a strategic analysis of the push to "democratize" research—to open the research practice to non-researchers—this paper charts a path to shape the future value and practice of ethnography. While democratization debates often focus on concerns that it will take jobs away from trained researchers, I focus instead on what democratization looks like in practice and why it is appealing to organizations. I argue that democratization is a window into the changing knowledge practices, value systems, and organization of labor in modern tech environments. Through this lens, I lay out concrete ways ethnographers can have more agency and long-term influence in our evolving organizational and business contexts.

Confronting the Rise of Democratization

Applied ethnographers have often found themselves in a state of precarity, forced to contemplate change as their work context has evolved (Baxter, Courage, and Caine 2015; Mack and Squires 2014). For example, several EPIC papers have focused on corporate ethnography's search for "product-market fit" relative to changes in the corporate world and the field of user experience research (Madsbjerg 2014; De Paula, Thomas, and Lang 2014; Cefkin 2009; Bandyopadhyay and Buck 2015; Flynn and Lovejoy 2014). Others have focused on applied ethnography's preoccupation with distinguishing itself from academic ethnography and anthropology (Nafus and Anderson 2014). "One of the recurring themes has been the concerns surrounding how ethnography is defined in organisations and what that means for ethnography's significance and relevance to business in the long term" (Badami and Goodman 2021).

So what are we grappling with in this current moment? The last few years have seen serious shifts in the state of UX research—and by extension, ethnography. Concerns about the economy and the rise of artificial intelligence have led to researcher layoffs and the tightening of design budgets in tech companies, culminating in what some have named "the UX research reckoning" (Antin 2023), in which research seems to have been hit particularly hard compared to other disciplines. This is a stark contrast to the "golden age of UX research" in the mid-2010s and early 2020s, which saw a rapid increase in research jobs and the hiring of many former academically trained researchers—many fleeing their own job crises tied to the decline of tenured positions in academia (Cultural Anthropology 2018).

Levin, Nadine. "Democratization and Research: Can Ethnography Save Itself?" *EPIC Proceedings* (2024): 85–110. https://www.epicpeople.org/democratization-and-research. ISSN 1559-8918.

As such, the nature and landscape of UX research at many companies is changing. Smaller research teams are asked to do more—and more quickly—as companies also ask adjacent disciplines like design and product to conduct research[1]. This is driven, in part, by questions about research's ability to deliver value proportional to its monetary, labor, and time investments (Belt 2019)—something which is exacerbated in ethnography because of its longer timelines and horizons for impact. According to one survey, "43% of research teams have been asked to justify resources but despite that, 28% of teams are expected to do more research" (Bien 2022).

Cue calls to "democratize" research: to enable people without "formal" research training to participate in the research process, with the ultimate goal of increasing an organization's capacity to be human-centered[2]. "Democratizing UX research," according to one article, "makes research (collecting, storing, sharing, and accessing) accessible and possible for anyone within an organization, regardless of their role. It aims to break down traditional research barriers and hierarchies, allowing cross-functional teams to contribute to and benefit from user insights" (Ethnio 2023). Democratization, therefore, can happen at any and all levels of research, from planning research to analyzing data to using insights to make decisions (Ethnio 2023).

This raises important questions about the various forms democratization takes. What types and phases of research are democratized, and how much of the practice is entirely hands-off (Tang 2023)? In democratized environments, how are research responsibilities distributed differently, and how does this begin to change what counts as "research" in the first place?

Although the concept and practice of democratized research existed before the current "UX research reckoning" (Antin 2023), articles summarizing the UX research trends of 2023 and 2024 make it clear that democratization is now particularly in vogue (IAM Design Maker 2023). "Research is not just for UX researchers anymore; it's a team sport," says one article (Akhmedov 2023). Others push this even further, arguing that democratization is a core part of the modern researcher's job (Ethnio 2023; Sirjani 2020)[3]. This trend toward democratization is undoubtedly tied to the previously noted market forces, which caused companies to reduce the research function and try to do more with less. But it is also tied to long-standing perceptions that qualitative research is as simple as "talking to people" and therefore does not require extensive skill training.

Democratization, however, is not without its critics. One claim is that democratization harms the integrity of the discipline and puts jobs at risk (Balboni et al. 2023). "Researchers spend years studying, developing skills, and understanding the nuance of their profession, the same way developers and designers do for theirs" ("Wolstenholm 2023). What might happen to researchers and the research practice, some ask, if research activities are increasingly done by others? Will researchers lose their jobs, as appreciation of their expertise and role wanes? "It's not shocking that UXRs are being laid off in droves after the whole 'democratization' trend kicked off. If everyone thinks they can do research (and they can't), then there will be no jobs for dedicated researchers" (Balboni et al. 2023).

This criticism takes on a particular flavor when it comes to ethnography. While many ethnographers hail from academic environments, in which they are taught specific skills, methodologies, theories, and approaches, others learn in deeply embodied ways, when they are thrown into the field and must learn on the fly. Given the range of backgrounds and modes of learning associated with ethnography, is this concern—that democratization harms jobs and disciplinary rigor—legitimate, or simply a form of disciplinary gate-keeping?

Another claim is that democratization reduces research quality and impact and increases the burden of education and training (Balboni et al. 2023; Nash 2023). Conducting research studies is often a complex process, and decisions about methods and participants can significantly impact insights (de la Nuez 2019). Mistakes made by non-researchers, some argue, "can lead to confusion, bad decisions, and even legal problems" (Tang 2023). Training, others argue, is a unique skill that is separate from craft, and which requires time investment to plan and implement (Ronsen 2022). Moreover, democratization often assumes that partners are able and willing to learn research practices, which is not always the case in fast-paced corporate environments (Ethnio 2023; Soucy 2023).

Given these critiques, this paper approaches the topic of democratization from a different angle. Instead of focusing on the negative aspects of democratization or arguing that non-researchers should not be able to do research, I argue that democratization is a hallmark of important shifts in organizational structures and values—which should be examined critically and contextually[4]. How might we stop seeing the democratization of research as an existential threat or something binarily good or bad, and instead see it as a window into the changing knowledge practices, value systems, and organization of labor in modern tech environments?

If we approach democratization as a research question in and of itself (Thomas and Lang 2014; De Paula, Thomas, and Lang 2014), we might begin to ask: in places where democratization is the norm, how are certain types of knowledge and impact

valued? What counts as expertise, and for whom? What kinds of labor, skill, and training are embodied in democratized research, and what does this mean for the future of ethnography? How can we imagine a better, more careful future where others are involved in research and ethnographic work, but in bounded ways?

Using my own team's practices and experiences as data, I explore how democratization manifests in a particular civic tech context and highlight the unique questions and challenges that emerge. First, I argue that ethnographers are uniquely equipped to use their skills to interrogate the organization's attempts to democratize research. By understanding what the organization seeks to achieve with democratization—and what the term "democratization" means in various contexts—ethnographers can help research deliver more value and more directly address organizational problems. In doing so, ethnographers can create feedback loops, redefining what "research" is and showing how it can have different types of impact.

Second, I argue that ethnographers can articulate new ways of working with and within organizations by approaching research through the lens of risk. Models of more "careful" democratization often advocate for drawing boundaries between tactical and strategic research, where non-researchers do carefully-scoped tactical research, leaving more strategic research as the domain of trained researchers. "Tactical evaluative research is the sweet spot for democratization. Every new feature should at least go through basic usability testing before launch" (Nash 2023). In this vein, strategic work is often seen as more complex and in keeping with traditional imaginings of ethnography, while tactical work is seen as a less skillful and impactful practice.

But this devalues the risk and impact that some tactical research entails. Talk to any usability expert, for example, and they will tell you how much skill and training is required to conduct impactful, high-quality usability studies! In this paper, I move beyond the tactical/strategic divide to show how ethnographers can identify the types of risk that might arise in various projects. This lens of risk can empower ethnographers to develop more contextual and varied working models for how and when to involve non-researchers.

Ultimately, this paper explores how ethnographic skills are needed in this moment of change and can act as a diagnostic and prognostic tool. How can ethnographic skills help us understand the root causes of efforts to democratize research, while also determining ways to intervene and shape what research is and does in the future? Democratization is happening, and, at least for the foreseeable future, is here to stay. Ethnography will not survive if we continue to see

democratization as a threat rather than something that is "good to think with." We have to grapple with trends like democratization seriously if ethnography—and the interpretive and relational value creation it foregrounds—is going to have a role in future organizations.

This paper asks: How can ethnography rearticulate the value and impact research brings and, in doing so, shift organizations towards a different model where democratization does not entail fundamental threats to the research practice and discipline? How can ethnographers have more agency in imagining their futures, which might entail different configurations of labor and practice as work environments themselves evolve?

A (Shortish) History of the Democratization of Research and the Opening up of Knowledge Practices

Where does this idea about democratizing knowledge come from? And what does it entail? Is it singularly tied to this moment of destabilization in the UX and ethnographic community, or is it part of something bigger and broader?

Tracing the history of democratization in organizational research and anthropology, one could argue it is tied to the emergence and growth of phenomena like user experience, design thinking, and agile product development over the last several decades. User experience, which emerged in the early 1990s as a way of pushing technology companies to think beyond computer interfaces and usability (Stevens 2019), advocated for involving users in the design process through iterative feedback.

Similarly, design thinking, which emerged in the early 2000s with an emphasis on empathy and human-centered problem-solving (Knemeyer 2015), contributed to the idea that anyone involved in product development could surface and engage with user needs. Moreover, design thinking opened up a space for research to become its own function, by driving a need for "insights" to help companies understand how designs impacted end users.

On the other hand, agile product development, with its focus on cross-functional collaboration and iterative testing, created a role for tightly-scoped research and insights to interface with engineering frameworks and timelines. Through its emphasis on "lean" product development, agile also gave rise to hybrid roles in which PMs were encouraged to own the entire product lifecycle—including research.

More recently, the push towards democratization has intersected with and been enabled by a variety of research tools that claim to increase research efficiency and speed. Such tools open up parts of the research process that have traditionally been

managed by researchers—like participant recruitment and data analysis—giving easier access to non-researchers. For example, remote unmoderated testing platforms provide quick access to participants, while AI data analysis software empowers anyone to analyze research data. Driven by notions that research (particularly of a more academic nature) takes too much time and effort, these tools aim to automate certain aspects of the research process, delivering insights faster and cheaper. But as these emergent technologies make things that formerly required craft and expertise available to everyone, they do not necessarily "democratize" the interpretive skills required to use those technologies in meaningful ways.

One could also argue that the democratization of UX research is part of a broader movement to make knowledge practices more accessible and open (Levin and Leonelli 2017; Kelty et al. 2015). Take, for example, the recent open science movements in the United States and Europe: these aimed to make scientific data and papers more widely available, to increase the transparency of scientific research processes, and to make knowledge from publicly funded research available through open-access publishing. Similarly, the citizen science movement has tried to open up the very practice of science to non-experts over the last several decades. This has given rise to many collaborative scientific endeavors (Rosas et al. 2022; Polleri 2020; Grace-McCaskey et al. 2019), which have also been enabled by the rise of open technology platforms that enable crowdsourcing and collaboration.

This turn toward "openness" has also extended to the social sciences, as academics have sought to make research more participatory and to foreground the experiences of those who are often the "subjects" of research, particularly under the broad banner of "participatory action research" (Participatory Action Research, n.d.). Here, communities involved in and affected by research are considered experts and are encouraged to produce their own insights, foregrounding their own lived experiences. Take, for example, the EPIC paper "Empowering Communities: Future-Making through Citizen Ethnography," which explores how the democratization of ethnography on a project dealing with high rates of youth suicide empowered local communities to generate, synthesize, and act upon findings (Badami and Goodman 2021). Drawing inspiration from feminist and post-modern studies, with their focus on structural inequality and power relations, these participatory approaches aim to decolonize knowledge and history. They question and challenge the power dynamics in more traditional and academic forms of knowledge production, arguing that these are extractive, giving little back to the communities from which knowledge is taken.

The turn to participation takes a particular flavor in design, with the rise of participatory design and co-design from the 1960s onward. These approaches aim to involve all stakeholders—employees, partners, customers, citizens, end users—in the design process, to ensure designs meet user needs and are usable (Wikipedia contributors 2024). Here, the public is invited to participate throughout the design process, from problem definition to design exploration (Bødker et al. 2022; Di Russo 2012; Asaro 2000). What stands out is how these movements to open up knowledge practices place less value on formal training and technical/academic expertise, often valuing and recognizing lived experience instead. However, in practice, more participatory approaches do not always break down the divide between stakeholders and end users; often, they simply end up enlarging the number of stakeholders exposed to the research and design practice.

Ultimately, democratization is not just about re-organizing how data is collected and interpreted. It reflects broader shifts in knowledge and value systems and the organization of labor. When it comes to the democratization of research more specifically, what does the push to redistribute certain types of power and agency say about the organization (Tang 2023)? How does the push toward democratization create different "epistemic cultures" (Cetina 1999) in different organizations? When research is democratized, what power is shared and what remains? What questions and outcomes are important, and for whom?

The Problem Space: A Reflection on the State of Democratized Research in the Author's Organization

If we see the democratization of research as a generative moment, as an invitation to leverage ethnography to look into and interrogate the dynamics of organizations (Madsbjerg 2014; Flynn and Lovejoy 2014), what does democratization look like in practice? How does digging into the specifics of how democratization is enacted—the motives, dynamics, outcomes—reveal how democratization is varied and contextual as a practice?

I began working for the County and City of San Francisco in 2021, about a year before the large tech layoffs I referenced at the opening of this paper began. After 4.5 years at a large social media company, I left a team of 1000 researchers and joined a team of 1: me. On my new team, San Francisco Digital Services (colloquially referred to as "SF Digital Services"), I was responsible for leading the research practice for a team of 50 (made up of a design practice of two researchers including me, two service designers, and three UX designers). I had to ensure research

contributed to the team's mission to scale government digital services across the city through partnerships, consultations, and the building of new digital tools.

This move was a conscious career decision. I knew that leaving a highly-resourced environment where I was surrounded by many like-minded colleagues would involve a good deal of adjustment. To this end, one of my core responsibilities was building an inclusive research practice; however, my new role came with a limited number of templates, resources, research tools, and no dedicated research operations team or person. The researcher who had previously held my role had created some resources and practices, but these were not widely adopted by the team and were seen as providing limited value. As a result, I had to personally carry out (and create processes for) all aspects of the research process, from participant recruitment to data storage. In summary, the research practice was not set up in a way to facilitate the scale articulated in the team's mission.

As I started to set up the research practice, one of the biggest adjustments I encountered was the push to "democratize" research in the organization. I was told (or perhaps strongly encouraged) to enable others, not just researchers, to lead and conduct research. This is a typical setup in civic and government organizations, which tend to be less well-resourced than private sector. However, this job marked the first time I had considered the term and concept. In academia, as I pursued a PhD in anthropology, knowledge was so specialized and resources were so guarded that the idea a non-academic could conduct research was preposterous. (After all, academia isn't called the ivory tower for nothing.) In the private sector, the quality and rigor of knowledge were highly prized, largely because UX research was fighting to be taken seriously in a data-driven organization. Moreover, sensitive research had been leaked to the press, leading to debates about the rigor of conclusions and ultimately leading to more of a rigorous process around the production of knowledge. Given my background and previous work experiences, the idea that non-researchers should participate in research was understandably alien.

As I digested my new team's democratization mandate, several alarm bells sounded in my head. Did non-researchers have the skills to navigate complex problem spaces, scope ambiguous research questions, and produce high-quality insights? Did they have the conceptual tools to generate impact and influence a complex organization like the government—a feat that even senior researchers would struggle to do? Moreover, why was research, not other disciplines, being pushed to democratize? Was research considered a low-skill activity if anyone could do it? What did that say about the value placed on research within the organization?

I was particularly sensitive to this line of thinking for two reasons. First, I had spent over ten years of my career studying and gaining academic accolades in ethnographic research. The suggestion that my expertise did not require training and skill was painful. Second, I had come from a private sector organization where ethnography and other qualitative research methodologies were often less valued and respected than quantitative methods (Levin 2019). Because qualitative research inherently deals with small sample sizes, it was not seen as "rigorous" or "objective" compared to quantitative methods like surveys or log data analysis. For example, when I tried to argue that insights from qualitative research could be generalized through careful sampling and by exploring underlying themes and factors, I received significant pushback from quantitative experts. They were so enmeshed in a statistical understanding of "representativeness" that they were affronted by the suggestion that qualitative research could extrapolate in certain, more inductive ways (Smith 2018).

I began to wrap my head around the idea of democratization at SF Digital Services—giving myself space to grapple with questions and concerns. I took stock of the previous research that had been done, to learn about the questions the team had asked, the type of impact they wanted to have and the ways they went about generating insights. A good chunk of the team's research had been done through a rolling research program: the researcher before me had set up monthly research sessions to generate qualitative insights. In this model, the researcher had created templates and guidelines, but had left non-researchers responsible for everything from participant recruitment to conducting studies to reporting on insights. In other words, the researcher had empowered non-researchers to do their own research, but had not provided guidance or oversight on execution and analysis.

As I dug into the decks and presentations that had been created, and as I talked to teammates who had been involved in the program, I saw how this approach was problematic. Many of the resulting slide decks and reports contained misrepresented and imprecise insights. The team had not considered how participant recruitment and sampling might affect their work. They had reported on qualitative insights in a way that did not carefully engage with sample size and composition and did not focus on the "why" of the findings. They approached the data through a falsely quantitative lens, saying things like "⅖ users like X feature". They also focused on user preferences instead of user behaviors. Ultimately, the team had not been coached on what qualitative research could and couldn't do. Their access to research templates and tools had not resulted in rigorous work.

These errors and misconceptions are common throughout qualitative research and were not unique to my team. Less experienced qualitative research practitioners

frequently try to make quantitative claims with small sample sizes, but because qualitative research does not try to and can never be statistically representative, such claims are not possible. In addition, less experienced qualitative practitioners often focus on user preferences ("I like design A better than design B") without digging into the "why" behind such preferences. This focus on preferences (rather than behavior) can also create inaccurate data, because what people say and do are often not the same. Often what matters more in small-scale qualitative data—particularly in usability studies—is observations about how features do or do not enable people to complete tasks. In previous roles, I had seen how such errors or misconceptions had led to bad product decisions: when researchers used small-scale qualitative research to conclude that one prototype was more appealing than the other, instead of focusing on how and why certain aspects of each prototype worked well or poorly.

While the rigor of the qualitative research was certainly a problem, there were several other reasons the rolling research program was limited in its impact. The questions had been selected to fill rolling research slots, rather than through the lens of riskiness or potential for impact, and therefore focused on small and non-urgent questions. The insights were scoped only to singular features, rather than overall patterns and behaviors. The solutions focused on organizational pain points rather than user needs. As a result, non-robust research data had potentially led to poor product decisions and outcomes, making it harder to "surface and evangelize…accurate and fact-based about user needs, expectations, and behaviors" (Carey 2019).

It is common, especially in less mature teams, for people to lack coaching on the importance of examining the tradeoffs with doing research. This is problematic for research for a number of reasons: requests can outpace capacity (raising questions about whether research is a good use of the team's time), and research effort may not always lead to impact. Moreover, research becomes a crutch for decision-making—something that teams turn to if they are unsure about the right path for a given product or project. In these cases, it can be beneficial to empower teams to say no to primary research and coach them instead on the range of approaches they can take to gain research-like insights, such as literature reviews, competitive/heuristic assessments, and stakeholder interviews.

I carried these questions and observations into my new research role, paying attention to how my colleagues talked about the importance of research and approached the craft of qualitative methods. Over time, as I carried out and

supported several research studies, I realized that the push to open up research to non-researchers glossed over critical skills that researchers brought, which could ultimately influence key decisions in the organization. Firstly, researchers cultivate the skills to frame problem spaces and wade through complexity and ambiguity through discovery and research design. We are taught the importance of finding hypotheses to test, and identifying where there is the greatest ratio of effort to impact. This is often done by treating research requests as research projects in and of themselves; spending extra time at the beginning of projects to ask clarifying questions often yields research that is more tightly scoped and more precise on the intended impact.

For example, several team members approached me for help scoping a research project on why editors of city websites weren't adhering to guidelines and rules around content best practices when they produced new pages in a content management system. The question was a good one, but I sensed that it was scratching the surface of the problem, and that we would have better insights and more impact if we framed the problem differently. I asked my team members to elaborate on why they thought editors were behaving in certain ways; this revealed that problems that manifested in the creation of content in a content management system were influenced by many upstream things, like onboarding and training. As a result, we expanded the scope of the research to explore the major challenges and pain points editors were experiencing throughout their editing journey. As a result, the team realized that compliance was not really the issue, and that a misalignment between editor skills and values was creating conflict instead.

Another skill researchers bring, which often goes unremarked upon, is the ability to carefully select and tailor approaches and insights. Researchers use multiple inputs—questions to be answered, resources and time available, potential for impact—to gauge which method works best for a given project. Once the project is going, researchers are trained to recognize and respond to various types of bias, such as the selection of certain participants or the disconnect between what people say and what they do. After data has been collected, researchers use their knowledge of stakeholders and the organization to frame insights as a response to key questions and to showcase a strong point of view, which ultimately helps the research resonate and have more impact. Having tools available does not lead to insights (Mitra 2020), because "Anyone can collect data—it's knowing how to collect it, what to do with that data, and synthesizing the results into valuable actions that's the hard part" (Carey 2019).

For example, when our team was just beginning to think about a visual refresh for SF.gov, I started a research project to explore what makes a city website feel like

a city. This research asked participants to compare their experience across several municipal websites, to reflect on what elements of each website reminded them of the city and enhanced or detracted from their overall experience. One path forward would have been to present the results in a straightforward way, outlining what was good and bad about each website and why. However, I knew the team needed other help and impact. They needed to understand what role aesthetics played in the overall experience of using a government website. For example, was usability more important than the look and feel of a website, and if so, what did that mean for future designs? Moreover, in the early stages of a redesign, the team needed high-level guidance about design values rather than insights at the feature level. Because I had the research skills not just to collect data, but to frame the insights in response to issues the organization was facing, the research had widespread impact.

Using Ethnography to Understand the Needs and Values Underlying Democratization

Given my concerns and fears, I initially resisted democratization. I doubled down on research quality and process, creating intake and review processes, as well as an 8-week "research curriculum" to teach non-researchers fundamental qualitative research concepts. But my resistance to democratization was met with resistance itself. My efforts to bring a focus on rigor and quality were met with skepticism and frustration. Team members didn't see problems with what had been done before; instead, they saw me as a gatekeeper of research, as someone who was preventing the team from connecting and empathizing with users. As a result, my team members didn't view me as a trusted expert. They felt forced to jump through seemingly unnecessary hoops when their process worked fine.

Pretty soon, I knew my approach—resisting democratization by pushing my team members to adopt my own standards and values around research rigor and quality—was not working. The overall volume of research decreased, and stakeholders stopped coming to me with questions. So I did what any researcher would do: I approached the phenomenon as a mini research project, using ethnography as a way to understand the root causes and themes, and to explore the historical and current research practice (Nash 2023; Knoll 2023; De Paula, Thomas, and Lang 2014; Thomas and Lang 2014; Bandyopadhyay and Buck 2015). What was the organization trying to achieve with democratization? What did it need and value?

What was it optimizing for? How had research been conceptualized and leveraged in the past? What was the culture of decision-making like?

I began with informal interviews with key members of the team, making sure to speak to people from a variety of backgrounds. I wanted to know about their past experience was with research, and how it had been shaped by the context in which they had worked. One teammate, who came from another established civic tech organization, explained how research had been a "team sport." For them, weekly team meetings had included video clips of user research, and widespread participation in research had increased the organization's connection to its end users. Another teammate from a small tech startup explained that gathering quick feedback on early-stage features had been standard practice for product managers because researchers were not present on the team.

All of this pushed me to ask important questions. What did my colleagues' desires around democratization reveal about their own needs, and the needs of the organization, when it came to research? My team wasn't an academic institution or a large tech company, so did it matter if a teammate produced a sub-par research report, or none at all? Turning my questions inward, how had my past experiences shaped my views? What other compromises was I willing or needing to make? What expectations were reasonable to place on others, when it came to following process and reporting on insights?

As a start, I explored how the needs around democratization were different in my current organization compared to the private sector. My new team was not focused on increasing speed and efficiency, at least not at the outset. The desire to democratize research was not driven by a need for faster insights; rather, it was born from a recognition that one researcher would struggle to support a team of 50, much less an entire city government. If the team's overall mission was one of scale, to provide internal tools and processes to level up city staff, how could one person support that? (There is also something meta here, in that the team's overall mission is one of democratization in and of itself: SF Digital Services saw itself as providing the tools and frameworks to help other city staff deliver digital services themselves. Perhaps this contributed to the idea that knowledge production could and should be participatory.)

Similar to the needs around democratization, the team's capacity for and approach to research was highly contextual and different from the private sector. In my previous role, stakeholders were often too busy to participate in (or did not place value on participating in) research studies. It was difficult to "bring people along" with the research, as I was constantly vying for their time and attention. On my team, because the pace of work was slower and more deliberate, stakeholders had more

time to engage with and even conduct research. They respected that researchers were as busy as other team members, and did not expect researchers to always carry out studies. However, I noticed that when non-researchers came up with and initiated research projects, they tended to focus more on methods than questions. In other words, they more intrinsically saw the value in talking to users of their products (compared to the private sector, where research is often an afterthought compared to A/B testing), but didn't necessarily know how to conduct or initiate research in an impactful way.

Another thing that emerged in informal interviews with colleagues was that not everyone had the same definition of "research." For example, tensions arose when a product manager wanted to do research with users to explore reactions to an early-stage prototype of a new content management system[5]. Was this research? Who were they talking to, what hypothesis and problem were they exploring, and what changes would be made? In this case, overlapping roles and terminology made the definition of research murky. In another example, a content strategist and researcher both wanted to use a card-sorting tool, which is often construed as a "research tool," but for different reasons. The researcher planned to use the tool to explore residents' mental models around groups of city services, while the content strategist planned to use the tool to get buy-in from city staff. Tensions arose when both the researcher and content strategist presented the outcome as "findings." Different levels of rigor had been applied to the methods and framing of the two projects, which created confusion around how each set of "insights" should be leveraged and used to make decisions.

"Research," I realized, was thought of and used much more broadly than I or my research peers would have intended (Carey 2019). It was synonymous with gathering quick input or feedback, generating empathy, building consensus, and doing community engagement (Democratic Society 2023). This took the shape of comments like, "I got feedback from some colleagues about X" or "We did some quick research to validate Y" or "We need to hear from our users more."

This is not to say that things like customer feedback aren't useful or valuable activities; they are "great for building empathy, sparking product ideas, and providing real-world examples of product usage… [but are not] representative or rigorous" (Nash 2023; Mateljan 2022). But as Jen Pahlka artfully describes, user research and things like "public input" are not the same (Pahlka 2024). Feedback is often unconcerned with bias around sampling and self-described behavior, and is focused on "listening" rather than exploring problems or proving hypotheses.

Over several years, what I learned was: construing research with these activities created confusion over goals and outcomes. Missing from these so-called "research" activities was a focus on testing hypotheses, methodically and rigorously approaching problems rather than solutions, and centering the pain points of users over the organization. But for my colleagues, these considerations were not important, because the activities they called "research" ultimately had different goals, like building consensus and generating empathy. For my colleagues, it wasn't important to be rigorous or to focus on quality, and as a result, my own emphasis on quality and rigor landed on deaf ears.

This was a pivotal moment for me. I realized that the organization didn't just need more clarity on who should do research. It was at a state of maturity (Bandyopadhyay and Buck 2015; Metzler 2020) where it needed help with knowing how to make decisions and create clear priorities—for which research is one of several valuable tools (Mazur 2023; Blom 2023; Dombrowski 2021; Yost 2016). As a result, I needed to redefine and clarify what "research" was and how this research vision could better serve the organization's problems, needs, and goals. And even more importantly, I needed to help the organization increase its capacity to make informed and careful decisions, while showing that research was one tool among many. Ultimately, my role transformed into one of organizational change, where I was attempting to exert influence on and change multiple levels of the organization, as a way to carve out a different role and set of practices around research.

Using the Lens of Risk to Deliver Impact with Care

As I engaged with this challenge—helping the organization improve its decision-making skills while also showing how research should and shouldn't play a role—I decided to use the lens of risk. I have drawn on this concept throughout my career, as I have struggled to handle high volumes of research requests. Seen through the lens of risk, the question becomes not "Should we do this research?" but rather "What is the risk of doing or not doing this research?" (Cuciurean-Zapan and Hammel 2019; Lalley 2019; UX Guys 2016).

This reframes the conversation around tradeoffs and encourages people to project into the future, to consider the possible outcomes that research (or a lack thereof) can lead to. By opening up the logic around research decisions, stakeholders are empowered to become better decision-makers across the board, using risk as an input alongside impact and effort. Instead of feeling compelled to conduct research on features before launch, out of worry or a need for validation, stakeholders can

assess the risk of decisions and courses of action, while also considering the various tools that are available to them to mitigate those risks (Belt 2019).

In this paper, I advocate for researchers and ethnographers to use the concept of risk not only to determine whether research should be done, but also how it should be done and by whom. The focus is on how opening up research can create different types of risk at various levels of the organization. This creates a clear set of considerations for researchers to guide how they include non-researchers in and scope research projects.

Below, I outline the various risks that opening up research to non-researchers can create, and provide examples of how this manifested in specific projects.

Risk to the Product and End Users

If research is a crucial aspect of product decision-making, what happens if insights aren't robust and lead to poor product decisions? If qualitative data is not interpreted with care, or if personal experiences and anecdotal observations are encoded as research, the wrong features might be selected and invested in. This may have adverse effects on end users, inadvertently making features harder to use, or creating features that do not solve actual user problems (and overall decreasing the value products provide to end users).

For example, in my previous role, I conducted research on settings for a social media app. At the time, it was customary for us to select a "random sample" of participants, which often included a mixture of age, gender, income, and occupation. This research often led to conclusions that participants wanted more control over their settings, which led the product team to create granular controls for various aspects of the settings experience. However, I prompted the team to think more carefully about sampling, and to consider how participants with lower digital skills could provide pivotal insights into the settings experience. When we included these participants in the research, we learned that granular settings created serious problems for people with lower digital skills. Participants had little awareness, knowledge, and experience with settings, and as a result, felt fearful and unempowered by granular controls, because these placed an increased burden on participants to learn about and manage complex user experiences—often through the language and framing of internal features.

Risk to the Research Practice

Given research teams are often small and understaffed, saying yes to one project almost invariably involves saying no to another. If non-researchers are not encouraged to use a rigorous process to determine whether research should proceed, they risk taking researchers' time away from other more impactful projects. Over time, the research practice risks having less impact and becoming narrower and less foundational—focused on answering specific questions rather than driving strategy—unless non-researchers are trained and encouraged to look beyond tactical, evaluative work. Moreover, there is a risk that research becomes a panacea for everything, diminishing broader problem-solving skills that can leverage other forms of expertise and approaches to decision-making.

For example, a team member approached me with a research project to evaluate changes to the page where editors would log into the content management system for SF.gov. The team wanted to "make sure" that the changes were not going to have a negative impact, so they assumed research could help answer that question. However, when I took them through research intake, I encouraged them to think about how significant the change was, and what bad outcomes they thought might arise. As we talked, they realized the change was not large enough to significantly disrupt the page's user experience. In the end, they decided not to do research, which freed up my time and their time to do other work.

Risk to Participants

A foundational part of training to be a researcher is learning about participant safety and ethics. This leads to practices like ongoing consent, verbal informed consent, anticipating questions that could bring up trauma, carefully asking about highly personal experiences, deleting data after a certain period of time, and more. We are often taught to put the safety and experience of the participant first, above all else, but what happens if we allow others without this training to do research with these populations? What potential harm can this cause to participants, and how might that reflect poorly on the team and the product?

Example: When our team began working with a Department that administered a permit for street vending, we decided to research the user experience of the permit form, because we knew that the people who needed the permit were from marginalized communities—according to the city partners and community organizations who helped applicants fill out the form, this group of people did not speak English, had lower rates of literacy, and also had lower digital skills. However, neither I nor the other researcher on the team were fluent in Spanish, so we had to

figure out who would conduct the research. A non-researcher on the team was fluent in Spanish, so she assumed she would conduct the interviews. However, as researchers we were highly concerned about participant safety: was it ethical or appropriate for someone without training in research ethics or trauma-informed research to work with marginalized communities? What would happen if the participant disclosed something highly personal or if a question brought up past trauma—would the non-researcher know to pause the recording or switch from questioning to comforting?

Risk to the Organization

When we do research that spans a wide range of topics and scopes, there is always the potential that we will conduct research on things that are political or sensitive—and which might lead to negative press, lawsuits, or government inquiry. This might entail recruiting participants to interact with a form for a competitive grant application before the application is available to the broader public. It might also entail researching something that carries legal risk to the organization, where if that research was leaked, the organization could be subject to lawsuits or other public inquiry. This type of risk is often less visible to individual contributors, given their more limited view of the organization, but it is still important.

Putting it all Together: Leveraging Risk in Decision Making

So how do we as researchers navigate these risks to make decisions about how to support our teams with research? On some teams, researchers advocate for guardrails, and for reducing the scope of work by non-researchers to tactical research. But on my team, I created frameworks and processes that opened up the decision-making process to non-researchers and taught them how to assess tradeoffs around risk, impact, and effort.

Through a process of research intake, we asked non-researchers questions like: "What would happen if we didn't do the research?" and "Will you be working with any sensitive populations?" and "What will change as a result?" We also assessed the positionality of non-researchers, exploring their past experiences with research and interviewing, their willingness and availability to learn, and their familiarity with and expertise in the problem space. This took the shape of a guided conversation, where non-researchers were asked to think through questions on their own, and then were encouraged to discuss them in the open later.

This encouraged the team to spend more time at the beginning of a research project than they otherwise would have, as an investment into thinking through potential positive and negative outcomes. It also increased their decision-making skills and empowered them to assess the best course of action—increasing the organization's overall skills. Over time, I saw my teammates' frustrations with what they had formerly perceived as gate-keeping fade, as they were exposed to the logic that led to nuanced decisions. Conversations shifted from "no you can't do this" to "here's what could happen if you do this." I also saw my teammates start to understand and empathize with the decision to have trained researchers conduct studies with marginalized populations or highly ambiguous problem spaces. Through the process of intake, they were exposed to the importance of experiences with ethics and trauma-informed research—and the potential harm that could come to participants. They were also exposed to the complexity of scoping and prioritizing impact for nebulous, ill-defined projects.

Ultimately, the process of research intake, and the surrounding conversations it generated, made clear the decisions about:

- Whether the research should proceed, or whether we should leverage other decision-making tools, like relying on external research or doing an expert evaluation
- Who should lead the research, and how closely the researcher should collaborate on various aspects of the work
- How narrowly or broadly the research should be scoped
- Any precautions that need to be taken or check-ins we need to have to go over specific topics or practices.

This is not a perfect approach, but it has allowed us to open up the research practice in a way that helps the organization increase its ability to make decisions while mitigating various types of risks. What this requires of us as researchers, however, is an investment in being present through the entire research process, not just at the beginning—as both players and coaches. We must learn to say no in a way that resonates and provides other options, and we must invest in continued relationship building with team members (often outside of the research realm, by embedding ourselves in teams), which buys us the goodwill to shape the outcomes of projects we do not always own.

Conclusion: Can Ethnography Save Itself?

As ethnographic practitioners, we are inherently familiar with the concept of democratization. Through the very act of doing ethnography, we accept that our expertise is complementary and often secondary to the expertise of participants, whose lifeworld we seek to elevate. Less familiar, however, is the attempt to democratize our own expertise. We reject (often in a knee-jerk sort of way) the idea that our stakeholders have equal experience and expertise with research ideas, methods, and practices.

So what happens, then, when we turn this phenomenon—this insistence that anyone can do research—into an object of fascination and inquiry? Ethnography is uniquely suited to this task. How might our skills help us navigate this moment of institutional and disciplinary precarity?

Ethnography isn't going anywhere. The title above is surely sensationalist. But my point is: simply pushing back on democratization, by insisting on the primacy of our own expertise, is self-defeating. Doing so does not address why democratization is often so appealing. As such, this paper encourages ethnographers to put aside our pride, to push against our urge to protect our discipline, and to instead see ourselves as organizational change leaders. The goal of this paper is not to argue for or against democratization, but rather to use the push for democratization as a way to interrogate an organization's needs and figure out ways for research to have more impact by satisfying those needs. "Research has been done, and will continue to be done, by people who don't have "researcher" in their title. It's imperative that we improve the quality of their work, rather than pretend that it doesn't exist" (Sirjani 2020).

Why does this matter? Why should we bother? The stakes here are bigger than our job security. If we do not insert ourselves into the democratization debate, shaping how it plays out in practice, research will have less impact. Poor decisions will be made. People might be harmed. What counts as research will not be shaped by those who have the most power to make sure research is impactfully and carefully done. How might we turn our fear into (guarded) curiosity by leveraging our ethnographic skills to ask questions like: how is "research" getting redefined and why? What counts as training and expertise, and why? What problems are truly being solved here, and why?

So what does "good" democratized research look like? Where is the most impact to be had as a researcher, and how—when considering how to intervene in democratized contexts? Through the lens of risk, researchers can make better

decisions about where to invest their time and energy and where to push back against stakeholder notions of research topics and practice. Researchers should be empowered to define democratization on their own terms, based on what the organization and end users need. As such, this paper is an invitation to consider how to leverage ethnography to intentionally shift and redistribute power throughout the research practice and process.

About the Author

Nadine is an anthropologist and UX researcher, with a focus on strategic research, improving equitable access to technology, and combining ethnography and quantitative research. She currently leads the design and research team at the City and County of San Francisco's Digital and Data Services team, whose work focuses on improving the digital experience with permitting and forms, as well as helping the city centralize and scale its digital resources on SF.gov. She has previously worked at Meta, UCLA, and the University of Exeter, after completing a DPhil in Social Anthropology at Oxford.

Notes

Many thanks to Kate Zykowski, Matthew Bernius, Cyd Harrell, Taylor Nelms, Jennifer Ng, and Andrew Symington for their thoughtful feedback and insights on this paper!

1. A turn which is, in some ways, reminiscent of how research was done in organizations before the "professionalization" of user research.

2. The term "human-centered" refers to the idea that people (customers, users, etc.) are at the center of business and creative processes. With a human-centered approach, teams are empowered to "design products, services, systems, and experiences that address the core needs of those who experience a problem" (DC Design 2017).

3. According to one survey, over two-thirds of designers and half of PMs do research, with that research skewing toward evaluative and qualitative work (Akhmedov 2023).

4. Here, I follow a similar approach as this paper (Cefkin, Anya, and Moore 2014), which takes the rise of more open and distributed work, and questions the conditions which give rise to it as well as the impact it has on work relations.

5. This product manager had come from a company where their responsibilities were porous—something that is common in smaller, less well-resourced organizations, where people must function as "jack-of-all-trades."

References Cited

Akhmedov, Nurkhon. 2023. "Top-3 UX Research Trends in 2023: What UX Designers and Product Managers Can't Ignore." *Bootcamp*. September 16, 2023. https://bootcamp.uxdesign.cc/top-3-ux-research-trends-in-2023-what-ux-designers-and-product-managers-cant-ignore-2a910bb7f7dc.

Antin, Judd. 2023. "The UX Research Reckoning Is Here." *One Big Thought*. May 4, 2023. https://medium.com/onebigthought/the-ux-research-reckoning-is-here-c63710ea4084.

Asaro, Peter. 2000. "Transforming Society by Transforming Technology: The Science and Politics of Participatory Design." *Accounting, Management, and Information Technologies* 10 (4): 257–90.

Badami, Sumant, and Sophie Goodman. 2021. "Empowering Communities: Future-making through Citizen Ethnography." *Ethnographic Praxis in Industry Conference Proceedings* 2021 (1): 282–302.

Bandyopadhyay, Kirsten, and Rebecca Buck. 2015. "From UI to UX: Building Ethnographic Praxis in a Usability Engineering Culture." *Ethnographic Praxis in Industry Conference Proceedings* 2015 (1): 156–143.

Baxter, Kathy, Catherine Courage, and Kelly Caine. 2015. "Evolution of User Experience Research." *EPIC Perspectives*. November 10, 2015. https://www.epicpeople.org/evolution-of-user-experience-research/.

Belt, Sara. 2019. "Accelerating User Research: How We Structure Insights for Speed at Spotify." EPIC Perspectives. October 28, 2019. https://www.epicpeople.org/accelerating-user-research/.

Bien, Heather. 2022. "Top 2023 UX Research Trends: Continuous Discovery, Democratization, Enable Increased Customer Focus." Spring blog. December 13, 2022. https://sprig.com/blog/ux-trends.

Blom, Ilse. 2023. "UX Research Prioritisation: Avoiding Low Impact Projects." *UX Insight*. January 23, 2023. https://uxinsight.org/ux-research-prioritisation-avoiding-low-impact-projects/.

Bødker, Susanne, Christian Dindler, Ole S. Iversen, and Rachel C. Smith. 2022. "What Can We Learn from the History of Participatory Design?" In *Participatory Design*, edited by Susanne Bødker, Christian Dindler, Ole S. Iversen, and Rachel C. Smith, 15–29. Cham: Springer International Publishing.

Carey, Ellen S. 2019. "Not Just Anybody Should Do User Research." *UX Booth*. August 14, 2019. https://uxbooth.com/articles/not-just-anybody-should-do-user-research/.

Cefkin, Melissa. 2009. *Ethnography and the Corporate Encounter*. Berghahn Books.

Cefkin, Melissa, Obinna Anya, and Robert Moore. 2014. "A Perfect Storm? Reimagining Work in the Era of the End of the Job." *Ethnographic Praxis in Industry Conference Proceedings* 2014 (1): 3–19.

Cetina, Karin Knorr. 1999. *Epistemic Cultures: How the Sciences Make Knowledge*. Harvard University Press.

Cuciurean-Zapan, Marta, and Victoria Hammel. 2019. "Designing Good Jobs: Participatory Ethnography and Prototyping in Service-oriented Work Ecosystems." *Ethnographic Praxis in Industry Conference Proceedings* 2019 (1): 514–32. https://www.epicpeople.org/designing-good-jobs-participatory-ethnography-prototyping-service-oriented-work-ecosystems/.

Cultural Anthropology. 2018. "Academic Precarity in American Anthropology: A Forum." Member Voices, *Fieldsites*. May 18, 2018. https://culanth.org/fieldsights/series/academic-precarity-in-american-anthropology-a-forum.

DC Design. 2017. "What Is Human-Centered Design?" *DC Design*. August 14, 2017. https://medium.com/dc-design/what-is-human-centered-design-6711c09e2779.

Democratic Society. 2023. "On a Journey to Democratising Research." *Democratic Society*. February 12, 2023. https://medium.com/@demsoc/on-a-journey-to-democratising-research-23d70acfa19.

De Paula, Rogério, Suzanne L. Thomas, and Xueming Lang. 2009. "Taking the Driver's Seat: Sustaining Critical Enquiry While Becoming a Legitimate Corporate Decision-Maker." *Ethnographic Praxis in Industry Conference Proceedings* 2009 (1), pp. 2–16. https://www.epicpeople.org/taking-the-drivers-seat-sustaining-critical-enquiry-while-becoming-a-legitimate-corporate-decision-maker/.

Di Russo, Stefanie. 2012. "A Brief History of Design Thinking: How Design Thinking Came to 'be.'" 2012. *I Think, I Design*. June 8, 2012. https://ithinkidesign.wordpress.com/2012/06/08/a-brief-history-of-design-thinking-how-design-thinking-came-to-be/.

Dombrowski, Roberta. 2021. "A Framework for Decision Driven Research." User Interviews blog. October 18, 2021. https://www.userinterviews.com/blog/a-framework-for-decision-driven-research.

Ethnio. 2023. "UX Research Democratization in 2023." June 18, 2023. https://web.archive.org/web/20231108161239/https://ethn.io/blog/ux-research-democratization.

Flynn, Donna K., and Tracey Lovejoy. 2008. "Tracing the Arc of Ethnographic Impact: Success and (In)visibility of Our Work and Identities." *Ethnographic Praxis in Industry Conference Proceedings* 2008, pp. 238-250. https://www.epicpeople.org/tracing-the-arc-of-ethnographic-impact-success-and-invisibility-of-our-work-and-identities/.

Grace-McCaskey, Cynthia A., Briana Iatarola, Alex K. Manda, and J. Randall Etheridge. 2019. "Eco-Ethnography and Citizen Science: Lessons from Within." *Society & Natural Resources* 32 (10): 1123–38.

IAM Design Maker. 2023. "Decoding the Future: Top Trends in UX Research Methods." *IAM Design*. December 19, 2023. https://medium.com/@iam.design/decoding-the-future-top-trends-in-ux-research-methods-a799c871df44.

Balboni, Katryna, Morgan Mullen, Olivia Whitworth, and Holly Holden. 2023. "The State of User Research 2023." User Interviews. 2023. https://www.userinterviews.com/state-of-user-research-2023-report.

Kelty, Christopher, Aaron Panofsky, Morgan Currie, Roderic Crooks, Seth Erickson, Patricia Garcia, Michael Wartenbe, and Stacy Wood. 2015. "Seven Dimensions of Contemporary Participation Disentangled." *Journal of the Association for Information Science and Technology* 66 (3): 474–88.

Knemeyer, Dirk. 2015. "Design Thinking and UX: Two Sides of the Same Coin." *Interactions* XXII.5 September-October 2015: 66. https://interactions.acm.org/archive/view/september-october-2015/design-thinking-and-ux.

Knoll, Alexander. 2023. "Research Democratization: Reframing the Conversation—from Training to Decision-Making." Condens blog. March 28, 2023. https://condens.io/blog/research-democratization-from-training-to-decision-making/.

Lalley, Joe. 2019. "What's the Worst That Could Happen? Reducing Risk through User Research." *UX Collective*. May 23, 2019. https://uxdesign.cc/whats-the-worst-that-could-happen-reducing-risk-through-user-research-8ebe7fbe2d57.

Levin, Nadine. 2019. "Ethnographic Agency in a Data Driven World." *Ethnographic Praxis in Industry Conference Proceedings* 2019 (1): 591–604.

Levin, Nadine, and Sabina Leonelli. 2017. "How Does One 'Open' Science? Questions of Value in Biological Research." *Science, Technology & Human Values* 42 (2): 280–305.

Mack, Alexandra, and Susan Squires. 2011. "Evolving Ethnographic Practitioners and Their Impact on Ethnographic Praxis." *Ethnographic Praxis in Industry Conference Proceedings* 2011, pp. 18–28. https://www.epicpeople.org/evolving-ethnographic-practitioners-and-their-impact-on-ethnographic-praxis/.

Madsbjerg, Christian. 2014. "Happy Birthday, Now Grow up." Keynote Address, Ethnographic Praxis in Industry Conference. https://www.epicpeople.org/happy-birthday-now-grow-up/.

Mateljan, Simon. 2022. "Talking with Users, Isn't Always User Research." *Bootcamp*. April 10, 2022. https://bootcamp.uxdesign.cc/talking-with-users-isnt-always-user-research-f812dcfa5b6b.

Mazur, Nikki Anderson-Stanier Nicky. 2023. "When Should You Skip Doing User Research?" Dscout, People Nerds Ideas. January 31, 2023. https://dscout.com/people-nerds/when-skip-doing-user-research.

Metzler, Brigette. 2020. "Leveling up Your Research and Research Operations: Strategies for Scale." *Ethnographic Praxis in Industry Conference Proceedings* 2020 (1): 203–17.

Mitra, Saswati Saha. 2020. "Undemocratising User Research." *UX Collective*. April 21, 2020. https://uxdesign.cc/undemocratising-user-research-6897e6f4cf72.

Nafus, Dawn, and Ken Anderson. 2006. "The Real Problem: Rhetorics of Knowing in Corporate Ethnographic Research." *Ethnographic Praxis in Industry Conference Proceedings* 2006 (1), pp. 244–258.

https://www.epicpeople.org/the-real-problem-rhetorics-of-knowing-in-corporate-ethnographic-research/.

Nash, Christopher. 2023. "The Do's and Don'ts of Democratization." Conversation hosted by Rally. April 13, 2023. https://www.rallyuxr.com/event/christopher-nash-on-the-dos-and-donts-of-democratization.

Nuez, Alfonso de la. 2019. "Democratization of UX Insights: What Does This Really Mean?" *UX Matters*. 2019. https://www.uxmatters.com/mt/archives/2019/08/democratization-of-ux-insights-what-does-this-really-mean.php.

Pahlka, Jennifer. 2024. "Public Input Is Not User Research." *Eating Policy*. May 17, 2024. https://eatingpolicy.substack.com/p/public-input-is-not-user-research.

Participatory Action Research. n.d. "Participatory Action Research: Resources for Researchers, Practitioners, Organizers, and Communities." Accessed May 25, 2024. https://participatoryactionresearch.sites.carleton.edu/about-par/.

Ronsen, Michele. 2022. "Is 'Research Democratization' Good? Who Benefits and What Should Be Considered?" Curiosity Tank blog. August 25, 2022. https://www.curiositytank.com/blog/is-research-democratization-good-who-benefits-and-what-should-be-considered.

Rosas, Lisa G., Patricia Rodriguez Espinosa, Felipe Montes Jimenez, and Abby C. King. 2022. "The Role of Citizen Science in Promoting Health Equity." *Annual Review of Public Health* 43 (April):215–34.

Polleri, Maxime. 2020. "Being Clear-Eyed about Citizen Science in the Age of COVID-19." *SAPIENS*. July 15, 2020. https://www.sapiens.org/culture/fukushima-citizen-science/.

Sirjani, Behzod. 2020. "Democratization Is Our Job." Yet Another Studio blog. July 22, 2020. https://yetanother.studio/blog/democratization-is-our-job.

Smith, Brett. 2018. "Generalizability in Qualitative Research: Misunderstandings, Opportunities and Recommendations for the Sport and Exercise Sciences." *Qualitative Research in Sport, Exercise and Health* 10 (1): 137–49.

Soucy, Kyle. 2023. "The Problem with Research Democratization That People Aren't Talking about." *Bootcamp*. June 6, 2023. https://bootcamp.uxdesign.cc/the-problem-with-research-democratization-that-people-arent-talking-about-827831c5ee38.

Stevens, Emily. 2019. "What Is the History of UX Design? A Brief Timeline." *CareerFoundry* (blog). July 12, 2019. https://careerfoundry.com/en/blog/ux-design/the-fascinating-history-of-ux-design-a-definitive-timeline/.

Tang, David. 2023. "Democratization Won't Save Research." UX Collective. June 19, 2023. https://uxdesign.cc/democratization-wont-save-research-d438b7b590b8.

Thomas, Suzanne L., and Xueming Lang. 2007. "From Field to Office: The Politics of Corporate Ethnography." *Ethnographic Praxis in Industry Conference Proceedings* 2007 (1), pp. 78–90. https://www.epicpeople.org/from-field-to-office-the-politics-of-corporate-ethnography/.

UX Guys. 2016. "De-Risking Decisions: User Research as Strategic Partner." *User Experience and Product Design Advice from UX Guys*. November 30, 2016. https://medium.com/user-experience-product-design-advice-from-ux-guys/user-research-as-strategic-partner-de-risking-decisions-46738719fa66.

Wikipedia Contributors. 2024. "Participatory Design." Wikipedia, The Free Encyclopedia. May 22, 2024. https://en.wikipedia.org/w/index.php?title=Participatory_design&oldid=1225173221.

Wolstenholm, Jack. 2023. "Democratization by Researcher, Not of Research." Great Question (blog). February 27, 2023. https://greatquestion.co/blog/democratization-led-by-researcher.

Yost, Gina. 2016. "Make Design Decisions with a Purpose." *UX Booth*. January 19, 2016. https://uxbooth.com/articles/make-design-decisions-with-a-purpose/.

Designing AI to Think with Us, Not for Us

STEFANIE HUTKA, *Sendfull LLC*

This paper establishes a framework and toolkit for designing Generative Artificial Intelligence (genAI) that address foundational challenges of these technologies and reframe the problem-solution space. The stakes are high for human-centered solutions: genAI is rapidly disrupting existing markets with technologies that exhibit increasingly complex and emergent abilities, and accelerate scale, cognitive offloading, and distributed cognition. The Problem-Solution Symbiosis framework and toolkit extends, rather than displaces, human cognition, including tools for envisioning, problem (re)framing and selection, interdisciplinary collaboration, and the alignment of stakeholder needs with the strengths of a genAI system. Applying the toolkit helps us guide the development of useful, desirable genAI by building intuition about system capabilities, developing a systemic understanding of emerging problem spaces, and using a matrix to identify if and when to offload tasks to the system. The framework is informed by systems theory, frame analysis, human-computer interaction research on current AI design approaches, and analogous approaches from spatial computing design research.

Introduction

ChatGPT, Open AI's chatbot, reached 100 million monthly active users within two months of its November 2022 launch, making it the fastest-growing consumer application in history (Hu 2023). This milestone is widely seen as a breakthrough moment for generative AI (genAI), a class of machine learning (ML) that generates text, image or sound output, based on user input of the same modalities (Google 2024). Countless companies across sectors, industries and geographies have since integrated genAI into their offerings, optimizing existing *foundation models* (i.e., models trained on a wide variety of unlabeled data that can be used for a broad range of tasks, as Murphy (2022) explains) or training new models for their use cases. Furthermore, a majority of senior executives across industries agree that genAI will "substantially disrupt" their industry over the next five years (MIT Technology Review Insights 2024).

Business leaders seek to "ride the wave" of this disruptive technology (Christensen and Bower 1995), pushing for "AI-first strategy" (Acar 2024). This push results in the application of genAI across use cases, regardless of whether the technology is well-suited to solve for the target audience's needs. Coupled with unique aspects of genAI technology that we will subsequently discuss, this push

Hutka, Stefanie. "Designing AI to Think with Us, Not for Us." *EPIC Proceedings* (2024): 111–132. https://www.epicpeople.org/designing-ai-to-think-with-us. ISSN 1559-8918.

creates new challenges for ethnography practitioners seeking to guide organizations towards building technology that is useful and desirable for people.

This paper seeks to provide ethnography practitioners with guidance on how to "ride the wave" of genAI, especially given the rapidly-evolving, emergent genAI solutions driving organizational strategy. We will begin with an examination of the unique challenges genAI presents for developing human-centered solutions and discuss genAI's potential to perform increasingly complex cognitive tasks, concluding that genAI urgently needs ethnography to build solutions that think with us, rather than for us. We then examine how teams currently design genAI systems, and discuss challenges and shortcomings of existing approaches to problem framing for genAI projects (Yildirim and Pushkarna et al. 2023; Yildirim and Oh et al. 2023), grounded in the practice of understanding and generating solutions for existing problems.

We posit that current approaches to guide the development of human-centered technology do not address a foundational challenge. Namely, genAI is a solution that unlocks new problem spaces, whereas current approaches focus on solving for existing problem spaces. We examine the reasoning mechanisms behind human-centered practitioners' current approaches to problem framing, and the tension between this framing and genAI's solution-first growth. We then offer a new framework to address the foundational challenge, helping reframe the interplay between problem and solution space to align stakeholders' needs with strengths of a genAI system. We offer three tools that follow from the framework, which practitioners can use to help teams build useful, desirable genAI solutions that think with us, not for us.

Why GenAI Needs Ethnography

There are two unique characteristics of genAI systems that simultaneously make human-centered solution development challenging, while also creating the potential for powerful – and potentially harmful – capabilities. First, with each new generation of foundation model, it becomes increasingly difficult to anticipate the system's behaviors *a priori* without ample hands-on experimentation. This challenge is rooted in three genAI system capabilities: *scale*, *homogenization* and *emergence* (Bommasani et al. 2021, De Paula et al. 2023; see Endnote 1 for further details).

Scale enables models to ingest massive amounts of data. Homogenization enables models to adapt to multiple tasks, modalities and disciplines, obviating the need to build separate models for each function. Emergence enables models to exhibit unprecedented and unexpected capabilities. For instance, when engineers

make a small change to one part of the model to better solve certain types of problems (e.g., math problems), new behaviors emerge as a result of the system's complexity (Klein 2024). These changes are not pre-programmed or trained into the model. The system may also "perceive" patterns that are not there (i.e., hallucinations), leading to unreliable output. Hands-on experimentation with genAI models has been proposed as an approach to develop intuition about the solutions' vast and often unpredictable capabilities (e.g., Mollick 2023, Walter 2024).

Second, as genAI systems' capabilities rapidly expand, so to do the range and complexity of tasks we can perform using these tools, and by definition, the potential problems that the technology could be applied to solve. We can examine ChatGPT's evolution as a case study (Open AI 2024a). The initial foundation model, GPT-3.5, was capable of natural language understanding and generation, with the ability to maintain context within shorter conversations. GPT-4, introduced in March 2023, offered more advanced problem-solving skills and more reliable responses, could retain context over extended conversations, and included multimodal abilities (i.e., processing and generating a range of data types beyond text, such as images and audio). GPT-4o, introduced in May 2024, honed its multimodal capabilities and added real-time functionality, with latencies comparable to human conversational turn-taking (Open AI 2024b, Stivers et al. 2009). This tool is not only capable of advanced cognitive tasks (e.g., thinking through a math problem), but also the ability to carry on a conversation approaching human capabilities and adopt a personality of the user's choosing.

Looking ahead, the technology industry widely accepts that artificial general intelligence (AGI) – roughly defined as "AI systems that are generally smarter than humans" (Open AI 2023) – is inevitable (Dilmegani 2024). Industry leader Open AI states, "as our systems get closer to AGI, we are becoming increasingly cautious with the creation and deployment of our models", indicating movement towards AGI as an end goal (Open AI 2023). Teams building genAI systems require ethnographers' guidance, not only for applying genAI where it is useful, but also extending, rather than replacing, human cognition in ways that are *desirable* to people using the technology. This requirement is especially urgent, given genAI's exponential, worldwide growth and evolution (McKinsey and Company 2023).

One could counter-argue that new technology has always enabled some amount of *cognitive offloading* (Risko and Gilbert 2016), enabling *distributed cognition* from us to our technology tools (e.g., Hutchins and Klausen 1996). For instance, outsourcing navigation in a familiar city to Google Maps, or using your phone's calculator to calculate a tip at a restaurant. We posit that genAI is already shifting us into a new realm of cognitive offloading given its increasingly advanced capabilities.

Ethnographers have the potential to help teams navigate questions about how much of our executive function we wish to distribute to a genAI system, and under what circumstances.

It is also important to highlight that the data on which genAI models are built can reflect cultural, social, ableist, gender, racial, ethnic and/or economic trauma (Sampson 2023), harming those who use genAI tools. For instance, Open AI's DALL-E 2 image generator showed people of color when prompted for images of prisoners, or exclusively white people when prompted for images of CEOs (Johnson 2022). Humans are also susceptible to *automation bias*, in which we tend to favor information from automated systems – even to the point of ignoring correct information from non-automated sources (Sampson 2023). This is particularly problematic given genAI's emergent behavior and potential for hallucination.

Best practices have begun to emerge for mitigating genAI's harm and bias risks, such as organizations establishing, *a priori*, what content the model will not generate, devising harms modeling scenarios, implementing adversarial testing, and heightened model performance monitoring. These approaches are well-documented in both the EPIC community (e.g., Sampson 2023; De Paula et al. 2023) and the broader ML community (e.g., Gallegos et al. 2023), and are therefore not a focus of the current paper.

Designing GenAI Systems: Current Approaches and Challenges

Ethnography practitioners have already been engaging with the development of genAI systems. Research on this engagement provides evidence that designing for complex, emerging ML systems like genAI presents three unique challenges (Yang et al. 2020). First, familiar user-centered design (UCD) approaches (e.g., sketching, and subsequently gathering feedback on, a low fidelity prototype of a user interaction, or conducting a "Wizard of Oz" evaluation; Klemmer 2002) have limited applicability, given the emergent nature of the genAI models' behavior. Relatedly, teams typically begin the design process with a solution in mind or in hand. If a team has the resources to train a foundation model from scratch, they are still likely seeking to "ride the wave" of genAI and know they want to explore using this specific technology for an application. More likely, the team is seeking to customize (e.g., fine-tune) an existing foundation model, meaning that there is a fully-fledged solution (i.e., foundation model) already available that the team is building from.

Early promising approaches to address these challenges include the development of prototyping tools that require minimal coding expertise, allowing practitioners

across backgrounds to experiment with AI and data (e.g., Carney et al. 2020). Teams have also reported success with regular, rapid experimentation with the technology solution, including a tight collaboration loop between product teams and engineers (e.g., ML engineers, data scientists; Yildirim and Oh et al. 2023, Walter 2024).

This experimentation approach overlaps with the second challenge – namely, genAI requires adaptations to how cross-functional teams collaborate – specifically, how product teams (e.g., ethnography practitioners, designers, product managers) and engineers building and/or customizing the foundation models (e.g., ML engineers, research scientists) work together. Members across teams have different ways of knowing (Hoy et al. 2023), including specialized language and workflows. Disconnects can quickly become amplified given the speed at which genAI technology develops.

Promising approaches to remedy these disconnects include ethnography practitioners and other product team members sensitizing ML engineers to user needs (Zdanowska and Taylor 2022) – for instance, using visuals and other *boundary objects* to help align cross-functional partners (Lee 2007, Star and Friesemer 1989), and close cross-functional collaboration between ML engineers and the product team to identify system capabilities (Yang et al. 2020). Some examples include a list of synthesized system capabilities (Yildirim and Pushkarna et al. 2023; Yildirim and Oh et al. 2023) and regular, cross-functional experimentation with an in-development system (Walter 2024).

The third unique challenge is a central tenet of this paper: genAI solutions unlock new problem spaces, standing in contrast to the human-centered playbook of solving for existing problem spaces. We have seen symptoms of this foundational challenge documented in human-computer interaction literature on how teams design AI-based solutions, manifesting in struggles practitioners face early in the design process.

Specifically, practitioners report needing support in "*getting the right design*" (Yildirim and Pushkarna et al. 2023; Yildirim and Oh et al. 2023). This phrase references Bill Buxton's (2007) distinction between identifying that your design solution is solving a valuable problem for people, versus refining its usability after the solution and intended use case are identified (i.e., "*getting the design right"*). Existing AI design guidebooks (e.g., Google 2019, IBM 2022, Apple 2023) tend to focus on solution refinement, leaving practitioners seeking guidance on early-phase problem framing and reframing, and subsequent ideating on solutions that are both useful and technically feasible (Yildirim and Pushkarna et al. 2023; Yildirim and Oh et al. 2023). The following section examines the reasoning mechanisms behind this foundational challenge.

Foundational Challenge: Problem-first Versus Solution-first Framing

GenAI is an example of a *disruptive* technology entering the market, displacing established markets and generating new ones (Christensen and Bower 1995). A famous past example is when Apple introduced the iPhone in 2007. This technology solution displaced mobile phone incumbents like Blackberry and Nokia, but was also the necessary precursor for mobile-first use cases and related business models to emerge – for instance, sharing photos via social media (e.g., Instagram), ridesharing (e.g., Uber) and mobile payments (e.g., Venmo). In the same way, genAI is starting to displace technologies we considered mainstream even mere months ago, and create new ways of accomplishing tasks – for instance, moving from a search bar to a prompt bar.

Disruptive technology – be it the iPhone or genAI systems – are, by definition, *solutions* that are then applied to use cases, rather than solutions developed around people's needs. Per Paul Graham, co-founder of startup accelerator and venture capital firm Y Combinator, genAI is, "the exact opposite of a solution in search of a problem. It's the solution to far more problems than its developers even knew existed" (Graham 2023). In other words, the solution (genAI) unlocks a number of *problem spaces* – all of the aspects related to understanding and defining a given problem people face (Simon 1969; Newell and Simon 1972).

The dynamic of a solution (genAI) unlocking new problem spaces creates a tension between human-centered design practitioners and genAI. We argue that this tension is rooted in how we approach problem framing in genAI solution development. Dorst (2011) examined the reasoning patterns that experienced human-centered design practitioners use to frame design problems. The following equation typifies a successful design solution:

$$WHAT + HOW\ leads\ to\ VALUE$$

In this equation, "WHAT" is an object, service, or system. "HOW" is a known working principle that will help achieve the value one aspires to offer customers ("VALUE"). The "WHAT", in combination with the "HOW", should therefore yield the aspired "VALUE".

When presented with a complex design problem, experienced practitioners tend to encounter the following equation, in which the only "known" variable is the assumption about the value that they seek to achieve. The solution ("WHAT") and the working principle ("HOW") that achieves the value are unknown.

$$??? + ??? \; leads\; to\; VALUE$$

This requires abductive reasoning, working backwards from value. Here, experienced practitioners adopt a frame, shown in square brackets in the equation below. The frame implies that applying a certain working principle ("HOW") will lead to the aspired value. The frame leaves the solution ("WHAT") to be determined.

$$WHAT \; + \; [HOW\; leads\; to\; VALUE]$$

Before designing a solution, the practitioner tends to search for what Dorst calls the *central paradox* – the crux of what makes the problem difficult to solve. This is done through searching problem space (e.g., via ethnographic research, sense-making). Once the practitioner has a clearer view of the central paradox, they progress to the solution ("WHAT").

In contrast, disruptive technology like genAI presents by *beginning* with the solution ("WHAT"), per the equation below:

$$WHAT + ??? \; leads\; to\; ???$$

Taking a conventional framing approach in this context will leave the practitioner with two unknown variables ("HOW", "VALUE"), untethered from the solution ("WHAT"). This may explain what is observed when human-centered practitioners working on an AI-based project report user needs that the AI system cannot help address (Yildirim and Oh et al. 2023). The current approach is, "*What solution could unlock this value?*", whereas we need to reframe to, "*What value could this solution unlock?*"

Adding further complexity to this reframing is that "WHAT" is difficult to define with specificity, due to genAI's emergent properties. Teams must acknowledge this uncertainty, treating the potential problems that a genAI solution is uniquely positioned to solve as a hypothesis. That said, there is evidence that grounding a cross-disciplinary team in a general inventory of AI system capabilities at the start of an ideation session can help the group generate ideas for applications that are both useful and technically feasible (Yildirim and Oh et al. 2023). In this same study, a traditional UCD approach to brainstorming, which starts strictly from problem space, led to generating ideas that were constrained both in usefulness and technical feasibility. Even ideas that were viewed as high value to end users required

extremely high system accuracy (e.g., predicting sedation dose for ventilated patients in an intensive care unit, Yildirim and Oh et al. 2023).

This examination of problem framing mechanisms highlights the tension between the current human-centered playbook and the solution-first framing that disruptive technology dictates. The subsequent section proposes a new framework to help solve for this tension, and offer a path forward to align an understanding of a genAI system's capabilities (i.e., solution space) and the needs of potential stakeholders of this system (i.e., problem space).

New Foundations: The Problem-Solution Symbiosis Framework

We offer the Problem-Solution Symbiosis (PSS) framework to help model the interplay between problem and solution space when building genAI solutions (Figure 1). We have previously demonstrated the dominance of solution space when working with disruptive technology, and how this can disrupt application of traditional problem framing if we assume problem and solution space are expanding in sequence, rather than in parallel. The PSS framework models this co-existence of problem and solution space. The arrows represent touchpoints at which each space can symbiotically inform and evolve the other. For instance, taking inventory of system capabilities can influence how we frame problem space (i.e., "what value could this solution unlock?" is predicated by an understanding of what the solution can do), and how subsequent work to understand problem space (e.g., ethnography) can inform development of the genAI system.

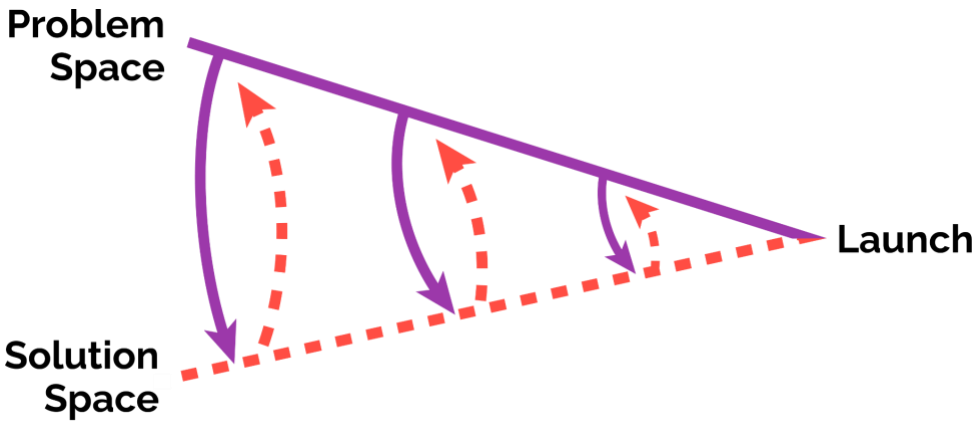

Figure 1. Problem-Solution Symbiosis (PSS) Framework. The solid purple line represents problem space. The dotted orange line represents solution space. The interplay of the arrows represents the symbiosis of problem and solution space as a team works towards launching a genAI solution. The arrows represent touchpoints at which each space can symbiotically inform and evolve the other. © Sendfull, LLC

As a function of time (i.e., going from left to right in Figure 1), the team developing the genAI solution can converge on a useful, desirable and feasible solution, mirroring traditional conceptualization of "controlled convergence" on a solution, such as those espoused by Pugh (1991) and Cross (1994). This convergence is typically marked by a product launch or other significant milestone. This paper focuses on early-stage framing and subsequent ideation, and therefore does not detail processes in later-stage or post-launch development. We encourage teams building genAI solutions to continue learning and iterating following product launch.

We posit that ethnography practitioners are well-situated to orchestrate the symbiosis of problem and solution space for genAI development. This orchestration requires the ethnographer's expertise to build coalitions between different perspectives between cross-functional team members (Hasbrouck, Scull and DiCarlo 2016), with the goal of developing a useful, desirable genAI solution that can think with us (rather than for us). We will outline three tools the ethnography practitioner can use to facilitate this orchestration, drawing from the author's lived experience in an analogous emerging technology space – zero-to-one spatial computing product development (Hutka 2021), systems theory, and human-computer interaction literature.

Tool 1: Build Intuition About GenAI System Capabilities

In the earlier examination of how teams currently approach genAI system design, we observed that regular, rapid experimentation with technology solutions, close collaboration between product and engineering teams, and practitioners experimenting with code-free prototyping tools were useful for understanding these emerging systems' capabilities. This process is analogous to *intuition building* that the author has applied when leading design research for zero-to-one spatial computing launches with large cross-functional teams (e.g., Hutka 2021).

Spatial computing describes a category of technologies in which computers understand people's contexts, such as reading hand gestures, body position, and voice, and can add digital elements into the physical environment that can be manipulated similar to real-world objects (Bar-Zeev 2023). Augmented reality (AR), which overlays digital content on the physical environment, often falls under this category. One example of a spatial computing technology for which the author led design research is Adobe Aero, an AR authoring application that enables creative professionals to build digital spatial experiences (Hutka 2021). We outline different activities (solo and with a cross-functional team) that the author has led in spatial computing design research, and offer genAI analogs that ethnography practitioners can lead.

Table 1. Intuition Building Activities Used in Spatial Computing Design Research, and Corresponding GenAI Analogs

Activity Participant(s)	Spatial Computing Activity Description	Benefit	GenAI Analog
Ethnography practitioner	Spend time using existing technology tools (e.g., computing headsets like Magic Leap and HoloLens; AR mobile applications), for tasks related to the team's product area.	Develop intuition about the AR medium's strengths and weaknesses in different form factors.	Spend time using foundation models, such as GPT-4o and LLaMA, and/or existing products built on these models. Enter prompts related to tasks in your team's product area.
	Try early versions of the product (i.e., "builds") shared by engineers.	Develop intuition about the in-development product's strengths and weaknesses; build shared language and rapport with engineers.	Partner with engineers to understand what they are building; interact with in-development systems as early as possible.
	Learn the basics of solution-adjacent tools. For example, learn the basics of Blender, a 3D authoring tool, which people frequently use to create content that will later be used in an AR experience built in a separate application.	Develop intuition about the product ecosystem into which AR fits.	If there is existing research on potential customers' workflows, spend time in primary tools that are likely to be used alongside, or obviated by, a potential genAI solution.
Ethnography practitioner and their cross-functional team	Human-centered playtesting sessions, in which cross-functional teams were invited to use the latest build, with specific goals and/or tasks, akin to a usability test. Outcomes and next steps are documented.	Foster cross-disciplinary dialogue and awareness around solution strengths and weaknesses. Document observations and next steps.	Conduct a playtest session where team members all try using their own prompts in a given generative AI system, as relevant to your team's product area.

Of these activities, cross-disciplinary sessions are uniquely valuable, due to their ability to organically build bridges across teams. For example, in such a session, one

team member outside of the engineering function says, "I noticed X was happening." An engineer in the same session responds, sharing why X may occur from a technical perspective, and takes an action item to investigate offline. This exchange not only moves development forward, but fosters exposure to different ways of knowing. Note, this activity should not be used as a substitute for adversarial testing, in which a team systematically and deliberately introduces inputs designed to test how the model behaves if exploited by bad actors.

Tool 2: Stakeholder Ecosystem Mapping to Understand Problem Space

Designing for a complex, emergent system (e.g., any genAI-based solution) involves numerous stakeholders beyond the end user, all of whom hold different values and incentives (De Paula et al. 2023). For example, if building a genAI tool for internal enterprise employees, the ecosystem may include the end user, their collaborators, managers, information technology administrators, customers of the enterprise company, and the team developing the genAI tools (e.g., the ethnography practitioner, ML engineer, legal practitioner, business leaders). When exploring what problems a given genAI solution can solve for people, taking this holistic view can help teams reflect on assumptions about intended audience and motivations for building the system, in addition to generating new opportunity areas.

In an analogous approach, the author has used such mapping to understand the dynamic nature of how designers in agencies collaborate on spatial computing projects. Primary research was conducted in the form of in-depth interviews, starting with end users (e.g., designers), to learn about existing goals, behaviors, workflows and collaborators.

Through snowball sampling, participants offered introductions to these collaborators when possible (e.g., creative directors, information technology managers, business leaders). The subsequent output of synthesis was a collaboration model, demonstrating how these stakeholders worked together, as well as the goals, values and pain points of each stakeholder. This stakeholder-sensitive approach provided the team with knowledge not only about the end user, but how upstream product adoption decisions were made. The outcome informed the go-to-market strategy, addressing upstream barriers to adoption faced by senior decision makers, as well as product requirements to meet the end users' needs. This approach adapts the service design practice of service blueprinting (Shostack 1982), in which a team maps the interactions between all stakeholders of a system. This differs from a UCD approach, which would more narrowly focus on the end user of the system.

The ethnography practitioner can take a similar approach on genAI projects. Following primary research to learn about the target audience, the ethnography

practitioner can map stakeholders, as well as their values and incentives. We recommend extending the aforementioned example to include the team building the genAI tool to promote reflexivity. For example, if a team is looking to fine-tune an existing foundation model to more effectively address customer service interactions for their company, they would consider: the customer (i.e., the end user; values: accuracy, efficiency; incentives: quickly resolve a problem); current customer service representatives (values: being able to focus on complex cases; incentives: skill development); business owners (values: cost reduction, scalability; incentives: increased profitability); internal product team (values: building useful tools; incentives: customer engagement and return use). This mapping exercise may reveal both opportunities for the product, but also expose risks (e.g., roles that genAI technology risks replacing) and new testable hypotheses for subsequent research. This activity also serves to bring interdisciplinary team members together, with the ethnography practitioner again serving as the "bridge" between disciplines.

Tool 3: The Cognitive Offloading Matrix to Identify Which – If Any – Tasks to Offload to GenAI

We previously posed the questions: how much of our executive function do we wish to delegate to a genAI system, and under what circumstances? To this end, we offer the Cognitive Offloading Matrix (Figure 2), which builds on human-computer interaction literature. The goal of this matrix is to help ethnography practitioners consider what cognitive tasks people may want to consider – or avoid – delegating to AI. While it can be applied within the context of current genAI solution development, it aims to serve as a durable tool as genAI capabilities increase and approach AGI.

The matrix has two axes: Unique strengths of humans and genAI systems (x-axis), and desirability of a given task (y-axis). The goal is to only leverage genAI systems for tasks that genAI is better equipped to solve for relative to people, *and* is something people want to offload.

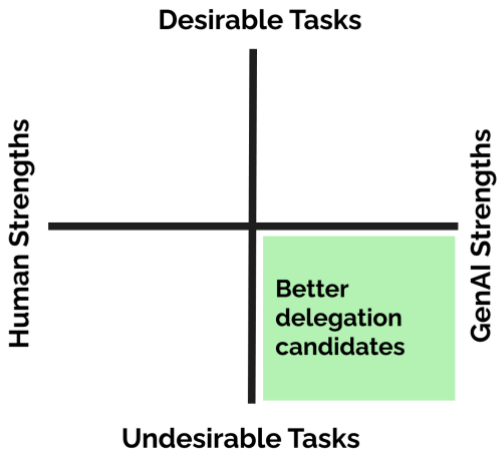

Figure 2. Cognitive Offloading Matrix. The matrix includes Strengths on the x-axis (human strengths at left, genAI strengths at right), and Tasks on the y-axis (desirable tasks at top, undesirable tasks at bottom). Tasks that fall in the bottom right quadrant (i.e., are undesirable and leverage genAI strengths) are better candidates to delegate to genAI (marked in green). © Sendfull, LLC.

To populate the x-axis, we will begin by examining a framework called the *Fitts List* (Fitts 1951), developed by psychologist and early human factors engineer, Paul Fitts. The Fitts List detailed the strengths of human versus machine capabilities. One example of human strengths is reacting to low-probability events, such as accidents (note how this activity is something at which machines, such as autonomous vehicles, remain poor). In contrast, machines excel at detecting stimuli beyond the capabilities of the human perceptual system, such as infrared detection (de Winter and Dodou 2011). Practitioners can refer to the People + AI Guidebook (Google 2019) for a list of "when AI is probably better" and "when AI is probably not better" as a starting point, though we caution that the lists are not specific to genAI technology solutions, with their properties of scale, homogenization and emergence.

Ethnography practitioners can adapt the Fitts List when working on genAI projects to identify potential areas where the technology can extend – rather than replace – human cognition. We can consider an example scenario in which a team is considering applying genAI to a music streaming tool. This activity can be done solo or in a workshop setting with a cross-functional team.

In this example, let us assume that a strength of a given genAI system was generating large-scale, real-time personalized recommendation. Let us also assume that a team has established (e.g., as informed by ethnographic research) that a human strength is making recommendations based on cultural nuances or non-obvious connections. These strengths can be mapped in a table, as shown in Table 2.

Table 2. Example of Fitts List for a genAI Music Streaming Tool

Human Strength(s)	GenAI Strength(s)
Making recommendations based on cultural nuances or non-obvious connections.	Generating large-scale, real-time personalized recommendations.

Populating the y-axis requires an understanding of a potential audience's workflows. If there is no evidence on the audience's workflows, this is cause for either engaging in primary research or treating tasks as hypotheses until otherwise investigated. This axis differentiates between desirable tasks and undesirable tasks to the potential audience. This concept is inspired by a survey investigating what activities on which people would like a robot's help (Li et al. 2022), as part of a larger effort to develop a human-centered benchmark for robots with AI capabilities (i.e., "embodied AI"). Respondents were asked, "*How much would you benefit if a robot did this for you?"* (ibid., p.3). The highest-ranked responses were: "wash floor", "clean bathroom" and "clean after a wild party." The lowest-ranked responses were: "buy a ring", "play squash" and "opening presents." These results supported that people sought to remain engaged in pleasurable and/or meaningful tasks, and only sought to outsource laborious tasks to robots, if given the opportunity.

We posit that these results offer relevant takeaways, when considering what problems to solve with a genAI system capable of increasingly complex "cognitive" capabilities. Considering the music streaming example. Let's recall that in this scenario, ethnographic research revealed that amongst people who used current streaming services, creating and sharing custom playlists for loved ones was an enjoyable and valuable activity. While we may have determined from our previous Fitts List that genAI is highly capable of recommendations for such a playlist, we can safely infer that this is not something people wish to outsource (Figure 3).

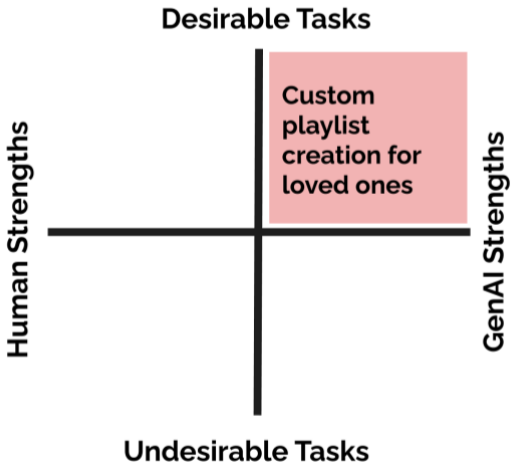

Figure 3. While 'Custom playlist creation for loved ones' is something genAI is highly capable of, it is also a desirable task for people. It therefore falls into the upper right quadrant of the Cognitive Offloading Matrix, and is not a good candidate for offloading to a genAI system, as indicated by the red square. © Sendfull, LLC.

Let us assume this same research study demonstrated that people want music to match their context, and that people are frustrated when the streaming service recommends the same low-energy music they want to hear when winding down in the evening, while embarking on an early-morning drive. Here, real-time generation plus contextual awareness could be something at which genAI excels. People are already open to algorithmically-generated playlists, and seek the added convenience of having music to match their context. This is a strong candidate for a "thinking" task to delegate to genAI (Figure 4).

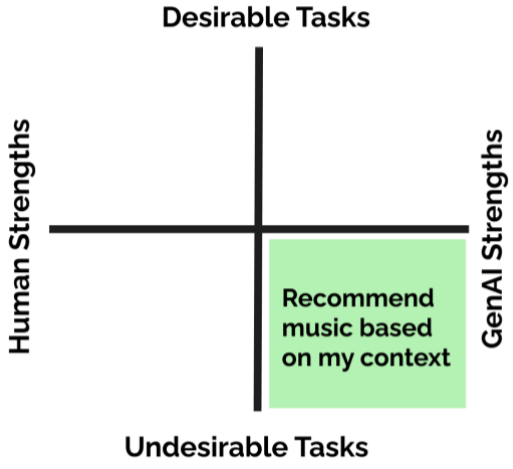

Figure 4. 'Recommend music based on my context' is something genAI is highly capable of, and a task people see as undesirable to perform themselves. It therefore falls into the lower right quadrant of the Cognitive Offloading Matrix, and is a good candidate for offloading to a genAI system, as indicated by the green square. © Sendfull, LLC.

General Discussion

The PSS framework proposes that ethnography practitioners move beyond a UCD approach by working *with* solution space early in the design process, and simultaneously, exploring problem space via a systemic approach to map stakeholders' values and incentives. This framework responds to observations that familiar UCD approaches often are of limited applicability in the development of genAI systems (e.g., Yildirim and Pushkarna et al. 2023; Yildirim and Oh et al. 2023), focused on how to identify and solve existing problems.

We posit that there are two sub-themes to tease apart regarding the limitations of a UCD approach. The first is that genAI is distinguished by unique characteristics (e.g., speed of advancement, complex, emergent behavior, coupled with being a disruptive technology) and therefore requires an updated approach (e.g., the PSS framework, and corresponding toolkit). The second is a more general observation that human-centered practitioners need to move beyond UCD – an observation that extends beyond the development of genAI systems. For example, Chesluk and Youngblood (2023) proposed a shift towards *user ecosystem thinking*, applying a systems-sensitive approach to industry ethnography to understand a broader range of human subjects and settings than UCD traditionally considers.

Forlizzi (2018) also advocated for moving beyond UCD to a more systemic approach. Forlizzi argues that we require stakeholder-centered design, accounting for "different entities interacting with and through products, services, and systems to achieve a desired outcome"; this is in contrast to designing "one thing for one person" (ibid., p. 1). The need for this shift beyond UCD approaches is a result of how computing technology has developed. In the earliest days of human-computer interaction, experts developed computer "programs" for themselves. Next came designing computers for others, as these machines became used in workplaces and homes. Consumer devices such as smartphones led to designing for entertainment and engagement, accelerating the development of user experience design. In the current age, we are building complex systems (genAI included) that require a systemic approach. The second tool proposed in this paper adopts this systemic view of problem space, grounded in this literature espousing a shift towards systematically building genAI solutions.

We recognize that systemic approaches can be met with friction when applied in industry settings. Product teams are often familiar with UCD language and methods, and therefore hesitant to adopt a broader view (e.g., Chesluk and Youngblood 2023). However, we are cautiously optimistic that the challenges of designing useful, desirable human-centered AI can accelerate the adoptions of systems-level thinking.

Relatedly, growing awareness of genAI "hype" (e.g., Chowdhury 2024, invoking the Gartner hype cycle [Gartner n.d.]), and the pitfalls of indiscriminately applying an "AI-first strategy" (Acar 2024) may increase business leaders' appetite for greater specificity around which problems are best suited for a genAI solution. The ethnography practitioner remains in a unique position to move human-centered AI forward given both their toolkits and mindsets, including the ability to build bridges between disciplines and perspectives.

Conclusion

Before *generating* new frameworks to guide human-centered development of genAI solutions, it is important to examine our *foundations* – for example, how we frame problems as human-centered practitioners, and identify foundational challenges associated with genAI development. Through this examination, we identified a foundational challenge, namely that current human-centered approaches focus on solving for existing problem spaces, and do not fully address the foundational challenge genAI presents – namely, a solution unlocking new problem spaces.

To address this foundational challenge, this paper proposed the PSS framework to reframe the interplay between problem and solution space when building genAI solutions. Relatedly, we shared three corresponding tools that the ethnography practitioner is well-equipped to apply. These tools included approaches to building intuition about solution space (i.e., genAI system capabilities) and stakeholder ecosystem mapping to understand problem space, as well as the Cognitive Offloading Matrix, to help identify if and when to offload tasks to genAI. We posit this Matrix can be used as genAI systems become increasingly capable of advanced cognitive tasks. We invite ethnography practitioners to apply the framework and toolkit offered in this paper, and adapt them to their needs. There are also opportunities to explore the application of this framework and tools to other emerging technology applications beyond genAI, such as solutions built on robotics (e.g., embodied AI) and computer vision technology.

About the Author

Stefanie Hutka (stef@sendfull.com) is the founder and Head of Design Research at Sendfull. She previously led research for 0-to-1 spatial computing product launches at Adobe and Meta, and established the design research function at AR start-up, DAQRI. She holds a PhD in Psychology from the University of Toronto.

Notes

Thank you to the anonymous reviewers of this paper submission for their valuable feedback, and Juliana Saldarriaga for their support and guidance as paper curator.

1. Scale, homogenization and emergence can be traced back to two techniques proposed in the seminal ML paper, Attention is All You Need (2017). These techniques were self-supervised learning, which enabled the ingestion of billions of data sources (e.g., documents) and multi-head self-attention, in which a ML model selectively chooses which input to pay attention to, rather than attending to each input equally (De Paula et al. 2023). These techniques scaffolded future genAI solutions like ChatGPT, and enabled scale, homogenization and emergence.

References Cited

Acar, Oguz A. 2024. "Is Your AI-First Strategy Causing More Problems Than It's Solving?" *Harvard Business Review*. Accessed May May 7, 2024. https://hbr.org/2024/03/is-your-ai-first-strategy-causing-more-problems-than-its-solving

Apple. 2023. "Human Interface Guidelines: Machine Learning." *Apple Developer Documentation.* Accessed May 10, 2024. https://developer.apple.com/design/human-interface-guidelines/machine-learning

Bar-Zeev, Avi. 2023. "What is Spatial Computing?" *Medium.* Accessed May 21, 2024. https://avibarzeev.medium.com/what-is-spatial-computing-48d5da22b09d

Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, et al. 2021. "On the Opportunities and Risks of Foundation Models." *arXiv* preprint, arXiv:2108.07258.

Bower, Joseph L., and Clayton M. Christensen. 1995. "Disruptive Technologies: Catching the Wave." *Harvard Business Review*, 43-53.

Buxton, Bill. 2007. *Sketching User Experiences: Getting the Design Right and the Right Design*. Amsterdam: Elsevier.

Carney, Michelle, Barron Webster, Irene Alvarado, Kyle Phillips, Noura Howell, Jordan Griffith, Jonas Jongejan, Amit Pitaru, and Alexander Chen. 2020. "Teachable Machine: Approachable Web-Based Tool for Exploring Machine Learning Classification." *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 1-8.

Chesluk, Benjamin, and Mike Youngblood. 2023. "Toward an Ethnography of Friction and Ease in Complex Systems." *2023 EPIC Proceedings*, 517-537. ISSN 1559-8918.

Chowdhury, H. 2024. "It Looks Like It Could be the End of the AI Hype Cycle." *Business Insider.* Accessed May 20, 2024. https://www.businessinsider.com/ai-leaders-worry-that-their-industry-is-based-on-hype-2024-4

Cross, N. 1994. *Engineering Design Methods – Strategies for Product Design.* Chichester, UK: Wiley.

De Paula, Rogério, Britta Fiore-Gartland, Jofish Kaye, and Tamara Kneese. 2023. "Foundational AI: Navigating a Shifting Terrain." *EPIC Learning and Networking Week 2023.* September 18, 2023. https://www.epicpeople.org/foundational-ai/

De Winter, Joost C. F., and Dimitra Dodou. 2014. "Why the Fitts List has Persisted Throughout the History of Function Allocation." *Cognition, Technology and Work,* 16: 1-11.

Dilmegani, Cem. 2024. "When Will Singularity Happen? 1700 Expert Opinions of AGI 2024." *AI Multiple Research.* Accessed May 10, 2024. https://research.aimultiple.com/artificial-general-intelligence-singularity-timing/

Dorst, Kees. 2011. "The Core of 'Design Thinking' and Its Application." *Design Studies* 32, no. 6: 521-532.

Fitts, Paul M. 1951. *Human Engineering for an Effective Air Navigation and Traffic Control System*. Washington, DC: National Research Council.

Forlizzi, Jodi. 2018. "Moving Beyond User-centered Design." *Interactions* 25, no. 5: 22–23.

Gallegos, Isabel O., Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2023. "Bias and Fairness in Large Language Models: A Survey." Computational Linguistics 50, no. 3. *arXiv:2309.00770v3*

Gartner. (n.d.) "Gartner Hype Cycle." Accessed May 21, 2024. https://www.gartner.com/en/research/methodologies/gartner-hype-cycle

Google. 2019. "People + AI Guidebook." Accessed May 7, 2024. pair.withgoogle.com/guidebook

Google. 2024. "What is Machine Learning?" Accessed May 12, 2024. https://developers.google.com/machine-learning/intro-to-ml/what-is-ml

Graham, Paul (@paulg). 2023. "AI is the exact opposite of a solution in search of a problem. It's the solution to far more problems than its developers even knew existed." *X,* Accessed May 7, 2024. https://x.com/paulg/status/1689874390442561536

Hasbrouck, Jay, Charley Scull, and Lisa DiCarlo. 2016. "Ethnographic Thinking for Wicked Problems: Framing Systemic Challenges and Catalyzing Change." *2016 EPIC Proceedings*, 560-561. https://www.epicpeople.org/ethnographic-thinking-wicked-problems/

Hoy, Tom, Iman Munire Bilal, and Zoe Liou. 2023. "Grounded Models: The Future of Sensemaking in a World of Generative AI." *2023 EPIC Proceedings*, 177-200. https://www.epicpeople.org/grounded-models-future-of-sensemaking-and-generative-ai/

Hu, Krystal. "ChatGPT Sets Record for Fastest-growing User Base – Analyst Note." *Reuters*. Accessed May 15, 2023. https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/

Hutka, Stefanie. 2021. "Building for the Future, Together: A Model for Bringing Emerging Products to Market, Using Anticipatory Ethnography and Mixed Methods Research." *2021 EPIC Proceedings*, 59-74. ISSN 1559-8918.

Hutchins, Edwin, and Klausen, Tove. 1996. "Distributed Cognition in an Airline Cockpit. *Cognition and Communication at Work*, 15-34.

IBM. 2022. "Design for AI." Accessed May 15, 2024. https://www.ibm.com/design/ai/

Johnson, Khari. 2022. "DALL-E 2 Creates Incredible Images – and Biased Ones You Don't See." *WIRED*. Accessed May 20, 2024. https://www.wired.com/story/dall-e-2-ai-text-image-bias-social-media/

Klein, Ezra. "How Should I Be Using A.I. Right Now?." *The Ezra Klein Show*. April 2, 2024. https://podcasts.apple.com/us/podcast/how-should-i-be-using-a-i-right-now/id1548604447?i=1000651164959

Klemmer, Scott R., Anoop K. Sinha, Jack Chen, James A. Landay, Nadeem Aboobaker, and Annie Wang. 2000. "Suede: A Wizard of Oz Prototyping Tool for Speech User Interfaces." *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology (UIST '00)*, 1–10. https://doi.org/10.1145/354401.354406

Lee, Charlotte P. 2007. "Boundary Negotiating Artifacts: Unbinding the Routine of Boundary Objects and Embracing Chaos in Collaborative Work." *Computer Supported Cooperative Work (CSCW)* 16, no. 3: 307–339.

Li, Chengshu, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang et al. 2024. "Behavior-1k: A Human-centered, Embodied AI Benchmark with 1,000 Everyday Activities and Realistic Simulation." *arXiv* preprint*,* arXiv:2403.09227.

McKinsey and Company. 2023. "The State of AI in 2023: Generative AI's Breakout Year." Accessed February 11, 2024. https://www.mckinsey.com/~/media/mckinsey/business%20functions/quantumblack/our%20insights/the%20state%20of%20ai%20in%202023%20generative%20ais%20breakout%20year/the-state-of-ai-in-2023-generative-ais-breakout-year_vf.pdf

MIT Technology Review Insights. 2024. "Generative AI: Differentiating Disruptors from the Disrupted." Accessed May 27, 2024. www.technologyreview.com/2024/02/29/1089152/generative-ai-differentiating-disruptors-from-the-disrupted

Mollick, Ethan (@emollick). 2023. "Or a cutting-edge anything-and-AI expert. Spend 10 hours just playing with a Frontier Model of your choice (ChatGPT Plus if you want to spend cash, Bing Creative Mode or Claude 2 otherwise) and you will learn to do something no-one else in the world knows." X, Accessed July 31, 2024. https://x.com/emollick/status/1717741987846037600

Murphy, Mike. 2022. "Design for AI." Accessed June 26, 2024. https://research.ibm.com/blog/what-are-foundation-models

Newell, Allen, and Herbert Alexander Simon. 1972. *Human Problem Solving*. Prentice-Hall: Englewood Cliffs, NJ.

Open AI. 2023. "Planning for AGI and Beyond." Accessed May 14, 2024. https://openai.com/index/planning-for-agi-and-beyond/

Open AI. 2024a. "ChatGPT – Release Notes." Accessed August 5, 2024. https://help.openai.com/en/articles/6825453-chatgpt-release-notes

Open AI. 2024b. "Introducing GPT-4o and More Tools to ChatGPT Free Users." Accessed May 13, 2024. https://openai.com/index/gpt-4o-and-more-tools-to-chatgpt-free/

Pugh, Stuart. 1991. *Total Design: Integrated Methods for Successful Product Engineering.* Wokingham, UK: Addison-Wesley Publishing Company.

Risko, Evan F., and Sam J. Gilbert. 2016. "Cognitive Offloading." *Trends in Cognitive Sciences* 20, no. 9: 676-688.

Sampson, Ovetta. 2023. "Pulp Friction: Creating Space for Cultural Context, Values, and the Quirkiness of Humanity in AI Engagement." *2023 EPIC Keynote.* https://www.epicpeople.org/pulp-friction-creating-space-for-cultural-context-values-humanity-in-ai/

Shostack, G. Lynn. 1982. "How to Design a Service." *European Journal of Marketing* 16, no. 1: 49-63.

Simon, Herbert Alexander. 1969. *The Sciences of the Artificial.* The MIT Press: Cambridge, MA.

Star, Susan Leigh, and James R Griesemer. 1989. "Institutional Ecology, Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology." *Social Studies of Science* 19, no. 3: 387–420.

Stivers, Tanya, Nicholas J. Enfield, Penelope Brown, Christina Englert, Makoto Hayashi, Trine Heinemann, Gertie Hoymann et al. 2009. "Universals and Cultural Variation in Turn-taking in Conversation." *Proceedings of the National Academy of Sciences* 106, no. 26: 10587-10592.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. "Attention is All You Need." *Advances in Neural Information Processing Systems*, 1-15.

Walter, Aaron. 2024. "Adobe Design Talks on AI and the Creative Process." Design Better Podcast (note: Podcast was a recording of an in-person panel, which the author attended and asked a question of the panelists during the Question and Answer period). Accessed June 26, 2024. https://designbetterpodcast.com/p/bonus-ai-and-the-creative-process

Yang, Qian, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. "Re-examining Whether, Why, and How Human-AI Interaction is Uniquely Difficult to Design." In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1-13.

Yildirim, Nur, Mahima Pushkarna, Nitesh Goyal, Martin Wattenberg, and Fernanda Viégas. 2023. "Investigating How Practitioners Use Human-AI Guidelines: A Case Study on the People + AI Guidebook." In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1-13.

Yildirim, Nur, Changhoon Oh, Deniz Sayar, Kayla Brand, Supritha Challa, Violet Turri, Nina Crosby Walton et al. 2023. "Creating Design Resources to Scaffold the Ideation of AI Concepts." In *Proceedings of the 2023 ACM Designing Interactive Systems Conference,* 2326-2346.

Displacing the Mind to the Body: Combining Methodologies to Address Socio-Ecological Crises

NORA MORALES, UAM *Cuajimalpa, México*
BLANCA MIEDES, *Huelva University, UHU, Spain*

This paper integrates arts-based methodologies into research and design, bringing new ways of knowing to the complex socio-ecological challenges we face today. New findings demonstrating that cognition is deeply rooted in bodily and social interactions highlight the need to shift our focus from the mind to the body. The authors provide a conceptual foundation and detailed instructions for practicing Social Presencing – which focuses on social dynamics – and Warm Data, which focuses on interrelationships within systems. Outcomes for participants included a clear evolution from stagnation and conflict to greater connection and potential for collective action, and the ability to face complex problems with a more collaborative perspective.

1. Introduction

This paper explores the combination of two methodologies—Social Presencing Theater (SPT) and Warm Data—to understand complex social systems and challenge traditional ways of knowing in ethnography and design. The combined methodology invites us to shift from a "systems thinking" perspective to a "systems sensing" experience to prepare for addressing global socio-ecological challenges. There are several reasons why combining these methodologies is possible and valuable: first, both argue that the foundation of successful activity in the social field arises from an engagement with the world and the back-and-forth interactions between the expert's "know-how" and the people's "know-what," sketching out a collective embodied approach. Second, they emphasize that meaningful action requires a deep change in perception as well as stronger social bonds among the actors involved; the dynamic relationship that emerges, as Shove, Pantzar, and Watson (2012) argue, reconstitutes the meaning of social practice.

In the face of global socio-ecological crises, traditional methodologies have proven inadequate in addressing the complexity of "wicked problems" (Buchanan, 1992). These challenges demand inter-systemic changes rather than sectorial solutions. Modern approaches to addressing these problems often lead to repeated errors or simplification of the necessary complexity, usually acting as corrections of behaviors instead of questioning what has caused human perception changes. Thus, there is a growing need to shift from linear strategies to more holistic, embodied, and relational approaches. As ethnographers and designers, we need to uncover new

Morales, Nora, and Blanca Miedes. "Displacing the Mind to the Body: Combining Methodologies to Address Socio-Ecological Crises." EPIC Proceedings (2024): 133–145. https://www.epicpeople.org/displacing-the-mind-to-the-body. ISSN 1559-8918.

strategies to collectively act based on a shared, deep, and first-hand understanding of the world, while shortening the gap between non-expert and expert thought—enabling connections between clients, researchers, and the external world (Roberts & Hoy, 2015).

At the heart of the expert ethnographer, there is always a practical expression that emerges from a set of techniques, theories, and heuristics that provide sophisticated cultural analysis and interpretation, which is later shared through communication tools, acts, and artifacts designed to influence, inform, and impact multiple audiences (Arnal & Holguin, 2007).

In this context, SPT, an embodied social arts-based methodology developed by the American-Japanese dancer Arawana Hayashi (Hayashi & Scharmer, 2021), has been used in multiple sectors and organizations to support processes of change. This approach is based on her extensive experience in oriental embodied practices across various social contexts combined with insights from Otto Scharmer's "Theory U" principles (Scharmer, 2018) at the Massachusetts Institute of Technology (MIT) which explores system thinking-sensing, organizational learning, and action research. SPT can be seen as an embodied collaborative tool that supports the possibility of achieving significant change and fostering connections among participants by applying deep listening practices, activating bodily wisdom to access the sensing world and transcending their narratives to described reality (Littlewood & Roche, 2004).

It is important to understand the meaning of the word 'theater' in SPT practice: it derives from the Greek *theatron*, a 'place of viewing,' and from *thesthai*, 'to behold.' It is essentially about making something visible, which takes us to the visual design field expertise, linking aesthetic and composition principles that guide the analysis to describe relationships created between the parts (shapes, materials, gestures, and bodies of individuals) and the whole (social body) to communicate and understand the essential quality of the practice.

In her book, Nora Bateson (2023) emphasizes that to improve our interaction with the world, we must first enhance our perception of the complexity we live within. Her concept of Warm Data, the other methodology we wish to combine, seeks to maintain the integrity of relational contexts within systems, providing a richer, more nuanced understanding of complex phenomena. Warm Data conversations capture qualitative dynamics and offer insights that quantitative data alone cannot provide, emphasizing the importance of interrelationships and contextual interdependencies, which are crucial for effectively addressing systemic challenges within living systems.

In this paper, we will outline how these methodologies, when applied across disciplines like dance, theater, architecture, and poetry, generate a sensory understanding and foster inter-systemic changes for various collectives (Hayashi & Scharmer, 2021). In the next section, we describe our combined methodology from a conceptual perspective. Afterwards, we provide detailed instructions for practitioners to apply it in their respective contexts. We also describe a workshop exercise in which the methodology was successfully applied. Finally, we discuss the workshop's results and provide general conclusions about efforts to combine methodologies in our practice.

2. Conceptual Foundations of the Combined Methodology

To understand SPT, we need to go back to the enactive approach to cognition. SPT proposes that cognition is deeply rooted in the dynamic interactions between the organism and its environment (Di Paolo, 2021; Di Paolo & De Jaegher, 2012, 2022; Sepúlveda-Pedro, 2024). It emphasizes autonomy, sense-making, emergence, embodiment, and experience. This approach holds that cognition emerges from the adaptive preservation of a dynamic network of sensorimotor structures sustained by continuous interactions with the environment and the body. By integrating the enactive cognition approach into ethnographic methods, researchers can recognize that understanding social and cultural dynamics requires more than observation and verbal analysis. It is essential to acknowledge that many cultures have long recognized the central role of the body in social dynamics—a recognition that Western thought has often overlooked. With SPT, we are now seeking to reclaim and integrate this understanding into our current context. The more we dive into these types of practices by exploring and expressing social experiences through embodied practices, the more genuine and powerful they become.

As a non-verbal arts-based methodology, SPT uses embodied practices to explore social dynamics and systemic challenges. It facilitates the perception and manifestation of social relationships and dynamics through body movements and group presence. Among the tools used to bring transformative innovation and change to organizations and institutions, SPT creates a new language that allows participants to shift their interaction from rational to emotional and express blockages (Stucks) and transitions (Shifts) within systems, fostering a deeper understanding of social fields and their dynamics. We emphasize the importance of embodied presence, movement, and collective inquiry drawn by this technique. When focusing on bodily sensations and the experience of being part of a dynamic social body, participants are able to infer collective responses and envision future

possibilities of any given situation. This set of embodied practices transcends language barriers and activates individual and collective consciousness, making it a powerful tool to address socio-ecological challenges.

On the other hand, Warm Data is defined as information derived from deep conversations about the interrelationships that connect the elements of a complex system. This type of data is trans-contextual because it captures the qualitative dynamics and relational patterns within a living system. Warm Data contrasts with traditional "cold" data by maintaining the context and interplay of system elements, which is essential for understanding and responding to complex challenges. By focusing on the relational and dynamic aspects of systems, it provides a more comprehensive and coherent understanding of the phenomena being studied (Bateson, 2023).

The Warm Data approach also integrates other domains in sciences, arts, and professional knowledge to create a qualitative inquiry into the integration of life. It challenges the mechanical reductionist view of environmental and social issues, promoting a more interconnected and relational perspective. This approach is particularly relevant for ethnographic research as it aligns with the enactive approach to cognition and resonance theory, both of which emphasize the importance of relationships and embodied experience in understanding complex systems.

We noticed that these two methodologies share a feature that can be best described with the concept of resonance. Resonance has a subjective and objective dimension, as it does not refer to cognition alone; it is also something that is felt and shared between the viewer and the viewed (Hayashi & Scharmer, 2021). Resonance awakens a sense and appreciation for the fullness of space, as seen in Hartmut Rosa's (2019) work with resonance theory, which focuses on the quality of relationships and the mutual vibrant response between individuals and their environment. Resonance is characterized by a meaningful connection, unlike alienated relationships which are marked by disconnection and repulsion. Following Rosa's suggestion on how to counteract the alienation caused by modern society's constant acceleration and expansion, we believe that the combined techniques of SPT and Warm Data provide us with a robust methodology to understand and address the complexity of socio-ecological challenges.

The combination of a conversational approach (Warm Data) with an embodied social art (SPT) can be conceptually described by four essential qualities:

1. **Relationality and Dynamism**: As an enactive approach, SPT emphasizes that human cognition is deeply situated in sociocultural contexts and that interactions within these contexts are dynamic and co-constitutive. Both the

subject and the world it encounters emerge from dynamic relationships and interactions, not as a priori given entities. Using gestures, movements, and brief sentences to express oneself opens dimensions to the quality of connection one can achieve, which is profound but not textual in nature.

2. **Embodied Cognition**: An embodied set of practices such as SPT deals with the premise that cognition is not just a mental process but is embodied and manifested through physical interaction with the environment. Warm Data, on the other hand, is not just about what we say or express but what we inscribe in objects, artifacts, and matter. Our combination provides a theoretical basis to understand how resonant experiences are bodily lived and how physical interactions with the world contribute to a collective understanding of a complex situation.
3. **Relational Ontologies**: By proposing this combined methodology, we suggest a relational ontological view where the subject and the world are co-original and mutually shape each other through processual performance. This emphasis on the co-constitution of mind and environment shows how interactions and relationships are fundamental to the formation of experience and knowledge.
4. **Conditions for Prototyping**: Finally, we stress the importance of studying the social conditions that facilitate social change by providing analytical tools to investigate and explore how specific social and cultural structures enable or inhibit meaningful interactions. This provides a deeper understanding of the conditions that favor new ways of exploring and relating to today's challenges.

We believe this combined approach can help ethnographers promote a more holistic and resonant understanding of social phenomena as complex living systems, fostering deeper connections and insights. It provides them with tools to explore new dimensions of reality and generate innovative responses to complex challenges. It also offers a powerful way to address the pressing socio-ecological crises of our time, moving towards healthier and more connected futures.

3. Applying the Combined Methodology

Applying our combined methodology means engaging participants in a series of performative practices to cultivate embodied awareness and social presence, followed by a structured conversation to gather data. The journey begins with the preparation and setting of the space for the stage. The main objective is to create an environment that encourages embodied exploration and deep conversation. It is better to have a

spacious room bathed in natural light, where comfortable seating arrangements invite participants to relax and feel at home. A set of carefully designed SPT practices must be ready for use, and recording devices should be discreetly placed to capture the rich dialogues that will unfold, always with the participants' consent. This setting welcomes a diverse group of participants—professionals, academics, students, and community members—each bringing their unique perspectives and energies.

Each session begins with a simple introduction to practices of simple movements for body and space awareness to set the tone. These practices are designed to help participants connect with their bodies and each other, grounding them in the present moment and preparing them for the journey ahead. The facilitators explain the goals and exercises of the session and introduce the theme or socio-ecological challenge to be attended.

After that, one of the practices we consider most appropriate for experimenting in organizational contexts is 4D Mapping. This practice consists of defining the specific situation to be addressed and identifying the main actors (usually in a preparatory session prior to the main session; the number of actors should not exceed 8 and must always include the Earth, the most vulnerable actor, and the one representing the highest future possibility). Participants then select their roles and engage in simple collective embodied arrangements, where they create a physical sculpture that represents a 'stuck' place or situation in the current state of the system. Afterwards, they are asked to slowly transition to a second sculpture representing a true move towards an 'unstuck' state. While Sculpture 1 is being formed, participants, acting in their roles, need to come up with a phrase that resonates with their role and enunciate it while taking their place in the sculpture. As Sculpture 2 is created, the participants develop another phrase, adding layers of meaning to the embodied exploration.

The session then transitions to an embodied conversation. This shift from embodied exploration to verbal articulation is gentle, allowing participants a few moments of silent reflection. Participants in the sculpture, as well as the observers, are invited to engage in a structured conversation guided by three main arguments: "What have you seen?", "What have you felt?", and "What have you done with your body?". These prompts encourage participants to share their visual observations, emotional and bodily sensations, and physical actions and movements with a brief intervention instead of a long speech.

The exercise thus evolves from SPT enactment to Warm Data conversations, highlighting the importance of capturing relational and contextual information and registering them in various forms. One form of registration is through log diaries

with trigger questions given to participants, which are answered in written form and shared during or after the session. Another type of registration is the graphic recording of the situation by a visual scriber. We also have the session recorded with audio and hand microphones for the main commenters that are later used for transcription, provided we have permission. Open-ended questions delve deeper into the relational dynamics and interdependencies observed during the SPT practices.

To conclude participants gather in a circle, discussing key learnings and insights. They are encouraged to consider how these insights can transform their personal and professional perceptions. The facilitators evocate the main themes and relational patterns that emerged during the conversation, encouraging participants integrate their experiences and carry forward the wisdom gained from the session.

This integrated methodology creates a holistic approach to understanding and addressing complex socio-ecological challenges. It not only deepens participants' perception of the systems they inhabit but also fosters a sense of connection and possibility, essential for meaningful action in the face of global crises.

4. Applying the Combined Methodology—A Success Story

In June 2023, this combined methodology was applied at a workshop in the Global Business Anthropology Summit (GBAS) in Mexico City. This workshop brought together 36 participants from diverse backgrounds, including entrepreneurs, researchers, activists, business leaders, academics, and students. A brief summary of the experience can be found in (Bueno et al., 2023).

The session began with SPT practices. The facilitators guided participants through a 4D mapping exercise to depict the "stuck" state of consulting firms facing socio-ecological challenges and to envision a "better" future. Participant roles used for the session were: earth, indigenous girl, shareholders, employee consultant, citizen activist, public regulators, and media. Each participant embodied their role, creating physical sculptures that visually and kinesthetically represented the current and desired future state. As they held their poses, participants shared phrases that resonated with their roles, adding layers of meaning to the embodied exploration.

Several SPT practices include speaking, and the reflection and dialogues of each session are a vital part of the learning. The GBAS session could be analyzed from the perception of change in the emotional and mental state reflected by the phrases enunciated by the participants before starting sculpture 1 and after making sculpture 2. Those phrases, as well as place arrangements and gestures observed in the participants' bodies, convey evidence of the transition experienced before and after the exercise. Below is a table comparing the phrases said by each role in the session:

Table 1. Two Stages in Phrases Enunciated by Participants

Participant role	Before (sculpture 1)	After (sculpture 2)
Earth	"Keep going"	"I'm waking up"
Indigenous girl	"I can't say nothing"	"I need help"
The best future possibility	"Please"	"There is a way"
Shareholders	"I can't see"	"I want to share"
Employee consultancy	"Advise"	"Support"
Activist citizen	"Hang in here"	"Together"
Public regulators	"Let's listen"	"Pay attention"
Media	"I won't see"	"I listen to everyone and everything"

The table shows the eight participant roles; the second column shows the phrases enunciated before making sculpture 1, and the third column on the right shows their final ones.

Following the physical exploration, the session transitioned into an embodied conversation. Participants were given a few moments of silent reflection to shift from physical to verbal articulation. They then engaged in a structured conversation guided by prompts such as "What have you seen?", "What have you felt?", and "What have you done with your body?". These questions encouraged participants to share their visual observations, emotional and bodily sensations, and physical actions, helping to bridge the gap between their embodied experiences and verbal expressions.

The session concluded with a reflection and integration phase. Participants gathered in a circle, reflecting on the session and discussing key learnings and insights. They considered how these insights could inform their professional and personal practices. The facilitators summarized the main themes and relational patterns that emerged during the conversation, helping participants integrate their experiences and carry forward the wisdom gained from the session. The exchange evolved into a Warm Data discussion that captures relational and contextual information. Open-ended questions delved deeper into the relational dynamics and interdependencies observed during the SPT practices.

Figure 1. SPT ludic workshop session in June 2023 at GBAS Mexico City, facilitated by the authors.

This workshop and others we have conducted provide valuable insights into the potential of these methodologies to address complex socio-ecological challenges. The results demonstrate significant shifts in the emotional and mental state of participants, highlighting the effectiveness of integrating embodied practices with relational conversations.

The transition observed in the participants' expressions before and after the SPT session reveals a profound impact on their perception and engagement with their roles. For example, the *earth* role transitioned from a state of relentless perseverance ("Keep going") to a state of awakening and heightened awareness ("I'm waking up"). This shift signifies a deeper connection to the environment and an enhanced perception of the surrounding complexities, aligning with the group resonance and mutual responses between individuals and their environment.

Similarly, the *Indigenous girl* role moved from a sense of silence and inability to express ("I can't say nothing") to an acknowledgment of need and openness to seeking help ("I need help"). This change indicates greater vulnerability and willingness to connect, which are crucial for fostering genuine relationships and collaborative efforts, resonating with the enactive approach to cognition that emphasizes the importance of social interactions in sense-making (Di Paolo & De Jaegher, 2012).

The best future possibility role illustrated a transformation from uncertainty and supplication ("Please") to clarity and optimism ("There is a way"). This shift reflects a newfound confidence in envisioning and pursuing concrete solutions, a critical aspect of addressing systemic challenges as suggested by Bateson's concept of Warm Data, which seeks to maintain relational contexts and provide a nuanced understanding of complex phenomena (Bateson, 2023).

The shareholders who initially felt blocked and visionless ("I can't see") evolved towards a desire for openness and sharing ("I want to share"). This transition suggests a movement from isolation to collaboration, emphasizing the importance of inclusivity in tackling complex issues, which is consistent with the goals of SPT to foster embodied presence and collective inquiry (Hayashi & Scharmer, 2021).

The employee consultancy role demonstrated a shift from a distant advisory position ("Advise") to active support and participation ("Support"). This change highlights an increase in empathy and commitment essential for effective collaboration and problem-solving, reinforcing the principles of enactive cognition that underline the embodied and relational nature of understanding (Di Paolo, 2021).

The *citizen activist's* role transitioned from persistence in challenging circumstances ("Hang in here") to a sense of unity and collective effort ("Together"). This shift underscores the power of collaboration and the importance of working together to overcome difficulties, also resonating with Rosa's concept of resonance, which advocates for a collective response to the crises of modern society (Rosa, 2019).

Public regulators moved from a passive willingness to listen ("Let's listen") to a proactive call for conscious and attentive action ("Pay attention"). This evolution indicates a heightened sense of responsibility and alertness crucial for effective governance and regulation, which aligns with the theoretical emphasis on the dynamic and responsive nature of socio-ecological systems (Sepúlveda-Pedro, 2024).

Finally, the *Media* role showed a dramatic shift from refusal and resistance ("I won't see") to complete openness and willingness to listen to diverse perspectives ("I listen to everyone and everything"). This transition highlights the importance of inclusivity and understanding in media practices, essential for fostering a well-informed and empathetic society, which is a key tenet of Warm Data's relational approach (Bateson, 2023).

6. Conclusion

Overall, the SPT and subsequent Warm Data sessions facilitated significant emotional and mental transitions from states of blockage, resistance, and isolation to states of openness, collaboration, and awareness. Participants exhibited a clear

evolution from stagnation and conflict to a greater sense of connection, mutual support, and potential for collective action. These changes suggest a positive impact on individuals' ability to face complex problems with a more integrative and collaborative perspective. The combined method not only brings clarity to nonverbal work but also sheds light on how we can communicate better collectively.

However, despite the promising outcomes, the implementation of this combined methodology presents certain challenges and limitations. The embodied practices involved in SPT can evoke strong emotions, which may be difficult to manage in some settings. Facilitators must be equipped to handle these emotional responses and provide a supportive environment for participants. Additionally, the dynamic and emergent nature of SPT and Warm Data conversations can lead to unpredictable outcomes. While this unpredictability is inherent to the process, it requires flexibility and adaptability from both facilitators and participants. Furthermore, incorporating these embodied and relational methodologies into traditional ethnographic practices may require a paradigm shift for researchers. It challenges the conventional focus on observation and verbal analysis, demanding a more holistic and embodied approach to understanding social phenomena. Therefore, it is crucial that facilitators possess all the necessary skills and capabilities to effectively implement these processes.

It is not debatable that the wellbeing of our planet is currently threatened by many crises, most of them manmade (Scarano, 2024). It is thus justified and fundamental that current changes are addressed collectively and with more diverse understandings of these complexities. New ways to approach data and to understand our presence in the world can emerge when we revisit and combine already existing methodologies.

The integration of SPT and Warm Data is a step toward a transformative approach to understanding and addressing complex socio-ecological challenges. This combined methodology could begin to shift the focus from the mind to the body and effectively integrate the context to a better understanding of social presence in relation to the system. The results from the workshop demonstrate that participants from different backgrounds and languages can experience significant awareness of their emotional and mental states, fostering deeper connections and encouraging a more holistic perception of their roles within systems.

The theoretical underpinnings of this approach are robust, drawing on enactive cognition and relational understanding to create a comprehensive framework for ethnographic research and practice. These theories highlight the importance of embodied experiences and relational dynamics, which are crucial for capturing the complexity of socio-ecological systems and fostering meaningful change.

However, the successful implementation of this combined methodology requires facilitators to possess a range of skills and capabilities to manage the emotional intensity and unpredictability inherent in these processes. Flexibility, adaptability, and a supportive environment are essential to navigate the dynamic nature of SPT and Warm Data conversations and to integrate these methodologies into traditional ethnographic practices effectively.

Finally, future research should continue to explore the applications of this integrated methodology in various contexts, refining the approach and evaluating its effectiveness in fostering systemic change. By embracing embodied experience and relational understanding, this methodology not only advances the field of ethnography but also proposes innovative forms of activism and social change, moving towards healthier and more connected futures.

About the Authors

Nora Morales is an academic from UAM Cuajimalpa, an information designer with a PhD in Social Sciences. She is careful to maintain curiosity about "mapping" tools to explore different realities and co-creation in this complex world. nmorales@cua.uam.mx

Blanca Miedes is an economist and researcher professor from Huelva University in Spain. She likes to think outside the brain, using the feelings and movements of our bodies in the physical space to comprehend more deeply and imaginatively. miedes@uhu.es

Note

Acknowledgments: We would like to thank Juliana Saldarriaga, our curator, for her commitment and dedication to this piece and all participants in the workshops that let us explore new ways to better connect to the essence of our humanity.

References

Arnal, Luis, and Roberto Holguin. 2007. "Ethnography and Music: Disseminating Ethnographic Research inside Organizations." In Ethnographic Praxis in Industry Conference Proceedings 2007, pp. 201–213. https://epicpeople.org/ethnography-and-music-disseminating-ethnographic-research-inside-organizations/

Bateson, Nora. 2023. Combining. S.I.: Triarchy Press.

Buchanan, Richard. 1992. "Wicked Problems in Design Thinking." Design Issues 8 (2): 5. https://doi.org/10.2307/1511637.

Bueno, Carmen Patricia, Consuelo González, Lucía Laurent-Neva, Maximino Matus, Nora Morales, and Ivonne Ramírez. 2023. "Lessons from the 4th Global Business Anthropology Summit." Journal of Business Anthropology 12 (2): 238–71. https://doi.org/10.22439/jba.v12i2.7063.

Di Paolo, Ezequiel A. 2021. "Enactive Becoming." Phenomenology and the Cognitive Sciences 20 (5): 783–809. https://doi.org/10.1007/s11097-019-09654-1.

Di Paolo, Ezequiel, and Hanne De Jaegher. 2012. "The Interactive Brain Hypothesis." Frontiers in Human Neuroscience 6. https://doi.org/10.3389/fnhum.2012.00163.

Di Paolo, Ezequiel A., and Hanne De Jaegher. 2022. "Enactive Ethics: Difference Becoming Participation." Topoi 41 (2): 241–56. https://doi.org/10.1007/s11245-021-09766-x.

Hayashi, Arawana, and Otto Scharmer. 2021. Social Presencing Theater: The Art of Making a True Move. Cambridge, MA: PI Press (Presencing Institute).

Littlewood, William C., and Mary Alice Roche. 2004. Waking Up: The Work of Charlotte Selver. Bloomington, IN: AuthorHouse.

Roberts, Simon, and Tom Hoy. 2015. "Knowing That and Knowing How: Towards Embodied Strategy." Ethnographic Praxis in Industry Conference Proceedings ISSN 1559-8918. https://epicpeople.org/knowing-that-and-knowing-how-towards-embodied-strategy/

Rosa, Hartmut. 2019. Resonance: A Sociology of the Relationship to the World. Medford, MA: Polity Press.

Scarano, Fabio Rubio. 2024. Regenerative Dialogues for Sustainable Futures: Integrating Science, Arts, Spirituality, and Ancestral Knowledge for Planetary Wellbeing. Sustainable Development Goals Series. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-031-51841-6.

Scharmer, Claus Otto. 2018. The Essentials of Theory U: Core Principles and Applications. First edition. Oakland, CA: BK Berrett-Koehler Publishers Inc.

Scharmer, Claus Otto. 2016. Theory U: Leading from the Future as It Emerges. Berrett-Koehler Publishers.

Sepúlveda-Pedro, Miguel A. 2024. "Sense-Making in the Wild: The Historical and Ecological Depth of Enactive Processes of Life and Cognition." Adaptive Behavior 32 (2): 117–36. https://doi.org/10.1177/10597123231190153.

Shove, Elizabeth, Mika Pantzar, and Matt Watson. 2012. The Dynamics of Social Practice: Everyday Life and How It Changes. 1 Oliver's Yard, 55 City Road, London, EC1Y 1SP, United Kingdom: SAGE Publications Ltd. https://doi.org/10.4135/9781446250655.

Estimating the Return on Ethnographic Investment: Finding More Human(e) Ways to Measure Our Impact

KATHERINE METZO, *Lowe's Home Improvement and Elemental Research*

This paper argues that more appropriate measures are needed to illuminate the full value to organizations of investing in ethnography. The author proposes a framework for defining qualitative and quantitative variables for impacts that unfold across both short- and long-term horizons, and through direct and indirect outcomes that are financial, organizational, and strategic. The author also argues that these variables must be defined in ways that are meaningful to both the organizations we work for and the communities they interact with. A case study of a past project demonstrates the multidimensional return on ethnographic investment.

Introduction

A perennial challenge for ethnographic researchers is to demonstrate the impact of our work. Whether it's a bullet point on a resume, a case study in a portfolio, or a talking point in a budget request, we're asked to describe the value our work brings to the business bottom line. Numbers drive the business world and while it is wise to speak the language of business, I want to challenge us as ethnographers to apply the same thinking we use in doing our research to defining our value to organizations in a more robust way. The output of ethnographic research cannot be reduced to a neat and tidy equation that generates a number that moves up and to the right.

A few years ago, Chad Maxwell challenged EPIC members to stop justifying and defending our work: "I'm done with that. I want all of us to be done with that. I want us to sit in a position of the offense and unapologetically show our craft, our approach, yields true, measurable impact, and unapologetically move forward with that" (2021). This essay is my attempt to help all of us do just that. In the spirit of looking back to move our discipline forward, I offer a framework to help us view our work in terms of return on ethnographic investment without being limited by financial measures. Because my specific focus is ethnographic research, the framework outlined here may not apply in every case as not every research situation or business context calls for an ethnographic lens. Conversely, for readers who are not ethnographers, I offer this—if you are investing your human capital into understanding and solving human problems through multi-method approaches, incorporating observation and a mix of structured and semi-structured methodologies, what I outline here likely still applies.

Metzo, Katherine. "Estimating the Return on Ethnographic Investment: Finding More Human(e) Ways to Measure Our impact." *EPIC Proceedings* (2024): 146–162. https://www.epicpeople.org/estimating-return-on-ethnographic-investment. ISSN 1559-8918.

"Return on investment" (ROI) is a success narrative that uses quantitative measures which can be easily converted into financial impact—improved conversion rates, increased customer base, profit, decreased error rates, saved time. None of these results happen through magic. They are the result of planning, evaluation, and tough choices. They are the result of dedicated people driving towards a goal. When ethnographers are engaged, our work comes early in the process and our qualitative insights are cautioned as "directional," as if direction isn't precisely what is needed for developing strategic vision and setting goals. This ROI narrative is so powerful that we've convinced ourselves of the imprecision of the value we deliver. We often search for more precision in speaking about and calculating ROI than we need to in order to create our own success narrative of ethnographic impact.

The framework I offer is a narrative framework that doesn't preclude ethnographers from offering a quantitative financial measure of our impact. Instead, I push us to understand which measures matter in both the short- and the long-term, acknowledging that these measures may not be the same. To truly measure impact, we cannot over index on short-term impacts and I invite us to reflect on how our impact unfolds over time as more than the sum of its parts. I begin with a reflection on a 19-year-old study to ground us in ethnography as the heart of the framework. While ancient history in business terms, enough time has passed to see both the tangible and intangible effects of this work and illustrate a more robust range of longer-term impacts that ethnography can bring to organizations. Following the reflection, I offer some myth-busting around traditional ROI and dig into how ethnography adds value beyond the short-term by preparing organizations for change through attention to liminal spaces and relationships. Finally, I offer my framework for how we can estimate the value of our work, encompassing both immediate impact and long-term impact, as well as sustainable impact, which is the amplified value we bring through insights and recommendations that foster greater organizational resilience.

Reflection on Habitat Charlotte Communities

In Spring 2006, I led a consultation project with Habitat for Humanity of Charlotte (Habitat Charlotte) which initially sounded simple—how can Habitat Charlotte make townhomes appealing for homeowners (Metzo 2008)? Given that I was using the project to train a group of students, a comparative study of Habitat affiliates already building townhomes made sense. Through multiple rounds of discussion with stakeholders, we learned more about the problem this solution was meant to address. Land values were rising rapidly in Charlotte—infill lots cost 10

times in 2005 what they cost in 1983[1]—but Habitat Charlotte wanted to sustain, if not grow the number of homes they were building each year. The root problem for these stakeholders was whether they were truly fulfilling their mission if service remained flat, or worse, declined.

The biggest surprise for us was that the organization had not talked to their own homeowners, so we did. While one group continued with interviews and site visits with other Habitat affiliates, a second group supported me in facilitating workshops with homeowners. We broke the workshops into 4 sections. The first three sections started with quiet reflection on a prompt before sharing with the group while the fourth section revealed townhomes as a proposed alternative to detached, single-family residences and asked which items on each list were most important in that context. The three prompts invited participants to imagine the interior, exterior, and neighborhood for their "ideal home." We ended the session with open discussion of criteria the staff and board should keep in mind as they identified ways to sustain their mission of building homes.

While we "validated" that townhomes were "acceptable," we uncovered more desirable alternatives that also met the financial constraints that Habitat Charlotte was trying to overcome. The idea of a new, single-family detached home was ideal, but based on location and neighborhood, homeowners were willing to consider renovated homes. Both townhomes and renovated homes allowed for greater flexibility in location—closer to family or work, better neighborhood schools, amenities, etc. A root problem for homeowners with the Habitat model was the heavy responsibility they felt to serve as role models in their communities. On reviewing her options through Habitat, one homeowner stated, "I was really pissed—like, 'how dare you bring me into these areas that are beat down and run down where the city hasn't done anything to improve the community." Another homeowner shared a more measured position: "I feel like they are choosing people to put in those communities so that we can become a community and help others . . . *I mean they don't let just anyone into Habitat*" (emphasis added). Throughout these emotionally charged conversations, homeowners landed on the idea of building an all-Habitat community—a place where they could more easily turn to each other for moral and material support, feel reassurance that each of the other homeowners went through the same rigorous screening and demonstrated their financial stability and personal responsibility before their selection as a homeowner. In the end, our research insights lead to the development of multiple new programs: Critical Repairs (2007); Townhomes (2008); and Habitat communities (2012).

Figure 1. View at dawn of cul-de-sac on Jimmy and Rosalyn Carter Project build for Meadows at Plato Price community, October 2023, Charlotte, North Carolina. Photo by author.

Why Traditional ROI Doesn't Work

We are all familiar with the saying "If you can't count it, you can't measure it," and that sentiment is at the core of the concept of ROI. When we choose to assign value to research, as with any other discipline that ultimately supports bringing new products and services to the market, we make some assumptions in order to craft a "measurement" of the output of our work. For User Experience (UX), Forrester (cited in Kucheriavy 2017) boldly asserts a return of $100 for every dollar invested in UX. Establishing this kind of blanket frame for a field that employs a deep bench of ethnographers has not prevented disproportionate layoffs of UX researchers in the early 2020s. Because business runs in quarterly and annual cycles, the myths below are taken for granted as meaningful measures of impact, but if we rely solely on these measures blanket assertions about ROI, we short change ourselves.

I am by no means the first person to offer a critique of calculations of ROI applied to research and this is not meant to be an exhaustive critique. In fact, some ethnographic impact can fit nicely into the traditional ROI assumptions. This more "traditional" model can work for more evaluative types of research and other business functions such as marketing spend, because that work typically comes later in the process and has a more linear relationship to the product or service that is being stood up. Where this model falls apart is when we consider how much ethnographic research directly or indirectly drives strategy, which is not always expected to have an immediate impact on the organization, but should always provide clear direction for longer term efforts. This type of impact is more difficult

to measure with a blunt instrument like ROI. For the purposes of this essay, I highlight several key myths of the ROI model: value can always be translated into a financial measure, it becomes apparent relatively quickly, and has a linear relationship to outcomes.

Myth 1: Outcomes Can Be Measured in Dollars and Cents

Most often, we see impact translated into a direct and discrete financial outcome or some other measure that can be converted into dollars and cents. Financial outcomes might include positive revenue from increased conversion rates, for example, or a net savings (e.g., positive margin) from decreased operating or labor costs. Put more simply, we built the right thing and the company makes money or we prevent a costly mistake and save money before building the right thing (and then make money). In design research, "build the right thing" is a truism and while we can never mitigate all risk, great research can bring us closer to the "right" fit as well as building in the ability to pivot quickly with research-based insights. My objective here is not to say it's never worth the effort of quantifying, only that even when it might fit it is limiting when applied to ethnography.

All research brings financial, temporal, and trust "costs" while the "return" is generally only talked about in terms of a financial benefit. Such measures fail to capture the rich and multi-faceted outcomes of ethnographic engagements. In the Habitat example, rather than a financial measure, the Charlotte chapter uses number of families served (e.g., homes built or restored) as a direct impact measure. So, while I could choose 3 years as a generous short-term window and define the ROI our research as benefitting over 10 additional homeowners, that doesn't account for variability of costs for renovation versus new construction or the shifting volunteer needs for adding more 2-story buildings. Such a measure would also completely ignore the intangible shift in perspective around creating a nurturing environment for new homeowners to support each other and feel supported by Habitat instead of feeling mostly responsibility for their neighbors. Collecting and counting outputs can never truly add up the human impact that ethnographic research provides by making sense of complex social relations and structures.

Myth 2: Impact Can Be Seen Fairly Quickly

We are generally asked to measure impact on a very short time horizon, but ethnography tends to inhabit different and multiple time horizons. We may have replaced the "ethnographic present" with a more dynamic view of time, extending our research into "futures," which makes measuring our work against quarterly key

results even more challenging. To even provide a rough arithmetic value for the Habitat project, I gave myself three years. What these narrow time views of impact miss are delayed wins and sustainability vs. inadequacy of solutions.

Much of the work within design research is evaluative, fast, and ideally iterative. While this work can be suited to the conventional approach to ROI,[2] ethnographers delivering these kinds of studies will also tap into more fundamental insights around mental models, user values or archetypes, integrating ethnographic practice into more structured methods. Guth (2022) addressed head-on the challenge of "research amnesia," where insights are forgotten or teams are "ill-equipped" to leverage research insights beyond discrete, time-bound frames. What if we revisited products or services built off of our research several years later to evaluate the staying power of solutions? Just as the knowledge about the communities we study doesn't "expire" at the end of the study, the impact of our research on future strategic decisions should be celebrated and measured.

Perhaps more troubling with the emphasis on short-term wins is the lack of sustainability in solutions that provide instant returns. Alan Cooper, one of the fathers of user experience, spoke to the emphasis on speed as problematic. He compared tech production to having all the ingredients for making a cake and a specification of an amount of time, but no definition of done. "All we know is that it is on time, but its success will be a mystery" (1999, 42). Consider the front-end experience that gets high marks from users and results in a 30-fold increase in leads. Now imagine this "win" lacks any back-end support for adjusting internal processes, whether technical or labor-related, to convert these leads to sales. Without broader support, a short-term win can plateau or reverse itself.

Myth 3: Insights are Directly Correlated to Business Outcomes

The way ROI is typically calculated is by drawing a straight line from outputs (e.g., research recommendations, design changes) to outcomes (e.g., increased profit, improved customer satisfaction). One reason ethnographers are reluctant to measure impact is that we know how many people and decisions are involved in the journey to a finished product, service, or campaign after our contributions. Even apart from the external costs and influences in delivering solutions to market (in the case of Habitat building townhomes, this includes increased overhead from training and insurance costs), researchers are generally part of a broader group of experts working towards the final goal. By investing in the development of three separate programs in 2007, Habitat Charlotte addressed the immediate concern around land costs, but also built in flexibility that supported the organization during the housing crash in 2008

because they could leverage a wider range of available land and housing stock. Rather than shying away from talking about outcomes, we need to articulate how the "crooked" line of ethnographic insights actually helps control for external factors by identifying risks, challenges, and opportunities early in the journey so that organizations can plan for change.

The Value of Investing in Ethnography

More traditional measures of ROI are insufficient for measuring the value of ethnographic research because they are limited by the assumptions behind them. Understanding ethnographic impact begs the question of whether, in these traditional models, we are measuring what matters. What if, instead of buying into the mythology of the market, we make explicit the value embedded in ethnographic praxis? In the following section, I will translate these values into more explicit rules that we might use to measure ethnographic impact. In this section, I call out the key foundations of ethnography as a method, an output, and a way of thinking. It is also an enduring conversation, as this anniversary of the EPIC community and twenty years of writing on praxis has illustrated. With its deep history and growing appeal, it's worth understanding what makes ethnography the proverbial Swiss army knife for making sense out of complexity and finding underlying cultural patterns that shape key behaviors.

One of my favorite definitions of our craft comes from Zora Neale Hurston: "[Ethnographic] research is formalized curiosity. It is poking and prying with a purpose. It is a seeking that he[/she/they] who wishes may know the cosmic secrets of the world and they that dwell therein" (1996[1942], 143). Hurston highlights that ethnography is simultaneously systematic and unconstrained by prior assumptions or hypotheses. It is a mindset through which we "decipher the unique cultural logics of human complexity" (Hasbrouck 2018, 5). It is also a deeply ethical practice. Ethnographers do not "seek to manipulate others for 'scientific' ends," but focus instead on the shared humanity between themselves and those whom they study (Atkinson 2015, 5). The richness of ethnographic practice comes from its manifold nature, both technical and fluid, at once embedded and anticipatory.

Our "formalized curiosity" reveals itself in the way we uncover webs of significance and answer the unasked questions. This is because ethnography was built for articulating meaning from ambiguity and identifying the tensions and opportunities inherent in the liminal space that leads to transformation. As our discipline has evolved, these representations are no longer static snapshots of the "other" but dynamic, reflexive, intersectional, and futures-oriented narratives.

Foundation: Webs of Significance

More than a specific method or set of methods, ethnography illuminates "webs of significance" within communities. From mountains to board rooms, ethnographers set out to articulate relationships, meaning, and value from, "a multiplicity of complex conceptual structures, many of them superimposed upon or knotted into one another, which are at once strange, irregular, and inexplicit" (Geertz 1973, 10). Ethnography is a wide-angle lens that embraces context and complexity that our stakeholders, understandably, want to filter out when defining their needs. In the field of user experience, for example, context setting for stakeholders is a crucial skill. While the aim of the work that follows our research is to build a digital or physical product, that product is not used in a vacuum. Humans and the products they interact with live in spaces. Unpacking this context for business stakeholders is fundamental to how ethnography delivers value, even before a project starts.

Robin Beers et al. (2011) talk about kicking off a first of its kind ethnographic research effort at Wells Fargo, which more than doubled the time needed for what became routine practice. Beers' entry experience at Wells Fargo parallels my own in conducting a foundational study of a contact center. As the first researcher in my product space and the only ethnographer at the company at the time, I spent as much time socializing the project proposal and gaining buy-in from director-level leaders on the potential value of the research for a major tech project (replacing legacy software) as I did collecting the data across three sites. Ultimately, what won over stakeholders was the need to understand the workarounds that our employees created to bridge gaps in delivering results for our customers and document when and why those efforts may still not be enough. Such data couldn't be collected by replaying quality assurance recordings.

Even before entering the field, socializing our approach does two things. First, it highlights the way in which ethnography can help unpack the complexity of the business problem. "We begin to see what has previously been overlooked and perhaps then discover what is truly possible" (Cliver et al. 2010, 231; see also Beers et al. 2011). In my case, the questions I raised in the research plan tapped into the key measure that matters for the business—customer satisfaction—by highlighting the way in which customers are "secondary users" of the software and the pain points inherent within them[3]. It's worth noting that even when we're brought in at the tail end of a project, ethnographers can bridge cross-functional gaps with a broader systems view (Sih, et al. 2018).

Design literature often focuses on "systems," but I posit that the systems we describe are primarily an artifact of our own analysis. Leaning into Geertz's phrase

"knotted into one another," the part of the "web" that is most significant is the situated relationship. To understand these relationships, Desmond suggests that we study "fields rather than places, boundaries rather than bounded groups, processes rather than processed people, and cultural conflict rather than group culture" (Desmond 2014, 548). The webs within ethnographic praxis are woven between stakeholder groups and the humans being studied, as well as within each of these respective groups. In my case, the relationship with stakeholders was built around identifying those unspoken questions with the customer experience as a particular type of filed nested within a process. In the Habitat example, we took people outside the confines of their home and property line to consider what makes a place a home, which unlocked both tensions and opportunities presented by Habitat Charlotte's vetting process. Our challenge is to simplify complexity without flattening it, providing enough context to highlight meaningful connections.

Foundation: Time Is about More than Length of Engagement

Classically, we are taught that ethnography involves long stretches of time at a field site, building relationships over months and years. But as far back as the 1960s, participatory methods, action research, and other applied approaches set the stage for kinds of brief but impactful engagements that characterize ethnographic praxis. Our Habitat engagement lasted about 12 weeks. Rather than focus on the amount of time spent on a research engagement, a more productive way to think about time in relation to ethnography is bridging the present to short- and long-term futures. We carry forward or re-unite stakeholders with what is known (Guth 2022) allowing us to go deeper or understand how things have changed (Charland and Hoffman 2017). Time may be more relevant as a factor in our work if we focus instead on things like timing and foresight.

Foundation: Liminality and Transformation

The organizations we work with may or may not be going through a period of conscious change. The problems they present to us, however, might show that they are on the precipice, perhaps wanting to maintain the status quo in the face of external forces or seeking to understand how their target audience is adapting to their own changes. Liminality is not restricted to ritualized movement through life stages or planned transformation. During any period of nascent change, we find that "Liminal entities are neither here nor there; they are betwixt and between the positions assigned and arrayed by law, custom, convention, and ceremonial." (Turner 1969, 95). The ambiguity and indeterminacy that Turner speaks of open up a

multitude of seeds of transformation. Whether the implicit or explicit, these liminal contexts are ones where ethnography can amplify the value of our work by making sense of the inherent ambiguity in periods of change.

Liminal state research inside organizations largely involves making explicit the tacit knowledge of a group. We are stating aloud what is "taken for granted" among the groups we study and report to. Ethnographers working in enterprise teams will sometimes hear "we already know that" from stakeholders. Our challenge is not to complain to our colleagues about our work being ignored, but to tie that tacit knowledge to the field of what is possible. Some researchers, including myself, tap into humor in delivering our insights, while others use workshops or targeted engagements to immerse stakeholders into the lived experiences of "users" they already "know" while exposing bias (Reksodipoetro and McCarter nd; McClard and Dugan 2017). For example, researchers at Intel wanted to embark on strategic research after identifying two dozen separate projects related to the theme of "child's play." Their business stakeholders were on the cusp of making decisions about their customers based on cultural biases that were limiting to the business, but the workshop served as a liminal space where leaders could reveal and reflect their biases, opening the way for the necessary foundational research to take place (McClard and Dugan 2017). The timing of this workshop also helped the authors reveal to stakeholders the degree of ambiguity in the space. A key for ethnographers is to both recognize or create openings for liminality and leverage timing to execute our work to deliver insights ahead of key decisions.

Framework for Evaluating Ethnographic Impact

While ethnographic research in the business world can run the gamut from short-term and one-off engagements to longitudinal and iterative studies, each study can be foundational and holds the possibility of creating critical shifts in our collective thinking about a problem. Ethnography is a seed of pure potential. Pure potential also means that not every ethnographic project is going to have a long-term impact. Like a garden, not all the seeds we plant germinate and even those that grow into plants may not thrive. Whether or not an ethnographic project bears fruit, it is worth considering how to measure the harvest.

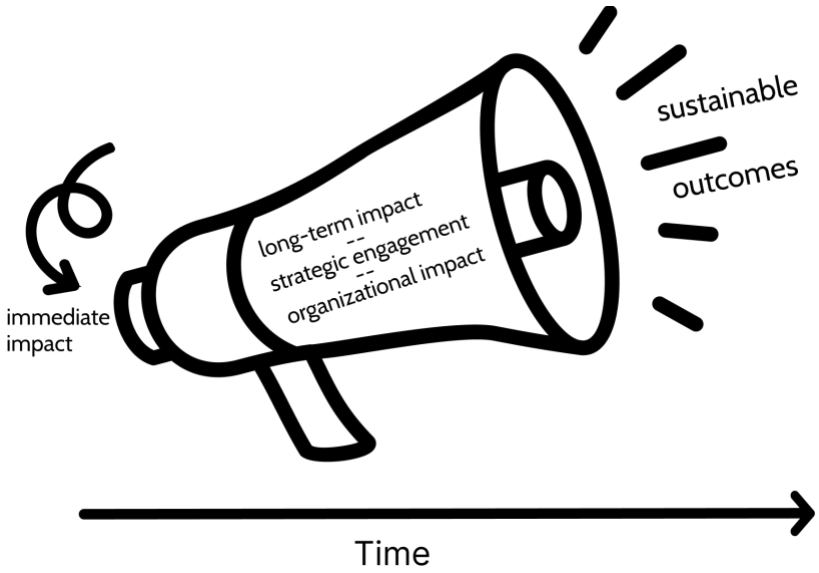

Figure 2. Visual model for understanding how ethnographic impact is amplified over time. Illustration credit: Frank DiTommaso

Determining a long-term impact of ethnography is, by necessity, a retrospective project. Measuring something that hasn't happened yet is a projection or a forecast but I want to simultaneously encourage us to plan for these elements at the outset of each new ethnographic project even as we reflect on past projects that we might measure now. At the core of this ethnographic impact framework is an immediate impact combined with a matrix of value adds that unfold over a longer time frame (see figure 2). While we might measure the cumulative impact of short- and long-term ROI using traditional measures, ethnographic praxis often amplifies these impacts into something that is beyond cumulative.

A seminal article on "scales of measurement" posits three steps to effective measurement (Stevens 1946). First, we must define the rules; this is, admittedly, the bulk of what this essay will accomplish. Next, the authors ask us to determine the mathematical properties of these scales and apply statistical operations to arrive at a final measure. I hope I will be forgiven for offering a more artistic than statistical treatment of ethnographic impact. If immediate and long-term ROI *can* be quantified using linear, short-term methods, we can use this as a partial measure. An overall measure of ethnographic impact might be the differential of the status quo subtracted from growth over time. In the Habitat example, this could be a differential of how many *more* homes were completed as a result of the three new programs implemented compared to the status quo. It could also be the sustainability of outcomes or the virality of an outcome that spreads to other organizations or sectors. In the Habitat example the townhome model would not be "viral" as it was

modeled on existing programs at other affiliates, while the Habitat Communities might be. Even in cases where a number can be assigned, we should not forget the context of the intangible value to the humans we study and design for. To break this down, I'll take each principle of the framework in turn (see Table 1), starting with the tangible elements.

Perhaps most obvious, ethnographic research in the world of business needs to address a compelling, immediate need. Regardless of any longer-term or longer-lasting impacts, we still need to deliver a tangible, measurable outcome for our stakeholders. It may not be easy to correlate outcomes to something like increased sales, but we may tie the delivery of robust, research-driven personas to key benchmarks that our product teams need to hit in identifying a target audience, feature set, and timeline. In one project, I worked with product and UX to flesh out a roadmap for a vaguely defined business "North Star" by providing a framing for pain points with processes and software according to pervasiveness and severity. This was a minor component of the overall project, but was a tangible way to identify "quick wins." My product owner added in a qualitative difficultly rating for engineering and we used this matrix to define a subset of pain points to solve for within our MVP product launch. By combining the vantage point of the long view with the reality of quarterly planning, we were able to prioritize which quick wins did the most to reduce friction while serving the longer-term organizational goals.

Table 1. Framework Principles for Evaluating Ethnographic Impact

Framework Principle	Description
Immediate impact	A direct outcome that can be measured using "standard" financial or other quantitative measures (e.g., increased revenue, conversion rates, clients served, improved CSAT or NPS).
Long-term impact	A direct outcome that may be measurable using a "standard" financial or quantitative measure, but is connected to future, rather than short-term, decisions made by the organization. It may be more reasonable to represent this impact with a qualitative narrative over a quantitate measure.

Organizational Impact	A direct or indirect outcome related to organizational change. Organizational impacts can be built into the business objectives for the research, or they may result from providing direction by answering unasked questions. While a quantitative measure may be possible, this is more likely to be a qualitative measure or narrative.
Strategic engagement	Engagement extends beyond transfer of final deliverables, such as developing product strategy or supporting program activation. Engagement can also be episodic and unplanned, such as workshops that identify gaps or reports that reduce "research amnesia." Strategic engagement is a binary (yes/no) variable.
Resiliency – Sustainability	Direct or indirect qualitative outcomes that represent the durability, longevity, or flexibility of both immediate and longer-term impacts over a broader time-frame.

Direct outcomes can happen across multiple time horizons. Even when an organization asks specifically for future-oriented research, our research insights are rarely aimed exclusively at a three- to five-year horizon. For such future oriented projects, we increase our impact by taking the opportunity to suggest quick wins or immediate first steps, like the mapping of employee pain points against severity and pervasiveness, described above. Likewise, projects that ostensibly focus on immediate needs may lend themselves to broader inquiries, such as McClard and Dugan's work on "child's play" (2017), or to multiple solutions across a longer time-horizon, as in the Habitat Charlotte reflection I began this paper with, where the immediate quick win was to leverage existing crews to start up the Critical Home Repair program, the anticipated short-term impact was townhomes, and the long-term impact was Habitat communities. In addition to the direct outcomes whose value can be approximated through more traditional ROI calculations, ethnography brings tangible (if not strictly measurable) and intangible values that can be talked about in comparative terms, such as the value experience of Habitat Charlotte homeowners who were afforded a wider array of choices to fit their desired lifestyle. Measures of impact that are the most meaningful can always be defined with precision, even if they cannot be quantified.

Organization change, like the opportunities to offer long-term recommendations, is not always part of the "ask." My corporate ethnography project with the pervasive and severe pain points was framed as a "lift and shift" of a primary tool, with

opportunities for "value add." In the beginning (and perhaps even now), our business partners on the project would not have said they were "betwixt and between"—they had a North Star vision. What marked that project as "liminal" was my manager seeing value in using ethnography as a route to identify both shorter and longer-term needs and site directors buying into the vision I set forth. The key to this principle is that ethnographers can leverage the liminality they've identified in an organization to drive new ways of thinking. Organizational shifts might be as small as hearing a term used in your research report be adopted by members of the C-suite. With the opening Habitat example, the organization was clearly at an inflection point and by digging into the meaning of "home" we were able to tie back ethnographic insights to the organizational mission, making insights around alleviating the pressure homeowners felt to "lift up" their neighbors less difficult for stakeholders to hear. Ethnographers must learn about what is at stake for the client beyond the business (financial) objectives.

Another principle of how ethnography provides a long-term impact is strategic engagement. If one were to think about these principles and their direct and indirect measures as variables in a multi-variate equation, strategic engagement would be a binary variable—we find a way to do it or we miss out on an opportunity to increase our impact. Working in-house certainly opens more opportunities for extended engagement, even if engagement is not always easy (see Beers et al. 2011; Guth 2022). The novel approach to pain points I implemented brought up multiple opportunities to work collaboratively on product roadmaps. Strategic engagement certainly doesn't require standing in through the ideation or execution of a new product or service, but is more about shepherding insights through the process and ensuring the voice of the beneficiary or consumer isn't drowned out by shareholders, efficiencies, or competing priorities. As a consultant, you may build in workshops or initial concepts as post-readout deliverables. In the case of Habitat, I remained active as part of an "advisory" group as the core team moved forward with approving building plans and volunteer safety requirements for multi-family units. Strategic engagement is a process of building communitas at the "edges" and "interstices" (Turner 1969, 128). It is about finding strategic allies and identifying ways to extend the conversation past the readout.

The final principle is resiliency and sustainability of outcomes. If we manage to identify key opportunities for positive (financial, social, organizational) change that address both short- and long-term needs, the impact is not simply outcome + outcome = impact. Change is the constant. Incremental improvements don't necessarily keep us moving towards a singular goal. If a research engagement can

help to set a strategic vision that is achieved over years, including additional research and planning, the initial research is not invalid. Rather, the direction set by foundational, strategic research lays out a path to a sustainable business model. Again, Habitat Charlotte didn't choose a single path forward and they didn't prioritize their hypothesized model above the other possibilities. Rather, they followed multiple paths that allowed the organization to remain responsive to dramatic market forces, like the 2008 housing crisis. The number of homes they built in 2007, 2008, even in 2012 is far less important than the sustainable, affordable housing model that they have built by honing in on community in both the literal and the figurative sense. If you ask me about the impact of my research in 2006, I would tell you that by shifting the mental framework of how they build, Habitat Charlotte created a flexible housing model that includes infill housing, renovated homes, townhomes, and entire communities, allowing them to nearly triple the number of families served annually even as Charlotte-area housing costs skyrocket. The quantitative impact is more homes and more families served. The ethnographic impact is operational flexibility, strengthened feelings of community, and increased levels of service.

Conclusion

Ethnographers often shy away from assigning value to our work, but as we move into a new decade of ethnographic praxis, it is time to embrace our value in concrete terms. Being concrete does not, however, mean we have to fall into the trap of over-indexing on short-term financial metrics. With ethnography, it is not an either-or choice between short-term gains and long-term value. Our craft is one of the few that is equipped to explore and define ambiguity, liminality, and multiple time frames, so we can use this skill built through our practice to define the immediate, long-term, and organizational impact measures that situate our work solidly in both tangible outcomes for our clients and the populations they serve. As time passes, it's useful to revisit these measures to understand how our work has prompted deeper organizational changes or created resilient solutions because of the ways in which ethnography illuminates potential challenges and opportunities. Most importantly, for the generation that will be the new torchbearers of our craft, thinking through impact at the start of a project, whether or not that vision is realized, can help frame the work for stakeholders and support the continued evolution of our field.

About the Author

Katherine Metzo, Ph.D. is currently a Lead UX Researcher at Lowe's Home Improvement. She has worked in academia, grassroots policy, customer insights, and design research. Finding mastery in the craft of ethnography and illuminating fundamental aspects of the human condition are the threads that weave these chapters together. katherinemetzo@gmail.com

Notes

The author is grateful to Matt Lane for support at multiple points during the writing process. The views expressed in this paper are solely those of the author and do not reflect those of Lowe's Home Improvement.

1. 1983 is the year Habitat for Humanity affiliate was formed in Charlotte.

2. Visit Sauro 2015 for key quantitative UX metrics. See Hazen, et al. 2020 for an example of how to apply quant measures to design research.

3. See Youngblood and Chesluk 2021 for a framework for studying indirect and possibly involuntary product experiences based on the physical and social environments we inhabit.

References Cited

Atkinson, Paul. 2015. *For Ethnography*. Los Angeles: Sage.

Beers, Robin, Tommy Stinson, and Jan Yeager. 2011. "Ethnography as a Catalyst for Organizational Change: Creating a Multichannel Customer Experience." *Ethnographic Praxis in Industry Conference Proceedings* (2011), pp. 61–76. https://epicpeople.org/ethnography-as-a-catalyst-for-organizational-change-creating-a-multichannel-customer-experience/

Charland, Carole and Karen Hofman. 2017. "Bringing Attention to Problem Solving and Meaningfulness at Work: How Ethnography Can Help Answer Difficult Business Questions." *Ethnographic Praxis in Industry Conference Proceedings* (2017), pp. 320–335. https://epicpeople.org/problem-solving-meaningfulness-at-work/

Cliver, Melissa, Catherine Howard, Rudy Yuly. 2010. "Navigating Value and Vulnerability with Multiple Stakeholders: Systems Thinking, Design Action and the Ways of Ethnography." *Ethnographic Praxis in Industry Conference Proceedings* (2010), pp. 227–236. https://epicpeople.org/navigating-value-and-vulnerability-with-multiple-stakeholders-systems-thinking-design-action-and-the-ways-of-ethnography/

Cooper, Alan. 1999. *The Inmates Are Running the Asylum: Why High-Tech Products Drive Us Crazy and How to Restore the Sanity*. Indianapolis: SAMS.

Desmond, Matthew. 2014. "Relational Ethnography." *Theory and Society* 43(5):547–579.

Geertz, Clifford. 1973. *The Interpretation of Cultures*. New York: Basic Books.

Guth, Kristen L. 2022. "Creating Resilient Research Findings: Using Ethnographic Methods to Combat Research Amnesia." *Ethnographic Praxis in Industry Conference Proceedings* (2022), pp 294–316, https://www.epicpeople.org/creating-resiliency-of-research-findings/

Hasbrouck, Jay. 2018. *Ethnographic Thinking: From Method to Mindset*. New York: Routledge.

Hazen, Rebecca J, Genny Mangum, and Tom Souhlas. 2020. "Scaling Experience Measurement: Capturing and Quantifying User Experiences across the Real Estate Journey." *Ethnographic Praxis in Industry Conference Proceedings* (2020), pp 117–126, https://epicpeople.org/scaling-experience-measurement-real-estate-journey/

Hurston, Zora Neal. 1996 (1942). *Dust Tracks on a Road: An Autobiography*. Harper Perennial.

Kucheriavy, Andrew. 2017. "How UX is Transforming Business (Whether You Want It to Or Not)." *Forbes*. January 23, 2017. Accessed 8/5/2024. https://www.forbes.com/sites/forbestechcouncil/2017/01/23/how-ux-is-transforming-business-whether-you-want-it-to-or-not/

Maxwell, Chad. 2021. "Demonstrating the Impact of Ethnography: Chad Maxwell on ROI and New Kinds of Ethnographic Value." *EPIC Perspectives*. Accessed 8/5/2024. https://www.epicpeople.org/demonstrating-the-impact-of-ethnography/

McClard, Anne Page and Therese Dugan. 2017. "It's Not Childs' Play: Changing Corporate Narratives Through Ethnography." *Ethnographic Praxis in Industry Conference Proceedings* (2017), pp 336–348. https://epicpeople.org/childs-play-changing-corporate-narratives-ethnography/

Metzo, Katherine. 2008 "Homeownership and Community Transformation: Habitat for Humanity in Charlotte, NC." *Anthropology News*, 49(9): 11,14

Reksodipoetro, Adri and Alexandra McCarter. Nd. "How Ethnographic Exhibits Can Shift Business Paradigms." *EPIC Perspectives*. Accessed 8/3/2024. https://www.epicpeople.org/how-ethnographic-exhibits-can-shift-business-paradigms/

Sauro, Jeff. 2015. "10 Metrics to Track the ROI of UX efforts." Measuring U blog (September 1, 2015) Accessed 8/5/2024. https://measuringu.com/ux-roi/

Sih, Brady, Hillary Carey, and Michael C. Lin. 2018. "Getting from Vision to Reality: How Ethnography and Prototyping Can Solve Late-Stage Design Challenges." *Ethnographic Praxis in Industry Conference Proceedings* (2018), pp. 457–465, https://epicpeople.org/ethnography-prototyping-late-stage-design/

Stevens, S. S. 1946. On the Theory of Scales of Measurement. *Science*. 103(2684): 677–680.

Turner, Victor. 1969. *The Ritual Process: Structure and Anti-structure*. New York: Aldine De Gruyter.

Youngblood, Mike and Benjamin Chesluk. 2021. *Rethinking Users: The Design Guide to User Ecosystem Thinking*. Amsterdam: Bis Publishers.

Ethnographers as Intermediaries: Plurality as a Double-Edged Sword in Web3

MRINALINI TANKHA, *Portland State University*
CHRIS ROGERS, *Independent Contractor*

This paper addresses a misalignment between market motivations and community needs in Web3, arguing that comparative ethnographic work can identify strategic opportunities in the industry. Comparing findings across different design research projects, we show how Web3 builders and founders are driving a plurality of projects and products and enabling cross-chain operability to build diverse blockchain-based ecosystems. Although this plurality is generative for innovation in the industry, it de-prioritizes fully fleshed out, end-to-end Web3 solutions to urgent social problems, and is increasingly displacing community-centered perspectives. We found that users desire a different kind of plurality—interoperable financial instruments across traditional banking and blockchain-based platforms.

Introduction

Web3–the industry that formed around the newest iteration of the World Wide Web using blockchain distributed ledger technology to enable novel transaction layers–is emerging from the most recent so-called "crypto winter." The anticipated crypto spring serves as a timely opportunity to address some key challenges underpinning potential growth and popular adoption of industry technologies and use cases. In this paper, we compare findings from three Web3 design research projects realized through industry-academic partnerships. We argue that builders and founders in this space contribute to—and strive toward—an expansive landscape full of projects and products constructing diverse ecosystems interconnected by new digital economic infrastructures (Crypto Research & Design Lab 2022). A significant increase of interest in the space over the last five to ten years produced an uncountable number of projects. This *plurality* is a highly sought after and admirable ambition for Web3 builders seeking to keep the industry diverse and decentralized. However, many of these projects remain disconnected from each other. Blockchain infrastructural projects emerging within the industry have begun to cohere around a direction of enabling cross-chain interoperability, or the functionality of digital assets such as cryptocurrencies and other digital tokens to move between various blockchain networks, which may eventually address some of the fragmentation undercutting smooth user experiences with Web3 tech. Meanwhile, as solutions develop to connect the patchwork of projects and networks, we find the plurality

Tankha, Miralini, and Chris Rogers. "Ethnographers as Intermediaries: Plurality as a Double-Edged Sword in Web3." *EPIC Proceedings* (2024): 163–181. https://www.epicpeople.org/ethnographers-as-intermediaries. ISSN 1559-8918.

frenzy ends up distracting from fully fleshed out end-to-end Web3 solutions to urgent social problems. The blockchain stack continues to grow more complex and further from on-the-ground experiences with the technology. While focusing on plurality is generative for more holistic exploration and innovation in the industry, it is also increasingly displacing user-centered or community-centered perspectives, as these agendas continue to be driven by a small number of people (Crandall 2019; Nelms et. al 2018), largely disconnected from users' most pressing needs–plurality is hence a double-edged sword. In 2024, industry market signals point to new use cases, such as DePIN–decentralized physical infrastructure, the application of decentralized technologies to run networks of physical hardware like IoT sensors, wireless infrastructure, or energy grids–further expanding the plurality of projects and accompanying networks and tokens that constitute the Web3 industry. Yet, end-to-end solutions remain un- or under-developed at key junctures in the user experience, often falling short of delivering on promises of lasting social, economic, or environmental impact for communities (Amirebrahimi 2016).

Our work compares findings across research projects we have conducted on blockchain and cryptocurrency at the Crypto Research & Design Lab (CRADL) (https://www.cradl.org/), the Crypto Council for Innovation (CCI) (https://cryptoforinnovation.org/), and Portland State University.[1] Both CRADL and CCI drive ethnographic research of industry participants to generate rich insights of human-centered experiences in the Web3 space. With the mission of "putting people at the center of crypto" CRADL has produced narrative research translating nuanced understandings of the complex realities of Web3 communities, from builders to people using products.[2] The Crypto Council for Innovation (CCI) too has conducted extensive primary research and created, a first of its kind, Center for a Digital Future: Impact Base (https://centerforadigitalfuture.squarespace.com/impact-base), a collection of Web3 projects actively creating social impact. In our work we have done research with crypto founders, builders and investors as well as end-users using a host of ethnographic methods—interviews, participant observation, financial diaries, payment mapping—to understand the alignments and departures in market motivations and community needs. Based on this comparative work, we propose viewing ethnographers as intermediaries "*who can sit between and at specific nodal points or critical junctures* to bridge, negotiate, translate, or convert different kinds of value" (Tankha and Dalinghaus 2020, see also Latour 2005; Lindquist 2015) for diverse Web3 stakeholders. Considering ethnographers as intermediaries entails thinking

across ethnographic projects, as we are doing in this paper, that can be invaluable in revealing silos, missed opportunities, misunderstandings, and the synergies and tensions in what different communities find valuable. Without industry ethnography, highly specialized spaces like this one are more likely to develop in chaotic echo chambers, where impact is idealized though rarely fully realized in substantial ways. Industry adoption of user experience research is a welcome tool to building more empathetic products. However, without rich ethnography it is insufficient to fully grapple with the broader scope of lived realities *beyond* the lens of the product. The blockchain industry would hence benefit from a reckoning and re-prioritization. Grounding research in understandings of individual and community ecosystems (Youngblood et. al. 2021) instead of just *incentives* is more likely to deliver impact—as opposed to endless design tweaks in pursuit of product-market-fit. Applying ethnography to the Web3 industry can provide pause and redirect our gaze away from technological hyperscaling to help generate precise points of intervention in the lived realities and financial pain points of communities. In doing so, ethnography could in fact help steward the crypto industry back to its initial commitments to disrupting the foundations of traditional banking and ensuring financial inclusion for marginalized and unbanked communities.

Plurality but out of Focus

Our ethnographic research with Web3 founders, builders, and people using products across the United States demonstrated an industry driving force embodied as a passionate exploration of plurality. Participants in the industry demonstrated holistic thinking about how to develop layered, multifaceted ecosystems spanning nearly every conceivable aspect of the pre-blockchain-existing world. This plurality manifests at different levels, shaped by the motivations and imaginations of industry participants and shaping the technologies and outcomes affecting individuals accessing these products; plurality of *projects and products*; plurality of *tokenomic design*; and a unique interconnectedness and plurality of *communities and ecosystems*, intertwined with developments towards cross-chain interoperability. Plurality in and between these levels however, produces a chaotic combination of possible outcomes. A still yet to mature industry lacking standards and best practices produces variety at each level that remain disjointed from each other. The infinite number of possible combinations between product direction, underlying tokenomic design, and community of participants being experimented with can be difficult to keep track of. Reliable metrics for measuring a project's potential sustainability are still being developed, and many unforeseen consequences have emerged from the

experimentation taking place. Without prioritization of concise, targeted use cases, potentially beneficial social, environmental, and economic tools are drowned out by hype-cycles that die in the trough of disillusionment.

Plurality of Projects and Products

An exceptional variety of projects was apparent from the beginning of participant observation fieldwork. An annually held Web3 conference in Austin, Texas provided the space for dozens of presentations, with speakers touching on all conceivable aspects of the industry. Predictably, many sessions focused on decentralized finance, blockchain governance, and technological design structures of blockchains and their accompanying tools. Sessions on the potential social impact of blockchain cited projects attempting to improve access to financial tools, remittance capabilities, and financial literacy of marginalized communities, such as the formerly incarcerated or experimentation providing indigenous communities with access to carbon asset markets. Although the projects in that particular session appeared to embody goals of equitable "financial inclusion," many other projects appeared to adopt this language without demonstrating a proven ability to deliver on those impact goals through last mile services. Other sessions featured networks boasting ecosystems with over 600 projects in development and ambitious goals of becoming the "largest on-ramp in Web3 history," or pointing to their communities exceeding half a million members. Blockchain gamification—the application of game theory and design along with tokens to incentivize activities such as learning objectives or industry participation–mining, social reputation tokens reimagining digital data management, and UI and Web design projects, like that of Brandon Eich, the creator of Javascript and founder of a web browser with data harvesting protections and crypto integration, all attracted audiences eager to learn about the future of Web3.

Our research also took us to Manhattan, where we attended an industry social event and observed projects sharing the Web3 dedicated coworking space to develop disparate projects side-by-side. There, Web3 gaming startups applying tokens to in-game exchangeable 'skins'—personalized character components—developed their visions alongside former traditional finance folks who identified the need for treasury management among Web3 projects, starting their own company to fill this need. While seeking research participants in the San Francisco Bay area we encountered a founder who, after immigrating to the United States, was working to establish his micro-tasking startup which utilized its own internal crypto economy to facilitate payments between participants. Analyzing industry publications and media

highlighted projects experimenting with derivatives in the DeFi sector and others committed to reconfiguring digital identity to provide individuals with more control over their data. Throughout the research process we heard industry participants voice concerns over centralizing effects, such as individual network dominance, calling for "strong competitors" they saw as keeping the industry diverse and innovative. Plurality seemed to emerge naturally as the industry expanded and participants discovered emergent needs, implementing legacy processes such as treasury management mentioned earlier and novel ones, like that of transaction ordering technologies, or MEV bots—computer programs coded and deployed to order transactions in the most efficient and value-return maximizing fashion.

Plurality of Tokenomic Designs

Tokenomics: Novel Assets and a Driving Force for Plurality

Cryptocurrencies are digital asset tokens created specifically for a particular blockchain network. Transactions in the network are carried out and recorded in these assets. In legacy electronic credit and debit processing a centralized authority, such as the Automated Clearing House, holds the sole responsibility for processing and verifying transactions. In Web3, community members volunteer to participate in this process with the possibility of being rewarded in cryptocurrency for their efforts. While various blockchain networks approach coordinating community members participating in this process somewhat differently, for a blockchain to be a decentralized technology this process is always carried out and confirmed by community participants and not a central intermediary. At the center of this process is tokenomics—a portmanteau of 'token' and 'economics'—the process by which these behaviors and functions are designed and incentivized through the application of digital tokens and blockchain ledgers.

Tokens, in this fashion, have no physical representation and exist solely as digital assets tracked on the blockchain. By rewarding specific behaviors with tokens, blockchain networks are able to incentivize the necessary functions to maintain a blockchain ledger through community participation. Although digital tokens like these make claims to monetary qualities, the reality is due to few opportunities for digital tokens to be accepted as a means of exchange beyond the limited sphere of the blockchain network they are native to, they only aspire to be considered as money in the way that fiat currency achieves. Blockchain ledgers, advocates point out, are cryptographically secure, transparent, and offer efficient economic and monetary control levers. As a public, decentralized technology the transaction history on the ledger is visible to anyone with a computer and internet connection.

Designing tokenomic systems is a complicated and intricate process. Designs vary significantly, from aspirational digital cash systems–like that of the Bitcoin network–to digital carbon asset representations. Bitcoin, while complex as a technology, functionally intends only to serve as a pseudonymous–accounts in the Bitcoin ecosystem are tied not to identities but to alphanumerical addresses–digital form of cash or cryptographically-secured currency. The public nature of decentralized blockchains and an industry motivation toward open-source technologies facilitated the creation of thousands of other cryptocurrencies to enter the market after Bitcoin, either by copying the design of Bitcoin directly or iterating on it in some way. Some of these projects have even surpassed Bitcoin according to certain metrics. Ethereum, for example, surpassed Bitcoin as the most used blockchain network over four years ago.

While conducting research into Web3 builder experiences, we encountered a multitude of blockchain network designs. Systems developed since the launch of Bitcoin intend to serve innumerable different purposes. Some of the tokenomic designs we observed during research included 'tokenized–represent as a digital token recorded on a blockchain ledger–carbon assets, wireless connectivity credits, complex investment vehicles, community membership, and certificates of ownership over real-world and digital assets like art or property, among many others. While the majority of ethnographic research examining the blockchain industry focuses near-exclusively on the 'gold-standard' crypto, Bitcoin, the truly expansive and far richer nature of the broader Web3 industry and its diverse participants and objectives are missed entirely.

Tokenomics Enables Plurality

The future imaginaire of Web3 promised to deliver a version of the World Wide Web characterized by 'decentralized' digital economies liberated from the control of centralized intermediaries and their monopoly over individual data and use of technology. Our research revealed significant challenges standing in the way of realizing this vision. Among the experts designing tokenomic systems, however, Web3 still held the possibility of far more than the popularly observed use of cryptocurrencies as speculative investments. These participants expressed nuanced, and seemingly counterintuitive, approaches to achieving the decentralized vision. A tokenomics consultant working with a popular firm articulated these considerations:

> It's important for projects to consider why they want to be decentralized and what their optimal level of decentralization is, and how important it is

> to reach a certain point of completion of the project before taking on the risk of opening it up to a wider group.

In working to achieve the Web3 imaginaire, passionate industry professionals viewed tokenomic systems and their accompanying communities as the key tools to realize a new, economically integrated and individually sovereign, World Wide Web.

Leveraging tokenomics enables Web3 builders to offer unique mechanisms through which communities can form and interact. Novel tokenomic incentives and touch points–such as NFTs–have given Web3 founders the ability to produce a plurality of ecosystems and communities through their various projects. Research revealed a number of challenges in effectively developing a sustainable tokenomic system. These challenges involved design considerations, community and individual actor behavior, and institutional influences, all of which complicate Web3 claims of decentralization and equitable access. Research involved participants from all over the world and included field sites on nearly every corner of the United States, where we observed Web3 participants pursuing 'plurality' and diversity in nuanced and unexpected ways.

Plurality of Tokenomic Designs

As blockchain networks expanded beyond the initial 'digital cash' use case they have become increasingly more complex. So, too, has the tokenomic designs behind these systems. The growth in complexity and number of individual token systems can make it difficult for participants to engage in these systems and move between them efficiently. Ecosystems featuring multiple tokens, each with specific utilities and characteristics, significantly increases the complexity and approachability of these systems for Web3 participants. Even the single token ecosystem of Filecoin, a decentralized blockchain network for file storage services, imbues their token with multiple functions. $FIL, the native asset of the Filecoin ecosystem, is used to pay for goods and services, retrieve stored data, reward validators for securing the network, and functions as a yield generating asset when 'staked'–the temporary forfeiture of exchanging the token. Tokenomic design attributes may catch traction and be adopted by other communities for their perceived success, while others seek to create functions and designs that improve upon established models.

In an example of the potential for tokenomic complexity, our field research led us to an early-stage carbon credit startup still in the midst of developing their tokenomic design. Their initial tokenomic system was to include three separate tokens to account for different stages in the carbon cycle. The complexity of this system included mandatory token holding periods, overlapping spheres of exchange,

and participant specific token uses. Mapping out this system with their team helped us to better understand the objectives and intentions behind the design, which appeared at first unnecessarily difficult to navigate. Tokenomic experts continue to explore the bounds of tokenomic design, refining and innovating an industry role that is central to the goals of the community at large.

Plurality of Communities and Ecosystems

Plurality is central to the future Web3 imaginaire. Achieving this vision, however, is complicated by the inability of digital assets to move seamlessly across various blockchain networks. Industry participants in our research were passionate about addressing these challenges through cross-chain interoperability—or the technological capabilities of blockchain networks to communicate with each other and facilitate the transfer of assets between them. Interoperability featured prominently across all research domains, from numerous sessions at the convention in Austin to podcasts and articles surfaced during discourse analysis. Convention speakers articulated frustrations with "siloism" and argued for interoperability as a means to reduce "tribalism" they characterized the industry by. One speaker called for greater measures in the security of interoperability technologies stating matter-of-factly: "We all know [the future is] going to be cross-chain." In another instance, a podcast guest explained the constraints of a single blockchain project attempting to be the sole processor of transactions globally:

> There is a misunderstanding that there is ever going to be a single chain that is going to manage all of the world's capacity or demand for blockspace. There is no such thing—[blockchains have] a finite amount of [available space to record new transactions].

Web3 'plurality' for industry builders and participants is not a singular goal, but rather, it is functional, even predicted to be technologically necessary for scaling the industry; it is experimental, with complicated and intricate technologies stretching the boundaries of what is possible; it is entrepreneurially savvy in that a greater variety of projects being built leads to greater involvement from the public, and thus more individuals participating, adding value, and strengthening Web3's goal of embedding itself in the future of digital environments; it is also an honorable goal for the genuinely equitably-minded; and it is sometimes a social signal for the performative actors destigmatizing their projects by riding the coattails of more altruistic endeavors. Industry builders clearly recognize friction exists. Whether the solutions

they produce are ultimately framed through user needs or company needs is yet to be seen.

The User in Focus

The dizzying levels of plurality in the Web3 industry seem to be increasingly out of touch with the large-scale urgent problems it initially purported to address. Providing speedy, transparent and unlimited cross-border payments and remittances have been one of the most critical use cases for blockchain and cryptocurrency, for instance. The industry has largely neglected to account for last-mile processes in cross-border payments, falling conveniently short of delivering the impact promised. This is in part due to regulatory constraints but also, due to a hyper focus on building a blockchain based ecosystem de-linked from existing banking infrastructure. Interoperability with the existing banking infrastructure is not a priority for Web3 developers and continues to be either non-existent or extremely expensive (Maurer et. al. 2013; Nelms et. al. 2018; Swartz 2018). Or as others have argued and our research confirms, cryptocurrency rails rely on more informal local networks for off-ramping, thereby hardly providing a new alternative global financial system (Rodima-Taylor & Grimes 2019, Tankha 2021). In the case of cross border payments, some Web3 founders also commented that it is hard to pitch the story of remittances to US investors that do not recognize market opportunities that don't reflect their experiences. For instance, in conversation with a startup providing blockchain based remittance services to Latin America, the founder described challenges in securing venture capital funding for the project due to lack of relatability. He noted:

> I pitched to 100 American funds, and probably 90 of them just didn't really have the full understanding of what we're doing [remittances for Latin American immigrants in the US] and why that matters. Because we get questions like 'Why don't they use Venmo?' If I have to explain why you cannot use Venmo to send money to Mexico then there's really no point in [trying to pitch this person]. Instead of a pitch it should be a classroom session. And my realization was, if you have never faced that problem, it's really hard for you to understand it…if you've never had to send money to someone in another country, then it's hard to understand why that matters.

In two user focussed projects on financial exclusion with tech entrepreneurs in Cuba and immigrant families and business owners in South Florida, we found that cross border transactions were in fact one of their main challenges. Our largely immigrant interlocutors in South Florida, used a host of financial services from traditional banking services (checking/saving accounts; credit and debit cards; Zelle)

to apps (Venmo, Cash App, Apple Pay) and formal and informal remittance services (Western Union, Moneygram, *mensajerias* or money couriers). Though they all encountered frictions with banking institutions in their everyday lives that made their trust in traditional banking fairly fragile, the most attractive use case for cryptocurrency was for making cross-border payments rather than for their daily expenses or as an investment, given its volatility. Many of them incurred high fees and had experienced fraud and delays in receipt of cross border payments. The alleged transparency and timely settlement of crypto money transfers was therefore promising. As part of the research design, the participants in South Florida were given money to load onto a cryptocurrency wallet to test out what value it would add to their existing financial ecologies. At the end of the study, we found that for our participants, off ramping or cashing out of the crypto infrastructure was necessary but expensive and inconvenient when making cross border payments and therefore did not make using cryptocurrency a viable option.

For instance, one of our interlocutors, Mariana, is a single mother of three children. Originally from Belize, she moved to Fort Lauderdale 10 years ago but has not held a steady job. She joined a community college to start the process of entering a nursing program, but had to quit when the pandemic hit. She has a terminal mother in Belize and has many medical bills she has to pay for her. She uses Western Union to send money to her brother in Belize who is in charge of her mother's care. She struggles with the high fees Western Union charges as well as the daily limits it sets on money transfers. Due to her mother's costly hospital bills, she often needs to transfer amounts of money much above the daily limits and since she is unable to send the entire amount, it causes disruptions in her mother's care. When she has attempted to send amounts above the daily limit, Western Union has flagged her payments and has to go through a long drawn-out process of proving that she isn't, in her words, "a scammer or fraudster." With no limits on daily transactions and purportedly less fees, cryptocurrency seemed like a great solution. The global cryptocurrency wallet being tested by the participants also came with a Visa debit card that was only issued to persons in the United States and not to users in other countries. Mariana had her brother in Belize open the same crypto wallet that she was using in Florida. She sent him money using the wallet and he received it almost immediately. She was pleasantly surprised at how fast the transaction was completed, unlike with Western Union where she sometimes had to wait for several hours or even a few days for her brother to receive the money. However, for the money to be used for payments in Belize, her brother had to cash out of the crypto wallet by

linking it to his bank account and in that transfer, he incurred high fees and taxes. To avoid these charges, Mariana thought about opening a crypto account for her brother registered in the US so that he would be issued a debit card which he could use in Belize. But he would then face a less favorable exchange rate as well as international transaction fees to withdraw cash from an ATM in Belize, leaving them back at square one. Hence, while the transactions were faster, the end-to-end processing of crypto payments ended up being more costly than sending money via Western Union.

Another one of our research participants, Terrence is a small business owner residing in Miami, Florida. He is a lawyer by trade but doesn't practice anymore. He moved to Miami to start a tobacco importing business and more recently diversified into chocolate, coffee and liquor. A majority of his suppliers are based in Latin America and Miami is the first port of entry. He talks passionately about dark chocolate and gadgets. His businesses took a hit during the Covid-19 pandemic due to supply chain issues, but they have now mostly recovered. Terrence tracks his finances very closely. This was something that he picked up from his mother. "My mother always said, 'check what you got and then see what you made.' I do it every day. I know what's gone out and what came in," Terrence told us. As a small business owner making regular cross border payments to suppliers in Latin America, he finds himself paying several fees. The simplest thing for him to do are wire transfers and the most complex is letters of credit. His experience with traditional banks has not been the most pleasant. "They charge a lot of fees! I can tell you stories. I am a small business man and these sums can put you out of business." Terrence would hence welcome an alternative to traditional banking in his business life. Because Terrence has several businesses that require transactions both domestically and internationally, he has issues with sending money overseas and also the naming structure of his entities. For instance, he has a cigar company with the word "Havana" in its name. Due to U.S. embargo restrictions against Cuba, every 3-4 months his transactions get flagged and Terrence has to produce evidence confirming that he is not sending money to Cuba. With his chocolate business too, the international wires he makes are often routed through Panama that he says is the "kiss of death" because Panama has a notorious reputation for money laundering. Transactions get held up for 30-45 days. To avoid these risks, he has several entities through which transactions are routed. For instance, he uses a family trust as a distributor for the chocolates so that "it looks like it's a nice little family-owned business." He also uses letters of credit when he doesn't have an existing relationship with the vendor. These letters of credit too are expensive to produce, sometimes they charge more than one percent of the transaction. Despite all the challenges

Terrence faces with making cross border payments he doesn't think cryptocurrency is the solution since there isn't a developed ecosystem. While he acknowledges that transacting in cryptocurrency could make global payments easier, he doesn't think he could convince his vendors in Latin America to open crypto wallets. He says that it is a question of getting comfortable with and trusting the technology and being able to convert more easily to more accepted forms of payment. Until then he doesn't see how cryptocurrency can be useful in his transactions. Letters of credit from banks, while cumbersome and expensive, guarantee and insure transactions and people trust them. For Terrence, there is a higher chance of "seeing whales with wings" than people having that level of trust in cryptocurrency.

Similarly, another small business owner Javier, originally from Colombia, is starting a business importing gold from a small network of miners in Bolivia. He co-owns the business with his father and some of his father's friends. He is very tech savvy and already has several crypto wallets and other money transfer apps. He needs to send large volumes of money to Bolivia for his business but is finding that this is becoming more difficult with completing transactions to pay his miners. Much like Terrence's case, his payments repeatedly get flagged for fraud and "suspicious activity." He then has to provide detailed explanations about where the money has come from and where it is going and if he doesn't the money is frozen. Unlike crypto transfer, banks and even Western Union are not operational 24 hours or on the weekend, so delays in processing that goes to the following business day do not capture the real time price of gold and exchange rates in a transaction. The miners he works with also live in remote areas and have to travel long distances to the nearest bank or Western Union. Another issue is that much of the gold business still operates in cash, both for the suppliers in Bolivia as well as the buyers in the US. As a result, Javier says that he would like to use cryptocurrency, especially to pay his miners, but since they don't live in cities, they often don't have reliable internet infrastructure, so getting them to transact in cryptocurrency is a non-starter. Moreover, these miners would need to cash out of the crypto ecosystem in order to pay for their other needs and as mentioned before this involves incurring high fees and taxes. Even though Javier is extremely keen on using cryptocurrency and understands its potential, there isn't a fully operational end-to-end solution to support his needs.

In their financial lives, people juggle multiple financial services and tools to meet different needs. The paired concepts of monetary ecologies (assemblages of instruments and technologies that together make up the world of value and exchange

in people's lives) and monetary repertoires (the diverse ways people might use, deploy and move between the different components of their monetary ecology) capture the complexity and plurality inherent in the lived experiences of finance (Maurer 2015; Tankha 2016; Tankha and Dalinghaus 2020). For people then, plurality is about having multiple interoperable financial tools between which they can, easily and cheaply, move, convert and generate value. Creating a blockchain based ecosystem with a plurality of tokens and products that are just interoperable across different chains, and not with existing financial services, falls short of addressing people's most pressing needs. Yet there is generally a lack of motivation for offramping when addressing strategies behind token implementation. For instance, a panelist at a Web3 conference illustrated the somewhat unspoken truth that for many companies in deploying tokens, the solitary goal of that asset is to drive and retain users *within* the network, describing tokens as "your customer acquisition strategy," and a "mistake [to think] the token is the business model." This sits starkly in contrast to remittance communities who are most in need of cash or other avenues to procure basic necessities. There is, therefore, a misalignment between how developers and investors are thinking about plurality and the kinds of plurality people and users desire in their financial lives. How then do you build ecosystems without users and how do you increase users without use cases for their existing challenges?

Moreover, the hype cycles of blockchain and crypto and their drive to create peer-to-peer decentralized economies divert from the fact that these ecosystems continue to be contingent on state legal institutions as well as informal local networks for off-ramping. In our work with tech entrepreneurs in Cuba, we found that many people on the island were transacting in cryptocurrency despite significant risks. Crypto exchanges based in the US or even elsewhere are bound by US sanctions and geoblock Cubans from operating on their platforms. Cubans, however, use virtual private networks (VPNs) to hide their location and IP addresses to access these exchanges. Our interlocutors explained that using crypto allowed them to access work and get paid in the global marketplace through platforms such as GitHub. Many of them used crypto wallets as an investment or store of value to protect from heightened inflation in the Cuban economy. However, transacting in crypto comes with risks because Cubans have no recourse in the event that they lose access to their accounts or forget their passwords and are asked to produce identity documents. Crypto wallets also cannot be linked to Cuban bank accounts, so Cubans either have family overseas to on-ramp or off-ramp their wallets with installments of cryptocurrency, or they sell/purchase cryptocurrency in person from local traders in Cuba in exchange for cash. Even so, crypto remittance services were started by tech

savvy Cubans but money was sent to personal crypto wallets of a network of programmers and software engineers that then made payouts to individuals or utility companies in either cash or through *Transfermovil*, Cuba's mobile application for making online payments. An entire informal network of cashing in/out of the crypto ecosystem has developed in Cuba, that is layered upon existing informal economy linkages, to complete transactions. Blockchain-based financial services do not rise above global banking restrictions and cannot provide alternate last mile services to financially excluded communities, such as those sanctioned by US geopolitics (Tankha 2021). In spite of the promise of creating more inclusive economies, crypto solutions are ridden with obstacles and risks for such communities.

Separating Noise from Impact

There are no doubt several challenges in the regulatory environment as well as in fostering trust in new financial infrastructures that impede the delivery of truly plural and interoperable solutions for pressing community needs. The lack of regulatory clarity around cryptocurrency leads to ambiguities about the legality of particular crypto platforms and tokens and therefore entails higher costs for specialized legal services to provide some assurance that they will not attract unwanted attention from the U.S. Securities and Exchange Commission or other regulatory enforcement agencies. Tokenomics is also highly complex and specialized, making it hard for companies to locate the expertise needed to design a truly sustainable ecosystem (Crypto Research & Design Lab 2022). The few that do possess the expertise command high premiums and are many times only accessible to networked and well-resourced projects. The industry also faces a serious threat from a high prevalence of scammers that impact mainstream perceptions of cryptocurrency, slowing adoption and making it difficult for honest teams to avoid destructive models that might harm their communities. Despite such frictions, we did encounter some builders working to overcome these challenges and create platforms more aligned with everyday realities and more devoted to ideologies of inclusion and shaking the foundations of traditional banking and finance. Part of this research involved creating the Crypto Council for Innovation's (CCI) Center for a Digital Future: Impact Base, a first of its kind collection of primary research on Web3 projects actively creating real world impacts. Developing this database through CCI included dozens of hours of research evaluating Web3 projects. Startups and established businesses alike were assessed by their reputations, business sustainability, and history of creating–or potential to create–positive impact, among other factors.

Founders from groups of historically excluded communities, in particular, found the tools that Web3 offers to provide opportunity to reshape their realities and those of their communities. Whereas Web2 had solidified in exclusionary ways, Web3 was seen as a promising avenue to pursue equitable access and individual autonomy. An Indigenous technologist, educator, and entrepreneur who works as a project manager for a community bank explained during an interview:

> There's such a long history of [indigenous folks] being subjected to bigger powers, and tech is one of those powers now. Figuring out a way to create an organization that has power that mirrors our community structure is incredibly important. We do have to disrupt and create something completely different that's on our own terms.

This individual was part of a team interested in developing a localized token economy that would encourage commerce and value circulation among the indigenous community while reducing potential leakage, or the tendency for value to be diverted from a particular economic system. Participants were passionate about creating new ecosystems formed by and for their own communities, not subject to the controls and exclusions they experience in Web2. Blockchain tokenomics and its ability to create bespoke economic tools was observed as a unique opportunity for founders of color to subvert the shortcomings of legacy systems: An entrepreneur and founder of a public cryptocurrency developed for the Global African Diaspora community told us:

> ...[As] Black people, we've been building a tremendous amount of social capital on platforms like Facebook and...our hands are tied—we can't really do anything with it. We're like slaves to the platform. Once we introduce the idea of blockchain...we're creating as a community...we can put our social capital there and it could get back to us in many different ways.

Through blockchain and tokenomic design, this project sought to encourage mutual support among the global Black community, fostering economic activity and building governance and collective decision-making power through blockchain voting mechanisms.

Other communities experimented with reimagining digital organizational structures without traditional hierarchies. A consultant who started one of the first decentralized organizations for designing token-economic systems for projects told us: "I always recommend starting with a community first, because then you have something proven and something tested, before you actually do the token." Careful consideration of the community emerged as such an important part of the design process that some tokenomic consultants we spoke to said the entire project might

end up being undermined without a clear community and a reflection of their values present in the tokenomic design. Sustainable applications of tokenomics to facilitate community formation emerged as a far more diverse and nuanced area than what was originally thought. When designed and deployed thoughtfully, tokenomics was revealed to be an essential tool for Web3 builders, opening up new possibilities for these projects including attracting participants to them, driving activity within the ecosystem, offering new models for project funding, and creating new forms of digital assets.

The foundations of the Web3 industry are actively being formed. Complex and intertwined relationships between the public and the various groups and individual actors involved in driving and guiding the industry mold what it will stand atop–or be perceived to be standing atop. Mainstream attention to Web3 has a displacing effect on communities and groups working toward generative and novel models of more equitable digital and economic realities for marginalized communities. Drowned-out by crypto-company implosions and token scam schemes, communities and builders (such as the ones mentioned in the previous section) are missed in mainstream or even industry self-assessments of what Web3 embodies today. Ethnographically-oriented researchers are uniquely equipped to shed light on individuals and projects working to build the foundations of more equitable systems through blockchain technology. Where impactful use cases are obscured, ethnographers can give voice to the beneficial projects unfairly grouped alongside grifts and carelessness.

Conclusions

In an industry like crypto that is innovating and expanding at a breakneck speed, comparative analysis across ethnographic projects illuminates disarticulations at different scales and between the goals and needs of diverse stakeholders and communities, but also helps chart new directions by archiving and uplifting projects geared toward lasting social impact. Ethnographers can, therefore, be valuable critics of broken systems as well as advocates of new forms of value creation. Academics and industry professionals have shown how ethnographers can serve as intermediaries translating between institutions and communities, interpreting and leveraging quantitative data toward humanistic ends (Levin 2019; Maiers 2018), bridging the "temporal mismatch" between futuristic industry thinking and the current everyday concerns of people (Van Leeuwen & Singh 2023). For the crypto industry, we have shown how comparative ethnographic approaches allow us to

critically interrogate notions of plurality in cross-chain blockchain ecosystems, celebrated by founders and builders, and show how they eclipse communities' needs for a plurality of interoperable financial instruments. Rather than stop there, however, we argue that ethnography can also be future facing in identifying where innovation can be targeted. Deep ethnographic work can be leveraged to distinguish strategic opportunities from the gale of silver bullet solutions, to deliberately and responsibly lay new foundations in the crypto industry, block by block. This work also entails addressing the role of politics and regulatory frameworks in the making of payment infrastructures, whether as protections for users or as exclusions tailored into the fabric of our global financial systems. As intermediaries then, ethnographers are tasked with not only interpreting social dynamics, critiquing institutions, and challenging power relations but also advocating for strategic innovation disrupting the existing foundations that fail to serve underrepresented communities. Such efforts could shift the tech paradigm toward developing platforms and products, not for the most economically viable customer, but for who it can benefit the most.

About the Authors

Mrinalini Tankha is Assistant Professor of anthropology at Portland State University with research interests in money, technology, and design. She has worked on projects in Latin America and the United States on financial inclusion, architecture and urban planning, user-interface research and board game design.

Chris Rogers is a social impact Web3 researcher working in the space since 2021. He holds a master's degree in cultural anthropology and a graduate certificate in business blockchain from Portland State University. His most recent work documented the experiences of Black and Latinx founders in the Web3 industry.

Notes

[1] While the individual research projects we did were conducted in these organizations, our arguments in this paper zoom out and put the findings from these different projects into conversation. Therefore, the arguments presented here do not necessarily reflect the perspectives of these organizations.

[2] This research, published as a number of in-depth reports, highlights the efforts of Web3 builders towards social and community impact, including Black Experiences in Web3, Investigating Generational Wealth, and Income and Wealth Creation in Web3.

References Cited

Amirebrahimi, Shaheen. 2016. "The Rise of the User and the Fall of People: Ethnographic Cooptation and a New Language of Globalization." 2016 Ethnographic Praxis in Industry Conference Proceedings, 71-103. https://www.epicpeople.org/ethnographic-cooptation/

Crandall, Jillian. 2019. "Blockchains and the 'Chains of Empire': Contextualizing Blockchain, Cryptocurrency, and Neoliberalism in Puerto Rico." Design and Culture 11:3, 279-300.

Crypto Research and Design Lab. n.d. Crypto Research & Design Lab. https://www.cradl.org

Crypto Research and Design Lab. 2022. "Black Experiences in Web3." Crypto Research and Design Lab. https://www.cradl.org/black-experiences-crypto

Crypto Research and Design Lab. 2022a. "Income and Wealth Creation in Web3." Crypto Research and Design Lab. https://www.cradl.org/new-income-wealth-web3

Crypto Research and Design Lab. 2022b. "Investigating Generational Wealth." Crypto Research and Design Lab. https://www.cradl.org/generational-wealth

Crypto Research and Design Lab. 2022c. "New Ecosystems in Web3." Crypto Research and Design Lab. https://www.cradl.org/new-ecosystems-in-web3

Crypto Council for Innovation. n.d. Center for a Digital Future: Impact Base. https://centerforadigitalfuture.squarespace.com/impact-base

Latour, Bruno. 2005. Reassembling the Social: An Introduction of Actor–Network-Theory. New York: Oxford University Press.

Levin, Nadine. 2019. "Ethnographic Agency in a Data Driven World." Ethnographic Praxis in Industry Conference Proceedings, 591–604, ISSN 1559-8918. https://epicpeople.org/ethnographic-agency-data-driven-world/

Lindquist, Johan. 2015. Anthropology of Brokers and Brokerage. International Encyclopedia of the Social & Behavioral Sciences 2, 870–874.

Maurer, Bill, Taylor C. Nelms, & Lana Swartz. 2013. "'When Perhaps the Real Problem Is Money Itself!': The Practical Materiality of Bitcoin." Social Semiotics 23(2): 261-77.

Maurer, Bill. 2015. How Would You Like to Pay? How Technology Is Changing the Future of Money. Durham and London: Duke University Press.

Maiers, Claire. 2018. "Reading the Tea Leaves: Ethnographic Prediction as Evidence." Ethnographic Praxis in Industry Conference Proceedings, 351–363, ISSN 1559-8918. https://epicpeople.org/ethnographic-prediction-evidence/

Nelms, Taylor C., Bill Maurer, Lana Swartz, and Scott Mainwaring. 2018. "Social Payments: Innovation, Trust, Bitcoin, and the Sharing Economy." Theory, Culture, and Society 35(3): 13-33.

Rodima-Taylor, Daivi & William W. Grimes. 2019. "International Remittance Rails as Infrastructures: Embeddedness, Innovation and Financial Access in Developing Economies." Review of International Political Economy 26:5, 839-862

Swartz, Lana. 2018. "What was Bitcoin, What Will It Be? The Techno-Economic Imaginaries of a New Money Technology." Cultural Studies 32(4): 623–650.

Tankha, Mrinalini. 2021. "Detained settlements: The Infrastructures and Temporalities of Digital Financial Transactions between the United States and Cuba." Economic Anthropology 8:1, 133-147.

Tankha, Mrinalini. 2016. "The Mobile Money Experience in Sub-Saharan Africa: Lessons from the Institute for Money, Technology and Financial Inclusion (IMTFI)." The African Technopolitan Magazine (Issue on "The Poverty of Development Strategy in Africa) 5, January: 96-105.

Tankha, Mrinalini & Ursula Dalinghaus. 2020. "Mapping the Intermediate: Lived Technologies of Money and Value." Journal of Cultural Economy 13:4, 345-352.

Van Leeuwen, Gijs & Abhigyan Singh. 2023. "Local Frictions in the Energy Transition: Design Anthropology for the Emergence of Energy Communities." Ethnographic Praxis in Industry Conference Proceedings, 277-294, ISSN 1559-8918. https://www.epicpeople.org/local-frictions-in-energy-transition-design-anthropology-for-energy-communities/

Youngblood, Michael, Benjamin J. Chesluk, & Nadeem Haidary. 2021. Rethinking Users: The Design Guide to User Ecosystem Thinking. Laurence King Publishing.

Ethnographic Perspectives on AI Chatbots in Customer Support

DAVID RHEAMS, *Atlassian*

This study on the integration of artificial intelligence (AI) and large language models (LLMs) presents key findings for the effective and ethical deployment of AI in organizations. It presents original research on the deployment of an AI-powered assistant at Atlassian that was designed to support engineers by summarizing technical documentation and generating example customer responses. Using ethnographic methods and computational analysis, the study examined the nuanced impacts of AI on labor dynamics, efficiency, and empathetic communication within technical support roles. Findings include the importance of understanding the socio-technical environment of AI integration and the ethical considerations associated with AI deployment in customer service settings.

1. Introduction

In March 2023, Atlassian's customer service team began to explore artificial intelligence (AI) and large language models (LLM) chatbot prototypes to reduce the effort required to solve technical customer support issues. The goal was to provide support engineers (SEs) with an AI-powered assistant that summarized and consolidated complex technical documentation while providing an example customer response. My role as a qualitative researcher was to help a team of developers and data scientists build the prototype AI assistant.[2] I designed a research plan focusing on participant observation to understand where the AI assistants fit into their daily work. However, at the time of the study, AI was perceived as a threat to customer service jobs. The newness of the technologies and even my role as an Atlassian employee advising senior leadership presented challenges to ethnographic research.

This case study breaks down the research challenges dealing with sensitive labor issues and AI technologies. It offers insights into how we developed guiding principles, conducted qualitative research, and combined ethnographic practices with quantitative analysis. This approach enabled us to delve into how our engineers use an AI assistant to bolster confidence in their writing, reveal their creativity, and deepen their empathetic communication.

The study ran through several prototype iterations from March through August 2023. By the time it concluded in August 2023, we had identified practices to guide the next stage of product development and developed a perspective on using AI in our customer service teams. Using an ethnographic method, we uncovered product requirements, discovered performance indicators, and cultivated trust with study participants.

Rheams, David. "Ethnographic Perspectives on AI Chatbots in Customer Support." *EPIC Proceedings* (2024): 182–194. https://www.epicpeople.org/ai-chatbots-in-customer-support. ISSN 1559-8918.

Ethnographic practices are essential to understanding AI tools' integration into complex work environments. Conducting interviews to determine what users value and how AI made them more efficient is only one part of the process, and it is insufficient on its own. Understanding the environments we inject AI into and adopting broader, more nuanced research approaches that force questions beyond operational efficiency is crucial. Experimenting with AI assistants in customer service environments is not as low risk as it initially appears. The impact on labor, training, and knowledge structures carries significant ethical and power implications for technical support teams. Writing off customer service as low-risk is not only shortsighted but unethical, given the profound effects AI can have on these workers' roles and livelihoods.

1.1 Customer Service as a Study Site for Organizational Ethnography

Customer service teams within software companies have broad perspectives because the role requires day-to-day interaction with customers, engineers, and management. Those teams capture insights into customer behavior, and they work to fill the gap between how the software works and how customers expect it to work. They are experts in a particular domain, with a deep knowledge of a company's systems and technical architecture. The result is a profession that blends the ability to communicate with a customer, understands abstract business problems, and has technical knowledge to identify software malfunctions. The breadth of their daily work makes them ideal candidates for organizational ethnography, as they encounter customer and coworker requests.

The unique technical support position within a company is one reason organizational anthropologists have studied technical support teams for decades. Julian Orr's *Talking about Machine*s identified technical support as a unique study site in 1990 when technical support was mostly field technicians going into physical business to fix copiers or mainframes. Orr describes the environment of the field technician this way: "The work of technical service involves the community of technicians, the community of users, and their respective corporate entities in addition to the machines, and it occurs in a public arena, the customer's place of business" (Orr 1996, 3). Each of the communities Orr identified was important as we considered how using an AI tool might impact customers and engineers, so we focused on an inclusive approach to capture the voices of these communities in our work.

We were especially attuned to the public discussions around AI, and we recognized that support engineers (like many others in the tech industry) were

sensitive to labor issues in an industry rife with recent layoffs (Lee 2024). If we wanted an accurate picture of how support engineers reacted to an AI assistant, we needed to ensure we were building tools designed to augment their work, not replace it. Most importantly, we required the research participants to understand our goal so that they could speak honestly about AI tools.

2. Study Design

We provided access to the prototype artificial intelligence assistant to 30 study participants. The program team invited participants into a private Slack channel where they interacted with the prototype by typing a message. A Slack automation delivered a response to the SE inside a threaded chat message. However, the prototype did not allow for a ChatGPT-style back-and-forth conversation. Instead, the bot provided answers in four steps inside the chat thread. First, it restated the question to ensure alignment. Second, it provided a short list of summarized sources to formulate the response. Third, it drafted a reply to the question, and fourth, it allowed the support engineer to provide feedback on accuracy. For example, when a support engineer asked, "What prevents a customer from editing a custom field in Jira?" the chatbot replied by restating the question, listing relevant technical documentation, and drafting a response.

The chatbot aimed to save time by eliminating lengthy searches on document repositories and public websites. Additionally, the chatbot could write simple SQL queries or code snippets to resolve issues. After reviewing the answer, the support engineer indicated the accuracy of the response with a *thumbs-up* or *thumbs-down* emoji.

These interactions generated detailed quantitative data on accuracy and time efficiency. We tracked the percentage of tickets where the AI assistant was used and the number of accurate versus inaccurate rankings to build insights on potential time-savings for support engineers. The research team also had access to every ticket that the support engineers worked on and flagged AI-assisted tickets. This extensive tracking produced a wealth of data illustrating when and how engineers interacted with the AI assistant. However, metrics alone cannot tell the whole story of adoption, SE happiness, efficiency, or accuracy. To capture the detailed experiences of support engineers interacting with the AI assistant, we employed ethnographic methods to uncover the nuances of these interactions, working with the 'thick descriptions' of anthropology (Geertz 1973).

2.1. Using Mixed Methods to Set a Baseline

Building thick descriptions requires a nuance and context of social interactions and behaviors. We were as interested in the subjective experiences, emotions, and attitudes towards the AI assistant as we were in the *time-to-resolution* of a support ticket or statistical accuracy. The team quickly recognized that we needed more contextual knowledge of the support engineers' daily work habits, training, and team interactions. Recognizing this gap, we understood that while participant observation and interviews would be the cornerstone of our qualitative research, preliminary work was necessary to build context and better understand the support engineer's specific environment. As we built the qualitative research plan, we arrived at a mixed methods approach, using techniques from both digital and organizational ethnographic practices. Working from Sam Ladner's *Mixed Methods* (Ladner 2019), where she builds on Bryman's (2006) ideas, we combined questionnaires, content analysis, participant observation, and in-depth interviews to build a comprehensive picture of the support engineer's workday.

We used questionnaires three times throughout the study (approximately one per month) to ask contextual questions before we started interviews to establish a baseline for people's perspectives on AI. The questions were a mixture of Likert and open-text, though we found the open-text fields to contain the most helpful information. We extracted anonymous quotes from these text fields and used them to expand the purview of our interview questions.

To further expand our context, we conducted a content analysis of 100 completed support tickets where the engineer used the AI assistant early in the research program. Krippendorff's (2018) comprehensive framework for content analysis guided the approach to systematically examining the textual data from customer support tickets. The analysis compared two components: The original text proposed by the AI assistant versus the text used by the support engineer in the ticket. We reviewed each ticket to identify linguistic patterns between human and AI text.[3]

The content analysis provided context for participant observation and allowed us to validate or disprove early theories. In combination with the ticket analysis, I also assessed the AI-generated responses by observing support engineers (SEs) interacting with the AI chatbot. This approach enabled the identification of patterns in SE behavior, specifically how they selectively incorporated beneficial aspects of the AI's suggestions while disregarding irrelevant or inaccurate parts. Consequently, this process allowed for the identification and categorization of the response

elements that were most valuable to the SEs, helping to corroborate other data sources and develop a vocabulary for coding interviews.

2.2. Interviews & Observation

Our data collection combined primary methods: semi-structured interviews and participant observation through shadowing sessions. Each session was conducted in a recorded Zoom meeting, though all participants were accustomed to Zoom as Atlassian is a fully remote workplace. The interviews were designed with structured questions but kept a conversational tone, allowing for adaptive variations of our questions to fit the natural flow of each conversation, making each interview slightly different. The tactic allowed the conversation to cover multiple topics, drawing out nuances we may have missed otherwise.

We accumulated approximately 17 hours of recorded sessions. While that is quite a bit of content for two researchers to collect and analyze over 12 weeks, we wanted to ensure we had participants from as many countries and Atlassian tenures as possible. To make the work manageable, we used Zoom's recording feature so that a single researcher could conduct the interview without the burden of taking notes.

For interview transcription and data management, we used Dovetail, a research repository platform, which facilitated automated transcriptions and allowed us to highlight and tag key segments of the conversations. The transcription tagging feature was instrumental in coding the data. For example, we tagged the moment in the transcript with "confidence" when a support engineer reported that an interaction with the AI chatbot made them feel more confident in their response.

This approach provided rich contextual and behavioral data that gave quantitative results added depth and revealed contractions between reported behaviors and observed actions. During the shadowing sessions, I was able to watch for these contradictions as the support engineers navigated their daily tasks—resolving tickets, responding to customer inquiries, and troubleshooting issues. My role was to maintain a minimal presence during observation, though I did encourage the support engineers to articulate their decision-making processes and describe their actions in detail as they went through their work.

We held shadowing sessions before the interviews so that I could ask follow-up questions during the interview session. The combination of recorded interviews and observation sessions allowed us to capture detail and nuance that would have been lost otherwise. However, using these features required substantial trust from our support engineer colleagues. One of the factors that made our participants willing to sit down, invite us into their work, and share potentially sensitive information in a recorded Zoom session was the clear guiding principles we shared with participants.

3. Guiding Principles

The program team prioritized writing guiding principles and incorporating them into our documents and practices at the start of the project. We created these principles to ensure we agreed to a specific set of ethical practices, and we wanted to build trust with every team we interacted with at Atlassian. We took a lesson from *Ethnography: Principles in Practice,* emphasizing the importance of building trust with participants and navigating the complexities of hierarchy and team dynamics (Hammersley and Atkinson 2019). We also worked with our internal Responsible Technology team to help craft specific principles for our research project.[4] Establishing these principles early and sharing them with our counterparts in other departments was one of the significant early successes of the program team. We aligned on the following principles: privacy, transparency, and reduction of the potential for bias.

3.1. Support Engineer Privacy

While the program team could not keep all study participants completely anonymous in a corporate environment, we were willing to share results with our internal colleagues. For instance, we removed individual and team names from public reporting on quantitative data. We took the same approach for recorded interviews and shadowing sessions (Murphy, Jerolmack, and Smith 2021, 47). We recorded the sessions but kept them stored on a platform with limited access (managers for study participants did not have access to recorded interviews). Prior to each session, we let the participant know that they were being recorded and that we required their explicit permission before quoting them. In each document, we also clarified that we were studying a feature or system, not an individual's skill. At the start of every recorded session, I stressed this aspect of our research to let the SE know that the conversation was private and would not appear in a performance review. We obtained an agreement from managers that non-participation in the study would not reflect poorly on a support engineer.[5]

3.2. Transparency

Along with allowing users to opt out of the study and limiting direct attribution to their interviews, the program team consciously decided to be as transparent as possible at every study step. Lack of accountability is a pervasive issue in AI systems (O'Neil 2016). Algorithms can create feedback loops that reinforce inequality and operate without transparency. In customer support, AI-driven decisions can be

opaque, making it difficult for customers to challenge or understand the outcomes. This lack of accountability can erode trust in AI systems. It also aligns with one of Atlassian's core values: "Open Company, No Bullshit." While no artificial intelligence system can ever be fully transparent, we strove to clarify how we applied AI and machine learning technologies, that our research process and findings were public, and that the tool provided as much transparency as possible.

The program team believed that we had an obligation to train people on how LLMs processed information to democratize access to the latest AI technologies and make decisions on how to incorporate them into their daily workflows. We created training videos and technical documentation that explain the prototype AI assistant and how to query a large language model to make the technology more approachable and less anxiety-provoking for new users. We hosted group training sessions on Zoom so that our support engineers could ask questions of the developers and researchers working on the AI assistant. The training sessions had the added benefit of improving the chances of a support engineer using the AI assistant. The program team received overwhelmingly positive feedback from these training sessions and videos.

In addition to training, we wanted to ensure transparency in the research methods and processes. Our process documentation and anonymized research results were open to anyone in our customer support teams.[6] We provided clear guidance on the actions the program team was tracking and how we were analyzing results. For instance, we tracked every ticket the SEs worked on during the study and almost every click and input on the prototype AI assistant. Each action was tied back to an individual user but anonymized in our reporting. To build trust, we let the SEs know which performance indicators we were tracking and why we decided to track them.

However, the team recognized that tracking performance metrics may make support engineers nervous in an environment that studies AI. Being transparent with our research participants meant acknowledging that AI has the potential to radically transform the way we work and making sure that the topic of labor was a critical part of the discussion. We encouraged SEs to share concerns or ideas in interviews and questionnaires so that we could address these questions and bring them into public company conversations.

3.3. Reduce Bias

Another way we sought to build trust with our teams was to reduce the potential for bias in our research. The past decade has yielded research and warnings that even seemingly small or insignificant product decisions may have an outsized impact on

the people who use the software (Gebru 2020; Buolamwini and Gebru 2018; Broussard 2018; Broussard 2024). The research program was assigned support teams for the prototype. We divided the teams into control and participant groups so that we could measure changes in behavior.

To account for potential biases, we included teams from multiple countries, people at different stages of their careers, and people with different levels of comfort with AI. We worked with support engineering managers and team leaders to ensure we built the study to incorporate as many perspectives as possible. We ended up with a group that spoke at least seven different languages from five continents. Some of our study participants were early adopters of technology, and others were less interested in artificial intelligence. Participants were also allowed to opt out of the study for any reason on a 'no questions asked' basis so that we could ensure that all employees were willing participants.

We worked with the data scientists on the project to determine the number of study participants and the length of time we needed to ensure statistically significant results for the quantitative A/B test. In an ideal world, the research team would have spent as much time with our control group (without access to the prototype) as we did with the group testing the AI assistant. Additionally, observing meetings where our control and test groups interacted with each other would have yielded context for conducting interviews and shadowing sessions. However, with limited time and a small team, we had to choose where to apply our efforts without risking the guiding principles we established.

When I spoke with the study participants, I built rapport by telling them about my background in technical support and the goal of the study. However, as I was speaking with study participants regularly, both in person and on Slack, I was concerned that I would inadvertently signal my positions on artificial intelligence, Atlassian products, or company goals and influence the participants. To alleviate my concern, I began interviews by acknowledging my position within the company and telling the study participant I was concerned about accidentally influencing their answers to interview questions. This approach opened up a conversation that I found valuable and resulted in some of our most surprising insights.

4. Outcomes

The guiding principles were one reason we achieved positive outcomes for the study. Not only did it help establish trust, but it also helped create a spirit of exploration and a willingness of our participants to try new things, give honest feedback, and help us understand their daily work life. The business goal of the study

was to determine the impact on efficiency, deliver a strategy for future AI assistant customer service tools, and gauge support engineer's happiness when using an assistant. At the end of the study, we had a series of documents that culminated in a clear strategy for the next iteration and a business case for further development in AI assistants for support engineers.

4.1 Efficiency

We collected enough quantitative data to have statistically significant results on efficiency measures and the time required to find answers to support tickets. Our results aligned somewhat with what others in the industry found in early 2023. Harvard Business School published a working paper that showed the effectiveness of AI on a customer support team. Working with Boston University and the Boston Consulting Group, the paper concluded that productivity increased for consultants who used an AI assistant by over 12%. They found that support consultants who were "below the average performance threshold" increased their productivity by 43%. While the specific efficiency gains for our cohorts of support engineers differed, the outcomes were comparable. In short, more experienced support engineers saw fewer gains in efficiency, while support engineers with only two years of Atlassian experience saw the largest bump in productivity (Fabrizio et al., 2023). Our results differed when we measured the productivity of employees with less than a year of experience supporting Atlassian products. The AI assistant was far less helpful to new hires because they needed to have the context to ask the right question and had to put additional effort into validating the assistant's response.

While we could answer the primary questions surrounding efficiency, the quantitative data alone did not explain the behavior of the study participants. The content analysis on customer support tickets found that the percentage of AI-generated text within tickets where we used the AI assistant was less than 20%. We rarely found more than a few phrases or bullet points written by AI in the final product, even when the support engineer flagged the AI answer as correct and helpful. SEs were using AI as a drafting tool rather than relying on it as the final output because writing a ticket is a source of pride for the team. While some SEs were motivated by the quality assurance reviews for Atlassian's support, most wanted to ensure that the customer response was their voice.

We broke a ticket response into four primary sections to understand where a human added the most value and where they used the AI assistant. We found that SEs typically wrote the salutation and restated the issue to ensure they understood the customer's request. There were no instances of AI contributing directly to these sections of a ticket. We found some overlap when the SE wrote directions to a

customer on how to use the product to take steps in a workaround. The SE also wrote all the following steps and salutations. When asked, SEs responded that the AI functioned as a way to write instructions and code snippets efficiently.

4.2. Empathy

Yet, when we watched how support engineers used the tool, some of our participants used AI to assist with translations. We documented three moments within the support process where AI potentially helped with translation. First, some SEs used AI to translate the customer's request if it needed clarification. Second, the SE would translate the customer's request into a more complete AI prompt by providing context or technical information. Some support engineers use AI assistants to translate concepts from their native languages into English (Atlassian's workforce and customers are global, though most customer support communication is in English).

In one instance, a support engineer from Brazil translated a phrase common in Portuguese into English to see if it lost its meaning during translation or risked being misinterpreted by the customer. The translation wasn't simply a word-for-word translation of a technical concept; it was a translation of a pun. The study participant wanted some of his personality to come through on the ticket to build a relationship with the customer. While he felt that the AI assistant's writing could not help him build rapport with a customer, they felt more confident using the AI tool as a sounding board for tricky phrases. The word *confident* surfaced in some capacity in over 80% of the interviews and shadowing sessions. While the direct text may have been discarded, the SEs did report that it helped them cross language barriers and feel better about their writing to show that they empathized with the customer's situation.

The moments where the AI helped the support engineer be more confident also sparked creativity when using the tool. When the AI could simply help consolidate documents and provide a quality summary, the SEs treated it like any other tool that helped them solve an issue. They were pleased but didn't stop and take any particular note. However, when the AI assistant was used more creatively, the support engineers stopped and noted how they were crafting customer communications.

4.3. Beyond Interviews

One of the more common refrains I've heard working on technology products is that "we must talk to the user," accompanied by a push to do user interviews. The impulse to talk to people is admirable and a practice we should continue. However,

one of the realizations I had doing this project was that simply conducting interviews is insufficient when dealing with sensitive topics like artificial intelligence and labor. Despite our early claim that we were "augmenting, not replacing" support engineers, it became increasingly apparent that the complexities of AI's impact on labor required a broad view of the technology beyond the chatbot we were testing. AI is already significantly affecting hiring and labor practices. For example, Intercom's leadership has emphasized the transformative potential of AI in customer service, predicting that AI agents will eventually handle the majority of customer queries, fundamentally changing the nature of human support roles (Adams 2024). Similarly, Microsoft's recent layoffs, attributed to strategic realignments in response to AI advancements, highlight how AI reshapes labor markets (Landymore 2024). These developments underscore that AI, much like past technological advancements such as robotics or the cotton gin, is poised to transform labor markets in unpredictable ways. Ethnographic studies on AI must employ methods that fully recognize and address the impact on labor. However, we found that we needed to create intentional spaces where our participants could freely discuss both the positive and negative aspects of AI, independent of their specific daily tasks, to uncover a broader range of perspectives.

Building a broad perspective also requires time. The additional time spent observing Slack channels, attending study participant team meetings, and having informal conversations was invaluable to the study's success. The openness about potential harms and willingness to slow down and watch participants wrestle with the new technologies provided more precise insights than we may have had otherwise.

About the Authors

David Rheams, PhD is a senior business architect at Atlassian and a lecturer at The University of Texas at Dallas in the Arts, Technology, and Emerging Communication department.

Notes

I am deeply indebted to the team at Atlassian who worked on this project. Andrew Clarke, Tom Albrecht, John Kim, Vineet Prasad, and Stephen Sifters were integral to this study and their expertise and wisdom cannot be overstated. Our numerous conversations and collaborations directly contributed to the development of the guiding principles and success of this research.

[1] Atlassian is a software company that builds tools for collaboration and productivity. Their core products are Jira, Confluence, and Trello.

[2] A program team in a software company is a cross-functional group responsible for planning, executing, and managing a project. Our team was composed of software engineers, data scientists, a program manager, customer support leaders, business analysts, and business architects.

[3] We broke the customer service tickets into their components to determine which part of the ticket was used for typical pleasantries, which parts were used for technical information, and which parts were additional customer information. We also compared sentence structure, tone, and directness between AI and human-written text.

[4] Atlassian has an amazing Responsible Technology team who I consulted with at the onset of the project. They were incredibly quick to help shape the specifics we needed for our research group and they play a vital role in development at Atlassian. They can be found here: https://www.atlassian.com/trust/responsible-tech-principles

[5] Some support engineers accepted our offer to opt out of the study because they had other projects to focus on. We risked biasing towards people who enjoyed AI technologies by letting people opt out, but the program team felt it was worth the risk.

[6] We ensured that the guiding principles were the same for both the qualitative and quantitative methods. In particular, our approach stemmed from the principle that our product was not designed to replace customer support professionals; stakeholders reiterated this principle in meetings and strategy planning.

References Cited

Adams, Paul. 2024. "There's No Going Back: AI-first Customer Service Has Arrived." Intercom. https://www.intercom.com/blog/ai-first-customer-service/.

Broussard, Meredith. 2018. *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge: MIT Press.

Broussard, Meredith. 2024. *More Than a Glitch: Confronting Race, Gender, and Ability Bias in Tech*. Cambridge: MIT Press.

Bryman, Alan. 2006. "Integrating Quantitative and Qualitative Research: How Is It Done?" *Qualitative Research* 6 (1): 97–113.

Buolamwini, Joy and Timnit Gebru. 2018. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." *Proceedings of Machine Learning Research* 81: 1–15.

Fabrizio, Dell'Acqua. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." Harvard Business School Working Paper 24-013. Harvard Business School. https://www.hbs.edu/ris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf.

Futurism. 2023. "Microsoft's Layoffs and AI's Impact on the Labor Market." Futurism. https://www.futurism.com/microsoft-layoffs-ai-labor-market.

Geertz, Clifford. 1973. *The Interpretation of Cultures*. New York: Basic Books.

Gebru, Timnit. 2020. "Race and Gender." In *The Oxford Handbook of AI Ethics*, edited by Markus D. Dubber, Frank Pasquale, and Sunit Das, 122–144. Oxford: Oxford University Press.

Hammersley, Martyn, and Paul Atkinson. 2019. *Ethnography: Principles in Practice*. London: Routledge.

Krippendorff, Klaus. 2018. *Content Analysis: An Introduction to Its Methodology*. 4th ed. Thousand Oaks: SAGE Publications.

Ladner, Sam. 2014. *Practical Ethnography: A Guide to Doing Ethnography in the Private Sector*. Walnut Creek: Left Coast Press.

Ladner, Sam. 2019. *Mixed Methods: A Short Guide to Applied Mixed Methods Research*.

Lee, Roger. 2024. "Layoff Tracker." Layoffs.fyi. https://layoffs.fyi.

Marcus, George E., and Michael M. J. Fischer. 2014. *Anthropology as Cultural Critique: An Experimental Moment in the Human Sciences*. Chicago: University of Chicago Press.

Murphy, Alexandra K., Colin Jerolmack, and DeAnna Smith. 2021. "Ethnography, Data Transparency, and the Information Age." *Annual Review of Sociology* 47 (April): 41–61. https://www.annualreviews.org/doi/10.1146/annurev-soc-090320-124805.

O'Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown Publishing.

Orr, Julian E. 1996. *Talking about Machines: An Ethnography of a Modern Job*. Ithaca: Cornell University Press.

Experts in the Loop: Why Humans Will Not Be Displaced by Machines when there Is "No Right Answer"

TOM HOY, *Stripe Partners*
JANNEKE VAN HOFWEGEN, *Google*

This paper explores the resilience of human expertise in an age of Large Language Models (LLMs) and establishes a framework and methodology for the ongoing value of human experts alongside conversational AI systems. Our framework identifies three themes that define the unique value of human expertise: curating knowledge, personalizing interactions, and offering a perspective. For each theme, we outline strategies and tactics used by human experts that AI struggles to meaningfully replicate at scale. We contrast this with complementary strategies that AI systems are uniquely positioned to offer. Finally, we provide design and engineering principles to guide product teams seeking to shape the implementation of this emerging technology in a way that augments, rather than automates, human experts across consumer domains. The framework and methodology are described in practice through a case study of travel planning—a domain where there is "no right answer".

In the second half of the 20th century, the field of Artificial Intelligence (AI) began developing what became known as "expert systems". Expert systems were hand-coded to automate narrow, repetitive tasks (Dreyfus 1987). These systems are useful when the "right answer" is universal and can be coded from textbooks and experts, such as a system making medical diagnosis for specific conditions, or a system designed to judge whether to approve a loan based on fixed rules (Buchanan et al. 2006; Forsythe 2001).

In the 1980s, new AI architectures structured around neural networks could both develop internal representations of knowledge and learn patterns in data. As computing power increased in the following decades, these innovations meant systems were not only defined by hand-programmed rules, but could operate as "learning systems" enabled by vast amounts of data. Tasks such as fraud detection and weather prediction could now be conducted more autonomously, with less need for humans-in-the-loop.

In recent years, these learning systems began to automate elements of more complex tasks such as spam detection, language translation and facial recognition. But these systems complemented or accelerated the work of humans rather than replaced them, especially tasks that required the expertise of "white collar" knowledge workers.

Hoy, Tom and Janneke Van Hofwegen. "Experts in the Loop: Why Humans Will Not Be Displaced by Machines when there Is 'No Right Answer'." *EPIC Proceedings* (2024): 195–223. https://www.epicpeople.org/experts-in-the-loop. ISSN 1559-8918.

Human experts remained uniquely suited to tasks that addressed ambiguous, emergent and complex problems that required novel, iterative approaches and original thinking, problems which may have nuanced answers, or no "right answers" at all. As Frey and Osborne summarized back in 2013, AI systems were still limited to relatively narrow tasks that enable "a programmer to write a set of procedures or rules that appropriately direct the technology in each possible contingency. Computers will therefore be relatively productive to human labor when a problem can be specified—in the sense that the criteria for success are quantifiable and can readily be evaluated."

This narrative shifted significantly in October 2022, when ChatGPT was launched. The adaptable, generative intelligence demonstrated by this new technology gave rise to claims that AI systems are, for the first time, capable of displacing human expertise in these more complex and creative domains. Speculation subsequently arose about the threat of automation to jobs previously believed to be immune to AI systems, such as copywriters, journalists, designers, customer-service agents, paralegals, coders, and digital marketers (Herbold et al. 2023; Lowery 2023).

To what extent, though, is human expertise being replaced in reality? Our research into the value of human expertise in one domain, travel discovery and planning, suggests we should be cautious about these broader claims. Travel is frequently used as an example to illustrate how conversational interfaces powered by Large Language Models (LLMs), such as Google's Gemini (Hawkins 2024) and OpenAI's ChatGPT (Mountcastle 2024), are transforming complex tasks and decisions that would have previously relied on human expertise. Travel shares a number of underlying characteristics with other domains that makes it a useful case study for exploring human expertise, and the extent to which it is automatable.

"No Right Answers" as the Contested Space for Displacement

Exploring travel advice reveals deeper truths about human expertise because travel is a domain in which there is "no right answer" (Pierce 2023). For each traveler there are multiple destinations, lodgings and itineraries that can address their desires and unique context. This makes decision making particularly ambiguous. As The Verge puts it "there's no page on the internet titled 'Best vacation in Paris for a family with two kids, one of whom has peanut allergies and the other of whom loves soccer'" (Pierce 2023).

Markets in which there is "no right answer" contain products that Economic Sociologist Lucian Karpik defines as "singularities" (Karpik 2010). Singularities are

unique products defined by three qualities: multidimensionality, uncertainty and incommensurability. Alongside travel other archetypal examples of "singularities" include luxury goods, wine, literature, music, movies, art, dining out, antiques, real estate and professional services.

Travel is a multidimensional product because it has many aspects and layers—from the taste of the local food to the comfort of a hotel room—that a traveler must try to assess. Travel is uncertain because, no matter how many user-generated reviews one reads, one can never fully know what the real experience of a vacation will be like in advance. And travel is incommensurable because no two vacations can be impartially compared: there is no objective basis on which to say a city break is better than a weekend in the mountains. Travel choices are contingent because they are a matter of taste, preference and circumstance.

This has important implications for how people make decisions about their vacation. Because there is "no right answer", the "market of singularities requires knowledge of the product that far exceeds anything necessary for the standard market" (Karpik 2010). To acquire this knowledge, would-be travelers rely on a variety of "judgment devices" to inform their decisions about travel: social networks, trusted brands and accreditations, critics and guides, expert and buyer rankings, and established sales channels (Jacobson and Munar 2012).

Because travel is an unknowable singularity, travelers must trust judgment devices to help them make their decision. This is where, historically, human expertise has been critical (Confente 2015). Alongside advice from friends and family (social networks), travelers rely on the inspiration and recommendations provided by travel agents, travel journalists, bloggers and influencers (critics and guides), as well as user reviews (expert and buyer rankings).

Today, human expertise published in the form of social media posts, user reviews or articles is often found intentionally by travelers via search, or discovered passively via algorithmic feeds. In both cases the results are tailored to the user in question. However, until now no digital product has been able to make this advice truly "personalized". For example, while Instagram may be able to predict the types of vacation a user might be interested in based on their profile and behavior over time, it does not know the full context and requirements of a specific trip, nor can it iterate recommendations via ongoing feedback.

A truly personalized product is, by definition, the most unique type of product possible because it is tailored to one customer (or customers). Karpik defines it as the "pure[est] form of singularity" (Karpik 2010). This is why, historically, many

travelers seek direct, personal advice from other people. To date, only other humans have come close to offering truly personalized recommendations.

For the first time, LLM technology might change this. Today it is possible to ask ChatGPT what the "Best vacation in Paris for a family with two kids, one of whom has peanut allergies and the other of whom loves soccer" is. And ChatGPT will confidently answer. This confidence stems from the large amounts of web data that a language model is trained on to "understand" the user requirements (family-friendly, peanut allergy, soccer) specified in their query, match it against the most relevant sources in its training data and synthesize an answer. In recent models such as Gemini, this black-box process is made traceable by the LLM supporting its generated answer with direct online sources, further adding to the sense of confidence of its prediction. Does this threaten the livelihoods of travel experts who earn a living by supporting travelers addressing these specific requests? And what, by extension, does that mean for human expertise in other "no right answer" domains?

A Multidisciplinary Approach to Understanding Expertise

The Google Search team is exploring the opportunities afforded by generative AI technology, including Google's own Gemini model, to better serve Search users. It was within this context that a cross functional team of Google product directors, designers and engineers tasked a joint Google and Stripe Partners research team with exploring the nature and value of human expertise in travel, as a case study to inform the new Conversational AI and generative search interfaces they were building. By understanding what expertise was uniquely human, Google could define the scope of emerging products, focusing on where a Conversational AI can add complementary value to existing forms of expertise.

The joint Google and Stripe Partners team consisted of qualitative researchers, linguists and data scientists. Together we designed a novel mixed-methods research approach combining ethnographic observation and in-depth interviews with NLP (Natural Language Processing)-driven analysis.

Our first decision was how to focus the study. We selected a single, versatile travel destination, Los Angeles (LA), as the focus of our study because LA is one of the most commonly searched-for and popular travel destinations in the US (Chang 2023), and offers a diversity of common travel-related activities (from museums to night life to beaches).

We recruited 9 "LA travel experts" ranging across three types of expertise. The first category were LA-focused travel agents who earned a living recommending and selling trip packages directly to travelers. We included travel agents in our sample

because we wanted to capture the value of "professional" travel experts who earned money directly from travel recommendation and planning. Second, we recruited social media influencers who frequently shared LA-related recommendations and content across their channels to their audience. We included influencers in our sample because we wanted to capture people who are earning money indirectly from their LA-related expertise. And finally we recruited LA locals; individuals who had lived in LA for at least 5 years, considered themselves local experts, and frequently shared advice informally to friends and family interested in visiting LA. We included locals because we wanted to explore the value of advice not influenced directly or indirectly by financial incentives.

We paired each expert with a tourist intending to travel to LA within 6 months and arranged a 30 minute "planning and advice session" between them via video call. In addition to recording their natural language conversation, we observed body language and tone of voice to identify pivotal interactions. Observing how experts and travelers interact was critical because, as Diana Forsythe argues in her renowned study of AI engineers attempting to acquire knowledge, experts "are not completely aware of everything they know" (2001).

Following the session we conducted an in-depth interview with the traveler to capture immediate reflections. Our qualitative research team then conducted a "grounded theory" style driven analysis of the 9 advice sessions and traveler interviews to surface common patterns across traveler needs expert strategies.

In parallel, our data science team ran NLP analysis of the sessions. This computational approach is complementary to qualitative observation and interviewing because it provides a detailed breakdown of the turn-based conversation between traveler and expert, allowing us to understand more precisely what characteristics of the language (outlined below) lead to the most successful dynamics. The approach is especially relevant here because we are trying to understand the nuances of conversational interaction, with a view to understanding its implications for a Conversational AI interface.

The NLP analysis was conducted across all interviews, resulting in thousands of expert advice-traveler response pairs which are analyzed at scale to quantify emerging patterns in our interviews.

The computational linguistic analysis included:

- Temporal analysis: (a) Spoken duration for expert and traveler (b) Average pause taken by expert/traveler before responding to traveler/expert (c) Intimate turns initiated by expert/traveler (d) Questions asked by expert/traveler

- Sentiment analysis
- Formality analysis
- Intimate sentence identification
- Specificity analysis

These analyses helped the team identify insights that had been missed in the focused qualitative work, and triggered the team to revisit elements of the interviews to deepen the analysis. Our data science team also helped the qualitative team understand the capabilities and limitations of emerging LLM technology and therefore develop hypotheses on which aspects of human travel expertise are unique, and which can be replicated and/or complemented by LLMs.

Following these analyses we hand-coded a "frontier dataset" to distinguish between instances of positive and negative expert practice, for use by engineers in model training and to improve the explainability of resulting product experiences.

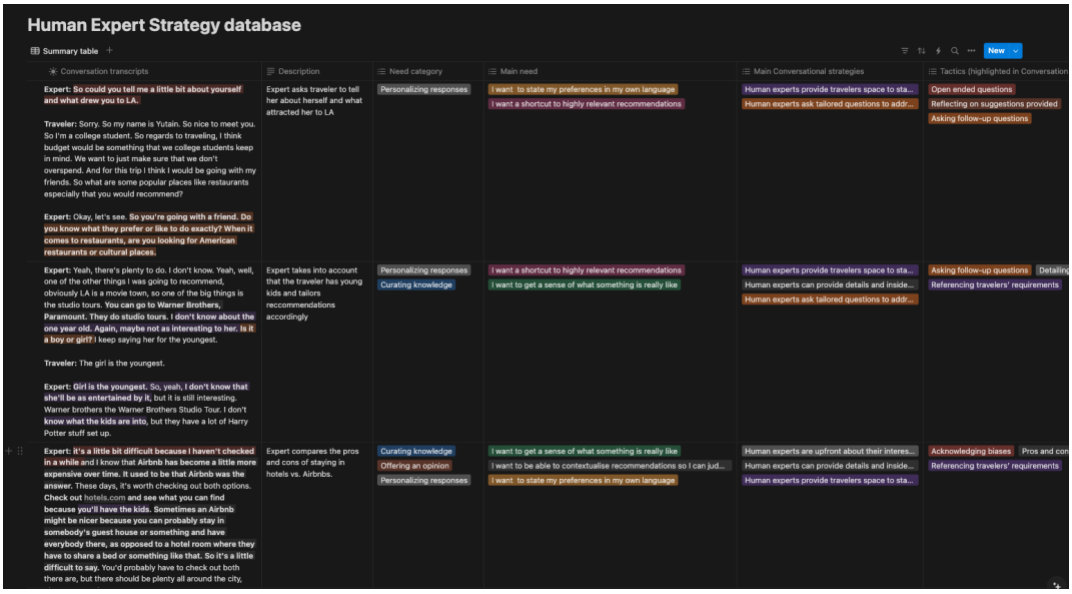

Fig. 1. Hand-coded "frontier dataset" which labels examples of expert best practice for model training and improving explainability.

New Foundations for Human Expertise

Our framework identifies what travelers value the most about human experts when discovering and planning what to do for their next trip. Second, it evaluates which aspects of human expertise an LLM-powered "Conversational AI" can meaningfully learn from, and what remains unique to humans. Here we define a "Conversational AI" as an interface designed to simulate human-like conversation

through Natural Language Processing (NLP) and Artificial Intelligence (AI) techniques.

The framework is broken into three themes that summarize important aspects of human expertise.

- **Curating knowledge:** how human experts obtain, filter and communicate their personal knowledge about a destination with travelers
- **Personalizing interactions:** how human experts shape conversations with travelers to ensure they are surfacing and addressing their specific needs
- **Offering a perspective:** how human experts provide clear direction to travelers by sharing their unique opinion

Within each theme we identified "strategies" and related "tactics" demonstrated by human experts in response to specific traveler needs.

Curating Knowledge

Travelers turn to human experts because they seek unique knowledge that they may not be able to find elsewhere. But travelers don't want experts to share everything. They hope experts will curate their personal knowledge in a way that fits in with their specific circumstances and needs. We identified four specific needs that experts addressed.

Table 1. Curating Knowledge: Human Expert Strategies and Tactics

Traveler needs	Human expert strategies	Human expert tactics
"I want to get a sense of what something is really like"	Humans can provide details and insider knowledge by drawing on their personal experience	Identifying as local
		Detailing specifics
		Personal photos and artifacts
		Pros and cons
"I want to discover unique recommendations"	Humans can suggest off-the-beaten-path options that are not easy to find elsewhere	Hidden gems, insider tips
		Cultural references
"I want to know what all the potential options are"	Humans present a wide variety of options and	Curated lists
		Offering alternatives

	arguments within a curated range	Suggesting what's popular
"I want to learn more about the destination"	Human experts adapt relevant facts and information to add depth and context	Sharing contextual knowledge
		Sharing latest knowledge

We learned that people traveling to LA for the first time want to get a sense of what being in LA is "really like" to help them make the right decisions. Take the example of the advice shared by Jeremy, an LA local who has lived in the city for 20 years, with Donny, a 30-something dad traveling to LA with his kids.

> It is a long standing battle amongst everyone in LA who's got the best tacos. So I would humbly submit that you try to find a truck for Leo's tacos. They're my personal favorite. They're always consistent. They make these amazing Al pastor tacos. It's a seasoned pork taco. They put a little pineapple on it. It's absolutely delicious. And of course, what kid wouldn't love getting food from a truck? Hopefully that's fun for them as well.

Following the advice session with Jeremy, we asked Donny to expand on why he particularly valued the recommendation of Leo's tacos.

> I like somebody telling me, "this is my favorite". And he told me why that one was his favorite: because they add a little pineapple to the tacos when they make it. So just being able to tell me this is what I enjoy and here's why I enjoy it, I feel like that was helpful and adding a nice personalized touch by saying, this one is my favorite: this is the one that I really enjoy.

Jeremy deploys the strategy of sharing insider knowledge by drawing on his personal experience. He starts by saying "[the best taco] is a long standing battle amongst everyone in LA." Here he is subtly establishing his credibility as a local, with 20 years of experience of living in the city (and presumably eating tacos!). This imbues his advice with a specificity that is both unique to his extensive experience and embedded within local social discourse.

Second, Donny appreciates the detailed explanation of why Jeremy prefers Leo's Tacos. When he talks specifics, such as the sprinkling of pineapple, Jeremy helps him home in and appreciate aspects of the experience. This enables Donny to visualize what to expect, cutting through the paralyzing "multidimensionality" outlined by Karpik, while at the same time situating this as part of Jeremy's personal experience and perspective. This, after all, is his "personal favorite".

The value of specificity was further reinforced through the NLP analysis of the interview transcript conducted by our data science team. Similar to the qualitative analysis, the computational analysis investigates which characteristics of Jeremy's advice led to an immediate positive response from Donny. To quantify the specificity of an answer, automated tools detect the presence of named entities (geo locations, named people, restaurant names etc.) in a text and by counting such instances, we can define and evaluate at scale how detailed an expert's recommendation is. We note that the classification of the sentiment of a text is often based on the presence of sentiment-specific lexicon, in this case the prediction relies on "enjoy". Doing this at scale (enabled by the automatic tools we used) across all interviews, we discovered that during the 30 minute advice sessions, moments where the expert became specific correlated with traveler "joy".

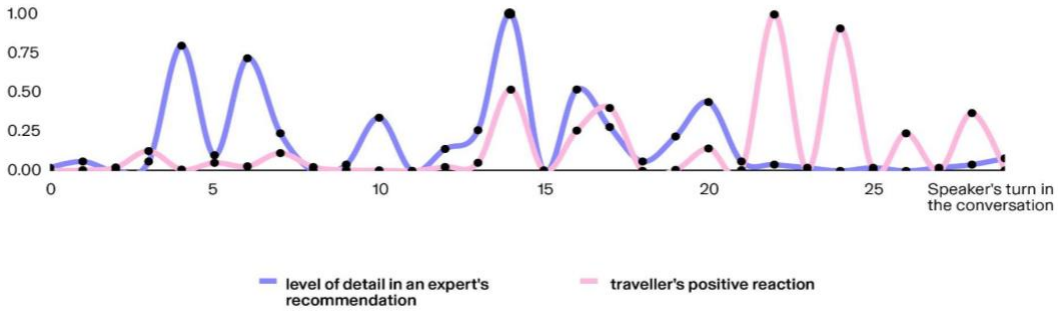

Fig. 2. NLP analysis shows experts' use of specific language partially correlates with the increase of travelers' joy score (joy is a subset of the positive sentiment).

A third and final lesson to draw from Jeremy is his referencing the relevancy of the recommendation for Donny's kids "what kid wouldn't love getting food from a truck?". Here Jeremy is calling back to the trip context Donny has shared, making the advice feel truly tailored and unique to his trip (which is the definition of personalization established by Karpik—see Personalizing Interactions below for more).

Other successful expert tactics we observed within "Curating Knowledge" included sharing personal photos and videos to showcase specific recommendations. This helped travelers visualize the experience more vividly. And seeing the expert they are talking to experiencing LA for themselves provides, perhaps, a form of vicarious enjoyment which feels more real and compelling than the photos of unknown strangers.

Elsewhere experts shared tips and tricks on how to get the most out of popular tourist destinations: where to park or which spot to have a picnic in. Travelers valued this because it was information they were unlikely to easily find, and if they did, were uncertain to trust it.

Curating Knowledge: Implications for AI Systems

The unique knowledge that human travel experts provide is anchored in their specificity as individuals. Expert knowledge is accrued through personal lived experience, creating a natural boundary between what is known and unknown. In the context of overwhelming information and the lack of an objective "right answer", these human limitations are often experienced positively by the traveler. Because it is impossible to seriously consider every option within a destination, the knowledge base of an individual expert provides a breadth of choice within a curated range, providing the traveler with a meaningful level of agency within boundaries. The near unlimited range of information available to AI systems, and the nature of their training, makes this aspect of human knowledge, especially that informed by lived experience, hard to meaningfully replicate.

Second, the specificity of lived experience also endows experts with knowledge that is unlikely to have been digitized, and therefore is not part of the data that language models are trained on. It exists within the "complex domain" of human culture, illegible to machine systems (Hoy, Bilal & Liou 2023). This "complex" knowledge can be collected from experts and translated into training datasets—as it was in this project—but this requires an intentional effort from product teams, and will not scale to the aspects of lived experience described above.

The fundamental challenge for Conversational AIs, despite their increasing capability and multimodality, is a wider epistemological weakness inherent in LLMs: the "symbol grounding problem" (Harnard 1990). As Bender and Koller explain, "language is used for communication about the speakers' actual (physical, social, and mental) world, and so the reasoning behind producing meaningful responses must connect the meanings of perceived inputs to information about that world" (2020). Because these statistical models have become deracinated from the world that produced the data to train them, their outputs necessarily lose the innate grounding of personal suggestions. And, more importantly, their training is based on the documented outputs of human culture that have been made legible via the internet and other datasets. The largely internal, mental process that an expert goes through to develop recommendations has not been digitized and therefore is not something

the models have access to. Put simply, once personal knowledge is removed from its human source, something meaningful is lost in the process.

These important limitations aside, there are aspects of the knowledge the AI systems have access to that are complementary to human expertise:

1. Data aggregation: AI can rapidly process and analyze large volumes of data from diverse sources, including travel reviews, booking patterns, and real-time information updates, at a scale far beyond human capability.
2. Global coverage: AI systems can understand and translate multiple languages simultaneously, allowing access to both local and international information about the destination.
3. Real-time information: AI can continuously monitor and integrate real-time data on weather, events, pricing, and availability across countless destinations, providing up-to-the-minute information.
4. Pattern recognition: AI excels at identifying subtle patterns and trends across massive datasets, potentially uncovering insights about travel preferences or destination popularity that might be invisible to human analysis.
5. Perfect recall: AI systems can instantly recall and cross-reference vast amounts of stored information, providing consistent and comprehensive knowledge without relying on impartial memories.

The following table demonstrates the way Conversational AIs can complement the unique human expertise of curating knowledge.

Table 2. Curating Knowledge: Human Expert and AI Complementary Strategies

Traveler needs	Human expert provides "personal insights"	Conversational AI provides "personalized information"
"I want to get a sense of what something is really like"	**Firsthand experience:** Humans can provide details and insider knowledge by drawing on their personal experience	**Aggregate for the specific:** AI can aggregate and analyze thousands of user reviews and experiences to provide a comprehensive overview of general impressions against a very specific query
"I want to discover unique recommendations"	**Uncharted gems:** Humans can suggest off-the-beaten-path options that are invisible online	**Trusted brands**: AI can identify lesser-known attractions or experiences validated by trusted sources
"I want to know what all the potential options are"	**Personal curation**: Humans present a wide variety of options and arguments within a curated range	**Comprehensive exploration:** AI can quickly compile an extensive list of all possible options from multiple sources and databases
"I want to learn more about the destination"	**Adapted to context:** Human experts adapt relevant facts and information to add depth and context	**Up to date news**: AI can provide comprehensive, up-to-date information on history, culture, climate, and current events, drawing from a vast range of continuously updated sources

Personalizing Interactions

One important expectation travelers demonstrated was that they would be able to share their travel needs discursively with the expert and receive highly personalized advice in response. We identified three strategies used by experts that made the interaction feel personalized, in response to three traveler needs.

Table 3. Personalizing Interactions: Human Expert Strategies and Tactics

Traveler needs	Human expert strategies	Human expert tactics
"I want to state my preferences in my own language"	Humans provide travelers space to state their preferences and explore them	Space for preferences
		Open ended questions and answers
		Referencing travelers' requirements
"I want a shortcut to highly relevant recommendations"	Humans ask tailored questions to address the depth of travelers needs and allows travelers to ask follow-up questions	Follow-up questions
		Assessing trip parameters
"I want to optimize the trip for me"	Humans aim to reduce friction in the itinerary	Maximizing the itinerary

A fruitful question to ask travelers following the advice session was "when did the expert not meet your expectations?". In highlighting negative experiences we learned what travelers hoped from the experience of talking to a travel expert.

For example, Yu grew frustrated with expert Ruby during their advice session because she didn't feel like she was afforded the space to truly express her preferences.

> Ruby (traveler): The conversation didn't flow as smoothly as I thought it would… I guess her answers are not open ended…you had to start another one… then start another one. I wanted it to flow from one topic to another.
> Interviewer: What could have made the conversation more flowing?
> Ruby (traveler): I think she should have asked me more questions after she's done answering so I can respond back. It should be a two way thing. But it was only me asking the questions.

Here Ruby is highlighting one of the key benefits of exploring a trip with another person: it is possible to have a back-and-forth conversation through which your requirements are explored, clarified and addressed. One of the key challenges of travel as an archetypal "singularity" (Karpik 2010) is its uncertainty. Travelers may have a general idea in their minds about what type of trip they want, but they don't fully know what's possible, nor whether those possibilities will meet their expectations in practice. As these possibilities emerge through an expert

conversation, they can visualize the trip more concretely and ask further questions, triangulating their preferences based on this feedback loop. As traveler Sheavon put it following their expert consultation, "Having a person there you get to narrow down your desires quicker. You got to do follow up questions that are answered instead of just generic websites leading you to ideas."

The most successful experts spent time at the beginning of the session asking multiple open ended questions to grasp the wider context of the traveler: who they were traveling with, past trips they enjoyed, typical activities they do on vacation, what lodgings they typically stay in. They then used this context later in the conversation to guide the traveler towards recommendations that fit with the picture they'd established, and, just as importantly, steer them away from the activities that didn't fit.

Experts explicitly referenced this context later in the session—like when Jeremy suggested Donny's kids would appreciate the Taco truck—to remind travelers they were taking their personal circumstances into account, and reinforce confidence that the item under discussion was ideal for them. They also used this context to suggest adjacent, complementary activities that might help to maximize the time spent in a particular place.

As Ruby intimates, the most successful advice sessions were not expert monologues, but those in which the traveler was afforded equal space to speak, provide feedback and clarify their own questions. This created a virtuous loop as their understanding of what was needed and possible converged through conversation. To explore and validate this phenomenon further, our data science team analyzed the extent to which both participants in the expert-traveler conversations asked a similar number of questions. This analysis further demonstrated that such conversations tend to result in travelers sharing more personal preferences, experiences, and a more positive sentiment towards the exchange. In the figure below, we highlight pair 7 where both expert and traveler are equally engaged in asking questions (both ask around 20% of questions across the duration of the interview which equates to 20 questions each in a 100-turn dialogue). When their dialogue is further analyzed using models trained for sentiment analysis and intimacy classification, the exchange proves even richer: the expert provides a safe space (measured by how positive their replies are) where the traveler admits to being a fan of visiting movie studios and consequently, receiving more personalized recommendations.

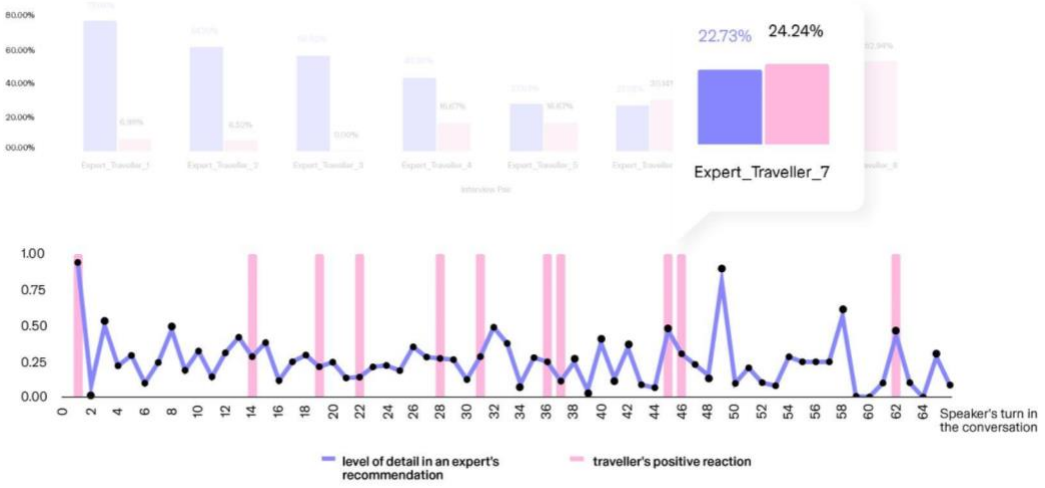

Fig 3. The first chart highlights the number of questions asked by the expert and traveler is similar (22.73%; 24.24%). The second chart shows that in the same pair there is a correlation between expert's probability of positivity (positive words used in sentences) and traveler's intimacy score (self-disclosure language used in sentences).

Personalizing Interactions: Implications for AI Systems

The personalized advice human experts provide is predicated, in part, on the social norms and cues inherent in interpersonal interaction. These norms help experts solicit information. When engaged, with intent, in a travel advice session travelers feel compelled to respond in detail to questions about their needs and context when asked directly for it by another person. Travelers also demonstrate a number of subtle cues about their engagement and interest levels through their body language, which the expert can use to "read between the lines" and interpret whether they are on the right track. A question remains around the extent to which a Conversational AI can successfully solicit and process the same fidelity of information.

Second, these social norms help skilled experts establish rapport. The most successful sessions evolved into meaningful conversations, filled with laughter, spontaneous asides and good will. When a traveler experienced rapport we observed that not only were they more likely to "open up" and share their underlying hopes and fears, but also they seemed likelier to view the advice being shared as being relevant and interesting to them. As we shall see in the final theme, "offering a perspective", because there is ultimately "no right answer" to travel, personalisation is as much about establishing confidence and trust around decisions as it is about algorithmically matching pre-existing preferences to trip items. As Forsythe explains "in everyday life, the beliefs held by individuals are modified through negotiation

with other individuals; as ideas and expectations are expressed in action, they are also modified in relation to contextual factors" (2001).

AI recommendation systems often face the challenge of creating interest and confidence in suggestions users are unfamiliar with. Data scientists call this the tension between "exploitation" and "exploration". Recommending something that is highly similar to another item that is already valued by a user is called exploitation. Exploiting known preferences provides effective personalisation in the short term, but the lack of novelty that arises can lead to repetitive experiences that erode the value of the recommendation system over time. Exploration, on the other hand, prioritizes new, unfamiliar items to a user that are at risk of not being valued at all, leading to a negative experience. But the upside of successful exploration is significant: the discovery of something invigorating and new expands personal horizons and makes the recommendation system more valuable to the user long-term.

If a Conversational AI can establish rapport like our experts, then it is more likely to convince them of these less predictable, higher value suggestions. It can explore as well as exploit. However, a big question mark remains whether a Conversational AI is able to establish this kind of rapport, both experientially (will the conversation be seamlessly intimate and responsive?) and philosophically (do travelers want to establish a relationship with a Conversational AI, and share personal contextual information with it?).

Research we've done to explore similar questions suggests that, currently, relationships with similar kinds of AI interfaces are more transactional: users want to cut to the chase when they're trying to solve a particular problem. However, there are elements that make a Conversational AI interface complementary to a human expert for personalized responses:

1. Availability and instant responses: access to travel advice anytime, anywhere, without waiting for human availability.
2. Comprehensive options: access to a huge database of trip items means that users can continue to ask for alternatives.
3. Effortless comparison and customization: easily compare multiple options and customize itineraries based on specific preferences, with the ability to quickly adjust parameters.
4. Privacy and reduced social pressure: explore travel options without the potential embarrassment or judgment that might come with asking certain questions to a human, and take as much time as needed to make decisions.

The following table demonstrates the way Conversational AIs can complement the unique human expertise of personalizing responses.

Table 4. Personalizing Interactions: Human Expert and AI Complementary Strategies

Traveler needs	**Human expert provides "meaningful conversation"**	**Conversational AI provides "infinite iteration"**
"I want to state my preferences in my own language"	**Seamless exploration:** Humans provide travelers space to state their preferences and explore them	**Unlimited prompting**: AI can enable travelers to take as much time as they need to share their thoughts, including across multiple languages
"I want a shortcut to highly relevant recommendations"	**Intuitive solicitation**: Humans ask tailored questions to address the depth of travelers needs and allows travelers to ask follow-up questions	**Endless alternatives**: AI has access to a huge database of options, enabling fast iteration to identify relevant trip items
"I want to optimize the trip for me"	**Context-specific suggestions**: Humans aim to reduce friction in the itinerary	**Unlimited filtering:** AI can integrate countless variables into suggestions to optimize itineraries

Offering a Perspective

Travelers appreciated experts who had a clear perspective, even when they didn't always agree with it. But they only listened to opinions when they trusted and related to where it was coming from. Experts exhibited three strategies that served to establish traveler trust in their opinions.

Table 5. Personalizing Interactions: Human Expert Strategies and Tactics

Needs	Strategies	Tactics
"I want to be able to contextualize recommendations so I can judge what is right for me"	Humans are upfront about their interests and biases to allow travelers to make their own decisions	Acknowledging biases
		Referencing shared interests
		Reflecting on suggestions provided
"I want to see the destination through a personal lens"	Humans let their personality and experience shine through to provide a unique lens on the destination	Personal storytelling
		Sharing opinionated suggestions
"I want to feel their excitement about the destination"	Human experts can share exciting content in an appealing format	Personal storytelling
		Positive and evocative language
		Personal destination photos

In a market in which there are "no right answers" people still seek a sense of certainty they are making the right decision. Which is why they seek the opinion of experts that they trust.

We noticed that the experts that travelers trusted the most tended to be those that they could personally relate to. When a traveler believed that an expert demonstrated a similar age, gender, interests and tastes (among other factors), they became more open to their opinions. We saw this process play out in the interaction between LA expert Jenny and traveler Christine, both of whom were female and of a similar age.

> Jenny (expert): I don't know if you've seen the photos of the Walt Disney Concert Hall. It's this amazing architectural piece. It's by Frank Gehry, who designed the Guggenheim in Bilbao. It's a really famous building, which is known for all the curvy architecture. And the concert hall in LA is really similar, so it's just like this amazing kind of metallic, silver, curved building. And even if you don't go to see a concert there, it's quite a cool building.

Subconsciously, or consciously, Jenny established credibility with Christine with these remarks. Not because of the specific recommendation in itself, but rather through what she revealed about herself in making the recommendation.

> Christina (traveler): Jenny mentioned Frank Gehry, who's an architect I know. I feel like most people don't know about him off the top of their head. So that caught my ear. I was like, okay, so I feel like I trust her recommendations more because she knows who that is. I feel like if she liked that, I can picture in my head what related work of his would look like. So maybe her recommendations are in a similar vein, style wise or visually, that would align with what I'm interested in.

By mentioning Gehry specifically, Jenny helped Christina contextualize her tastes in general. This not only made that specific recommendation more valuable, but had the effect of throwing all her other suggestions—before and after—into a different, more relevant light.

Conversely when the traveler didn't find the expert relatable, this undermined even potentially relevant suggestions. Expert Leroy, intuiting that traveler Beth was seeking "off the beaten path" suggestions, proactively disclosed his perspective "If you want to go to Disneyland, I'm probably not that guy, but if you want to know about unusual stuff I am." However, following the discussion Beth wasn't convinced:

> I couldn't really tell his age, but I'm assuming he's older than me. He still took into consideration that I'm younger, but then he mentioned concerts and that's pretty stereotypical. I feel like if I had someone younger, and if it was like a girl and we'd bond about shopping or something. And then she would give more of those inside scoop on where to find a vintage handbag.

Because Beth did not relate to Leroy, she read his recommendations as being "stereotypically for young people" like her, rather than emerging from his own personal experience. His opinions were therefore inauthentic.

Authenticity and self-disclosure were highly valued by travelers when conversing with experts. Our NLP analysis across conversations demonstrates that when experts shared their personal experiences and engaged in self-disclosure then travelers were more likely to use positive words in response. As seen in an earlier section, these inputs often take the form of positive feedback to the expert's advice such as "enjoying" or "looking forward" to the suggested experience made by the expert.

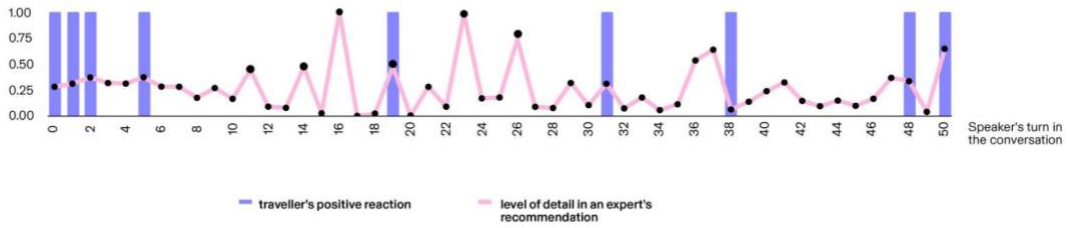

Fig. 4. The chart shows that there is a correlation between expert's intimacy score (self-disclosure language used in sentences) and traveler's probability of positivity (positive words used in sentences).

Successful experts also imbued these personal opinions with excitement, passion and positive language. Lewis, a traveler, loved Charlotte's recommendation of theme park Six Flags because of the passionate way she recounted her own experiences of the rollercoaster "it was absolutely insane!". Our NLP analysis further demonstrates how travelers respond positively when experts express excitement and passion.

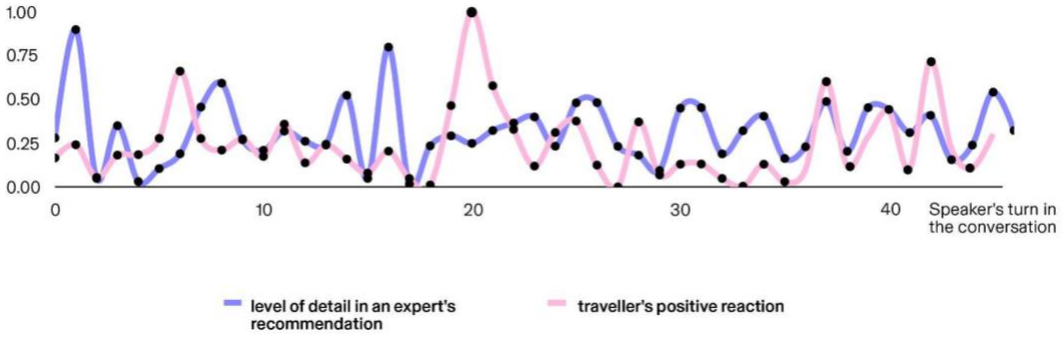

Fig 5. The chart shows that there is a correlation between expert's probability of positivity (positive words used in sentences) and traveler's probability of positivity (positive words used in sentences).

Offering a Perspective: Implications for AI Systems

The power of opinions is that they are singular; they represent the specific, authentic perspective of an individual. This situates recommendations, enabling them to be assessed in the context of their particular source. Travelers are able to pick up explicit and subtle cues about the expert to understand where the opinions are coming from, and triangulate them against their own tastes and preferences. As Forsythe summarizes: "what counts as knowledge is highly situational but is also a matter of perspective and cultural background" (2001).

It is true that Conversational AIs can represent specific characters, showcase particular characteristics, and reveal the sources of their information. But currently it is much harder for a traveler to situate themselves meaningfully against a Conversational AI because the sources of information and training are neither singular, nor embodied. Without a clear mental model of who or what they are talking to, travelers may find it hard to assess the perspective at hand.

Travelers experienced opinions as more powerful when they stemmed from the personal experiences of the expert. And if the traveler personally relates to the expert in question, then the opinion becomes even more compelling. In these moments the expert acts as a proxy, experiencing the destination vicariously on behalf of the traveler, reducing the feeling of uncertainty (Karpik 2010) in the process. Opinions that do not stem from personal experience are less authentic. Experts that did not naturally relate to the traveler in question had to defer to "second hand" suggestions, based on their general knowledge, or the experience of others. These might be highly relevant recommendations, but tended to be less trusted by the traveler.

Conversational AIs cannot experience a destination themselves, so their generated "opinions" are necessarily second hand and often a reflection of the data they were trained on. They can, of course, refer the traveler to examples of people who have experienced the destination in question (highlighting user-generated reviews, for example) but they cannot speak on behalf of those people, nor tailor the opinions of those people to the specific context of the traveler. This observation echoes the findings by NLP scholars Wu and colleagues who have highlighted that the contested reasoning capabilities of LLMs are vulnerable to degradation the greater the distance from the use case they are tested on is, as it is often the case for hyper-specific, personalized recommendations (2024).

Offering a perspective is, perhaps, the hardest type of human expertise to replicate. But there are strengths that a Conversational AI can leverage to complement and build on human perspectives:

1. Transparent source attribution: AI can clearly cite its sources for each opinion, allowing users to understand the origin and credibility of the information. This transparency can build trust.
2. Aggregation of diverse perspectives: AI can synthesize opinions from a wide range of sources, presenting a balanced view that incorporates multiple perspectives, which may be more comprehensive than a single human's opinion.
3. Personalization based on similarity: AI can identify and present opinions from travelers with similar preferences or demographics to the user, making the recommendations more relatable.

4. Quantitative backing: AI can support its opinions with relevant statistics and data analysis, adding a layer of objectivity to subjective recommendations.
5. Unbiased presentation: AI can present opinions without personal biases that might influence a human expert, potentially offering a more neutral starting point for travelers.

The following table demonstrates the way Conversational AIs can complement the unique human expertise of offering perspective.

Table 6. Offering a Perspective: Human Expert and AI Complementary Strategies

Traveler needs	Human expert provides "authentic beliefs"	Conversational AI provides "aggregated opinions"
"I want to be able to contextualize recommendations so I can judge what is right for me"	**Self-disclosure:** Humans are upfront about their interests and biases to allow travelers to make their own decisions	**Transparent sourcing:** AI can specifically reference the source on which its opinion is based upon
"I want to see the destination through a personal lens"	**Personal experience:** Humans let their personality and experience shine through to provide a unique lens on the destination	**User-generated content:** AI can surface relevant experiences from similar types of travelers
"I want to feel their excitement about the destination"	**Self-expression:** Human experts can share exciting content in an appealing format	**Aggregated experiences:** AI can demonstrate the relative excitement travelers-in-general feel towards something specific

Generalizing the Framework to Other Domains: Implications for Design and Engineering

For ethnographers exploring the value of AI within specific domains, our framework provides a useful starting point for clarifying the role Conversational AIs can play in relation to human experts. The implications of the framework span both the design of end user experience (what is increasingly termed AIUX), and the training and engineering of the underlying models themselves.

The following provides some specific principles that can be applied to any "no right answer" domain into which a Conversational AI interface is being deployed.

Knowledge: Deliver "Personalized Information" Not "Personal Insights"

Table 7. Knowledge Design Principles

Conversational AI design principle	AIUX / design implications	Training / engineering implications
Aggregate for the specific	**Do ask:** How can aggregated data address highly specific queries, without presenting it as a singular answer? **Don't ask:** How can the Conversational AI share specific, personal insights?	**Do ask:** How can personal experiences be better labeled to correspond to specific queries? **Don't ask:** How can we generate synthetic anecdotes or experiences that present a secondhand version of the real thing?
Leverage trusted brands	**Do ask:** How can we clearly attribute information to specific trusted sources while maintaining a seamless user experience? **Don't ask:** How can the Conversational AI curate recommendations itself?	**Do ask:** How can different sources be labeled to correspond to specific queries and user types? **Don't ask:** How can the Conversational AI access firsthand knowledge that hasn't been shared online?
Enable infinite exploration	**Do ask:** How can we provide a balance between curated recommendations and user-driven exploration? **Don't ask:** How can we arbitrarily create limited lists that reflect the limits of individual knowledge?	**Do ask:** How can different options be labeled for quality and relevance against specific contexts? **Don't ask:** How can we artificially expand the dataset to cover all possible scenarios without proper validation?
Surface live information	**Do ask:** How can live access to different feeds enable cutting edge, up to date information? **Don't ask:** How can we assume what live information is relevant to the individual without clear signals?	**Do ask:** How can we implement real-time fact-checking mechanisms to ensure the accuracy of live information? **Don't ask:** How can we prioritize recency over accuracy or relevance in information retrieval?

Interaction: Provide "Infinite Iteration" Not "Meaningful Conversation"

Table 8. Interaction Design Principles

Conversational AI design principle	AIUX / design implications	Training / engineering implications
Encourage additional prompting	**Do ask:** How can we design an interface that encourages free-form expression and prioritizes questions over answers? **Don't ask:** How can we replicate the situated understanding of cultural context that human experts bring to language interpretation?	**Do ask:** How can we improve NLP to accurately interpret preferences that are expressed colloquially? **Don't ask:** How can we program the AI to have personal travel experiences in different cultures to make it more relatable?
Leverage the signals	**Do ask:** How can we present a wide range of options while quickly narrowing down to the most relevant based on user interaction? **Don't ask:** How can we mimic the intuitive ability of human experts to read between the lines and infer unstated preferences?	**Do ask:** How can we optimize algorithms to rapidly filter and prioritize options based on user inputs and behavior patterns? **Don't ask:** How can we replicate the human ability to draw on experiences for nuanced recommendations?
Offer unlimited filtering	**Do ask:** How can we create an interactive interface that allows users to easily adjust and visualize different optimization parameters? **Don't ask:** How can we simulate the empathetic understanding that human experts use to balance conflicting preferences in trip planning?	**Do ask:** How can we develop algorithms that consider multiple variables simultaneously to create truly optimized itineraries? **Don't ask:** How can we program the AI to have the contextual awareness and flexible problem-solving skills of seasoned travel experts?

Perspective: Share "Aggregated Opinions" Not "Authentic Beliefs"

Table 9. Perspective Design Principles

Conversational AI design principle	AIUX / design implications	Training / engineering implications
Provide transparent sourcing	**Do ask:** How can we design an interface that clearly presents the sources and context for each recommendation? **Don't ask:** How can we replicate the human expert's ability to intuitively share personal biases and interests?	**Do ask:** How can we implement a system to accurately track and cite sources for all information used in recommendations? **Don't ask:** How can we program AI to develop its own interests and biases based on a particular perspective?
Match to user generated content	**Do ask:** How can we create a user interface that effectively presents and filters user-generated content based on traveler similarities? **Don't ask:** How can we mimic the unique personality and experiential narrative that human experts bring to their recommendations?	**Do ask:** How can we develop algorithms to match user profiles with the most relevant user-generated content? **Don't ask:** How can we give the AI its own personality and set of travel experiences to draw from?
Aggregate sentiments as well as information	**Do ask:** How can we design a way to present aggregated excitement levels visually or interactively to engage users? **Don't ask:** How can we make the AI genuinely feel and convey personal excitement about a destination?	**Do ask:** How can we develop sentiment analysis tools to accurately gauge and represent traveler excitement levels from various data sources? **Don't ask:** How can we program the AI to have authentic emotional responses to destinations?

Conclusion: Exploring the Space between "Autonomous Agent" and "Extended Self"

The debate about the future of AI computing interaction is divided between two visions. The first, articulated by OpenAI CEO Sam Altman, is that of AI as a "super-competent colleague that knows absolutely everything about my whole life, every email, every conversation I've ever had, but doesn't feel like an extension" (Varanasi

2024). This is the classic vision of AI as an agent with an independent identity and the capacity to act autonomously.

The second vision is articulated by Apple CEO Tim Cook. For Cook, the "Apple Intelligence" platform is about "empower[ing] you to be able to do things you couldn't do otherwise. We want to give you a tool so you can do incredible things." (Indian Express 2024). This is the more incremental vision of AI as an embedded feature that gives users new capabilities, in line with Steve Job's vision for computing as providing a "bicycle for the mind". In this view AI is experienced as an extension of the self, rather than as an autonomous, separate entity.

These distinct visions provide us with a useful spectrum for thinking about the future of expertise. On one end we have the OpenAI vision of the Agent and, on the other, the Apple vision of the Tool. Applying these two concepts to a conversational experience in travel, the Agent approach benefits from its autonomy, and is therefore perhaps better positioned to challenge, inspire and act on behalf of travelers as a human expert would. On the downside, the inevitable personification of such an Agent, especially when experienced through a conversational interface, may trigger expectations of human expertise, and the "personal insights", "meaningful conversations" and "authentic beliefs" that we have discovered people value. The fact that these expectations are unlikely to be satisfied for the reasons outlined above may create disappointment and confusion.

The more conservative vision for AI articulated by Apple will not trigger these unrealistic expectations, but could suffer from the opposite problem. If an AI is only experienced as an incremental Tool in the context of an existing product journey, then user expectations will be constrained by the pre-conceived limits of that specific context.

Technologist Matt Webb nuances the emerging AI landscape further by adding "Chat" as a third pole, defined as the capacity to "talk to the AI as a peer in real-time" (Webb 2024). For Webb, Tools, Agents and Chat are distinctive user modalities for interacting with AI, each enabled by a different aspect of the underlying technology. An AI that works alongside a user as a "Co-pilot" incorporates elements of all three modes.

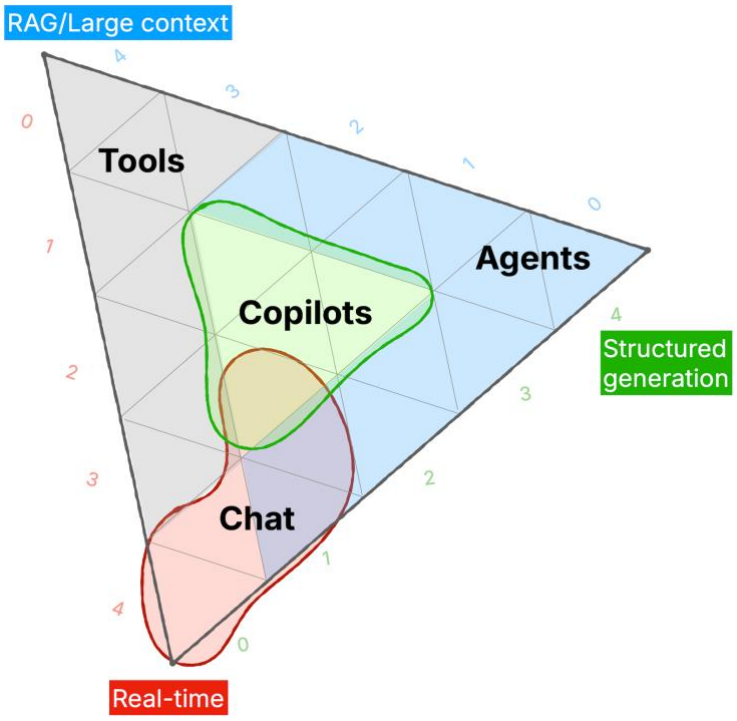

Fig. 6. Matt Webb distinguishes between different AI user experiences, enabled by three core capabilities.

What this increasingly complex picture surfaces is the lack of adequate mental models to help people successfully navigate these emerging modalities, and, by extension, what to expect from AI-driven expertise. This presents a significant opportunity for ethnographers. The travel case study has helped us identify the underlying, complementary qualities that AI can provide alongside human expertise. What is less clear, however, is how these complementary qualities should be experienced in domains outside of travel, including through what modality.

The design principles outlined above are useful starting points for product teams seeking to build new conversational experiences across consumer domains, especially those for which there is "no right answer." However, our experience is that the implications for user experience (AIUX) and engineering will be specific to the context in question. This requires teams of ethnographers and data scientists (using combined methods like those outlined in this paper) to carefully explore the cultural and behavioral nuances of the domain they are designing for, and convert this understanding into outputs that are legible to engineers. As with any type of product, the difference between success and failure will be defined by the value it provides to the end user, not the raw potential of the underlying technology.

About the Authors

Tom Hoy is Partner and co-founder of Stripe Partners. His expertise lies in integrating social science, data science and design to unlock concrete business and product challenges. The frameworks developed by Tom's teams guide the activity of clients including Apple, Spotify and Google. Prior to co-founding Stripe Partners, Tom was a leader in the social innovation field, growing a hackathon network in South London to several hundred members to address local causes.

Janneke Van Hofwegen is a UX Researcher on the Real World Journeys team on Google Search. She employs ethnographic and quantitative methods to learn how people are inspired by and decide to embark on experiences in the world. She holds a PhD in Linguistics from Stanford University. Her linguistics work focuses on understanding human behavior and stylistic expression through the lens of language and group/individual identity.

References Cited

Bender, E., & Koller, A. 2020. Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 5185-5198). Online: Association for Computational Linguistics.

Buchanan, B. G., R. Davis, and E. A. Feigenbaum. 2006. "Expert Systems: A Perspective from Computer Science." In *The Cambridge Handbook of Expertise and Expert Performanc*e, edited by K. Anders Ericsson, Neil Charness, Paul J. Feltovich, and Robert R. Hoffman, 87-103. Cambridge: Cambridge University Press.

Chang, L. 2023. "Most Visited American Cities." Condé Nast Traveler. Available at: https://www.cntraveler.com/story/most-visited-american-cities.

Confente, I. 2015. "Twenty-Five Years of Word-of-Mouth Studies: A Critical Review of Tourism Research." *International Journal of Contemporary Hospitality Management* 17(6): 613-624. https://doi.org/10.1002/jtr.2029.

Dreyfus, H. L. 1987. "From Socrates to Expert Systems: The Limits of Calculative Rationality." *Bulletin of the American Academy of Arts and Sciences* 40 (4): 15-31.

Forsythe, D. E. 2001. *Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence.* Stanford: Stanford University Press.

Frey, C. B., and M. A. Osborne. 2013. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114: 254-280.

Harnad, S., 1990. The symbol grounding problem. Physica D: Nonlinear Phenomena, 42(1-3), pp.335-346.

Hawkins, A. J. 2024. "Google AI Gemini Travel Assistant Now Handling Hotel Bookings." The Verge. Available at: https://www.theverge.com/2024/5/14/24156508/google-ai-gemini-travel-assistant-hotel-bookings-io.

Herbold, S., A. Hautli-Janisz, U. Heuer, Z. Kikteva, and A. Trautsch. 2023. "A Large-Scale Comparison of Human-Written Versus ChatGPT-Generated Essays." *Scientific Reports* 13 (1): 18617.

Hoy, T, Bilal, I.M, and Liou, Z. 2023. "Grounded Models: The Future of Sensemaking in a World of Generative AI." *EPIC Proceedings* 2023: 177-200.

Karpik, L. 2010. *Valuing the Unique: The Economics of Singularities*. Princeton, NJ: Princeton University Press.

"Tim Cook Explains Why Apple Calls Its AI 'Apple Intelligence'." 2024. Indian Express. Available at: https://indianexpress.com/article/technology/artificial-intelligence/tim-cook-explains-why-apple-calls-its-ai-apple-intelligence-9389541/.

Jacobsen, J., and Munar, A.M. 2012. "Tourist Information Search and Destination Choice in a Digital Age." *Tourism Management Perspectives* 1: 39-47.

Lowery, A. 2023. "How ChatGPT Will Destabilize White-Collar Work." The Atlantic. Available at: https://www.theatlantic.com/ideas/archive/2023/01/chatgpt-ai-economy-automation-jobs/672767/.

Mountcastle, B. 2024. "Using ChatGPT for Travel Planning." Forbes. Available at: https://www.forbes.com/advisor/credit-cards/travel-rewards/chatgpt-for-travel-planning/.

Pierce, D. 2023. "The AI Takeover of Google Search Starts Now." The Verge. Available at: https://www.theverge.com/2023/5/10/23717120/google-search-ai-results-generated-experience-io.

Varanasi, L. 2024. "Sam Altman Wants to Make AI Like a 'Super-Competent Colleague That Knows Absolutely Everything' About Your Life." *Business Insider*. Available at: https://www.businessinsider.com/sam-altman-openai-chatgpt-super-competent-colleague-2024-5.

Webb, M. July 19, 2024. "The AI Landscape." *Interconnected.* Available at: https://interconnected.org/home/2024/07/19/ai-landscape.

Wu, Z., L. Qiu, A. Ross, E. Akyürek, B. Chen, B. Wang, N. Kim, J. Andreas, and Y. Kim. 2024. "Reasoning or Reciting? Exploring the Capabilities and Limitations of Language Models Through Counterfactual Tasks." In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (Volume 1: Long Papers), 1819-1862. Mexico City, Mexico: Association for Computational Linguistics.

“I Pick My Poison”: Agency and Addiction in the Age of Subscriptions

IVETA HAJDAKOVA, Stripe Partners
MORGAN WILLIAMS, Stripe Partners

This paper illuminates key aspects of the subscription economy and offers a framework for understanding customer decision making and willingness to pay. Subscriptions are aimed at becoming embedded in our lives, habits, and behaviours, and our in-depth research explored how people see themselves as agents in the subscription economy, how they value subscriptions, and how they navigate a complex market characterized by ever-changing bundles of features and deals. We argue that a new kind of customer agency emerges from the complexities and uncertainties of the subscription economy—one that recognises dependencies and even addictions to subscriptions and navigates their activity and passivity to produce desired outcomes. The paper describes 4 key factors that play a role in people's evaluations of subscriptions. Building on the work of Michel Callon, we also describe subscriptions as attachments and reframe 'willingness to pay' to better reflect the complexity of subscriptions. The paper is based on self-funded research (not commissioned by any client) conducted by Stripe Partners to improve our approach to willingness-to-pay studies.

> It's very difficult to be someone my age in this day and age and not be addicted to something. I just pick my poison and I would rather watch a TV show than get sucked into social media. (Archana, 21, US)

In the summer of 2023, Stripe Partners conducted self-funded (not commissioned by any client) research with 12 participants across the US and the UK into the role that subscriptions played in their lives to help our data science team innovate willingness to pay studies. Our goal was to produce a framework for understanding people's decision-making around subscriptions and to identify factors that impacted their willingness to pay. Across paid and free subscriptions—from streaming services to digital storage and meal boxes—we were interested in how people thought of the value of subscriptions and how they navigated the market and the ever-changing bundles of features and deals.

In our conversations about subscriptions, participants talked about their passivity and dependency on subscriptions and their unconscious behaviours, such as forgetfulness, laziness, poor money management, or being prone to seek dopamine hits. Our youngest participant, 21-year-old medical student Archana from the US,

Hajdakova, Iveta, and Morgan Williams. "'I Pick My Poison': Agency and Addiction in the Age of Subscriptions." EPIC Proceedings (2024): 244–253. https://www.epicpeople.org/agency-and-addiction-in-subscriptions. ISSN 1559-8918.

understood her passivity as a form of addiction, and saw her agency limited to "picking her poison":

> It takes so much longer on Instagram for me to get enough of that dopamine hit to stop versus these streaming services, they provide that so much more quickly and, to an extent, yeah, okay, maybe I am addicted to streaming platforms or whatever, but I think it's very difficult to be someone my age in this day and age and not be addicted to something. I just pick my poison and I would rather watch a TV show than get sucked into social media, ... I might be scrolling on Instagram, I might see some things that are entertaining, but I'm also going to see a lot of like gosh, climate change is killing our planet, gosh Trump is ruining our democracy. Streaming provides a little bit of escapism in a way that social media and YouTube cannot. So if I have to be addicted to something, I would rather be addicted to streaming services.

Archana was not alone in recognizing her dependency. Other participants talked about being "hooked" or needing to "break up" with their subscriptions, reflecting on the habits they had formed. Our inquiry into their willingness to pay for various subscriptions revealed a subjectivity characterised by its awareness of their behaviours and habits, and by the need to make conscious decisions about their unconscious habits. This subjectivity emerges from the complexities and uncertainties of the subscription economy, under which more and more services are aimed at becoming embedded in people's lives, habits, and behaviours.

In this paper, we aim to explore how the subscription economy impacts people's agency and subjectivity, how they see themselves as agents in these new markets and how that affects the value they attribute to various subscriptions. We bring together several different perspectives and "schools of thought":

1. We look at subscriptions from a business perspective, aiming to understand how digitally driven subscription models bring value to customers and how customers establish this value.
2. We draw upon theory and especially the work of Michel Callon whose analysis of the relationship between market agents and markets helps us unlock how people think of the value of subscriptions.
3. We bring a critical perspective to shed light on how the subscription economy impacts customer agency.

Bringing these different perspectives together will help us make sense of the inherent uncertainty of subscriptions and decision making around them and better understand the current phase in the subscription economy and the impact it has on customers.

The Subscription Economy

While subscriptions are not new, the subscription economy has been enabled by digital technology and fueled by vast amounts of personal data. This has led not only to a wider socio-economic shift from ownership of products to access to services, but also to new types of relationships between customers and service providers, in which personal data is used to generate value across multiple channels and surfaces, enabling all aspects of life to not only be measured and ranked, but also predicted (Fourcade and Healy 2024). The resulting relationship we have with the digital services we use impacts how we see ourselves, how we conduct ourselves in our daily lives, and how we navigate our aspirations and futures.

In his 2018 book *Subscribed,* Tien Tzuo explains that the subscription economy represents a transformation in the way business is transacted. The defining characteristic of the subscription economy, Tzuo argues, is the shift from the product-centric business model to a customer-centric business model (Tzuo and Weisert 2018). As a result of this transformation, says Tzuo, a new type of consumer emerges, one who prefers access to services to ownership of products. To cater to this new consumer, businesses need to leverage new tools, build new knowledge about their customers and sustain the process and development of a customer relationship over time (Tzuo 2015).

Tzuo's framing of the subscription economy summarises how many businesses currently think about their products and services and the relationships they want to build with their customers. The subscription model underlies many business tools aimed at personalisation, retention, localisation, customer journeys and monetisation. At the heart of these new customer relationships are mechanisms that help attract and retain customers. While these mechanisms are not unique to subscriptions, they become more urgent for both businesses and customers, because subscriptions require more ongoing interactions with digital interfaces, often leveraging and generating vast amounts of data.

Fast forward to 2024, Tzuo writes:

> But a year ago, a new phrase entered our lexicon: subscriber fatigue. Consumers are overwhelmed with the multitude of subscriptions they're expected to maintain. Some businesses discovered that they have more SaaS applications than employees. Everyone is reassessing the value they get from these products. Profits are down. Churn is on the rise.

The Three Phases in the Subscription Economy

According to the 2022 FT Strategies report, we have entered a third distinct phase in the subscription economy:

> Starting in the early 2000s, the first phase took around 20-plus years in which digital subscriptions gradually took hold as a strategic lever for businesses across most industries, from Media and Lifestyle, to Financial Services, Gaming, and many more. It was followed by a second phase of marked acquisition growth during the COVID pandemic, which helped spike the already healthy consumer appetite for subscriptions. We are now entering a third phase, or what we see as the retention phase (Adelman 2022, 3).

The report explains that the current phase is characterised by heightened consumer awareness driven by the economic downturn and the cost of living crisis. While consumers still value subscriptions, the expectation is that they will be more reserved and likely jump between subscriptions (Adelman 2022, 3). Tzuo shares the concern. In his 2024 article, he writes:

> It's not looking too good for subscriptions right now: news sites are dropping their paywalls, streaming services are turning to ads (and into cable), and there's even noise of SaaS companies abandoning the model.
>
> What is going on?

Yet, new services, products and experiences are still being offered as subscriptions or memberships, including, for example, restaurants (Goldfield 2024) or financial services that help users consolidate their subscriptions (e.g. Rocket Money or Finanz Guru).

In this current context, businesses need to think of new ways to retain customers and new monetisation and pricing strategies. Businesses want to understand their customers' willingness to pay, however, as our paper illustrates, subscriptions are

complex and uncertain products, and from the customer perspective, their value is difficult to establish.

The New Business Model

As business models, digitally delivered subscriptions offer unique benefits to both businesses and subscribers: they offer businesses a predictable revenue stream and deeper customer insights, while providing consumers with convenience, cost savings, or access to a broader range of products and services (Ughade 2024).

Tzuo explains that the old product-centered business model was aimed at selling more products and this was achieved by getting the products into as many sales and distribution channels as possible (Tzuo and Weisert 2018, 19-20). The new customer-centric business model represents a shift from these "linear transactional channels to a circular, dynamic relationship" with the customer. Technology enables companies to identify various channels across which they can meet their customers' needs as they emerge (Tzuo and Weisert 2018, 20).

An ongoing relationship also has implications for pricing and revenue. According to Tzuo, the traditional business model does not differentiate between recurring and nonrecurring revenue (Tzuo and Weisert 2018, 179), whereas it is precisely recurring revenue that is the cornerstone of the subscription economy.

The new business model enables powerful relationships between businesses and their customers by leveraging various mechanisms—social, financial, technical and behavioural—that foster engagement over time. At the core of these mechanisms is data. Data can be used for many purposes, including targeted marketing, ads, or segmentation. Many subscription models are enabled and empowered by large amounts of personal data harnessed and leveraged at scale. The data is used to provide better personalisation and create more powerful algorithms, which creates value for the individual customer as well as for the wider customer base. However, this data also generates new forms of capital for businesses, advertisers and third-party data providers.

When coupled with digital products and personalisation, subscriptions, such as for streaming platforms, benefit from vast and engaged user bases who generate data used to provide better recommendations. In this context, however, the relationship between the individual user and a large digital platform is an asymmetrical one, albeit often disguised as reciprocal. As Fourcade and Healy argue, in these relationships,

the role and value of personal data is obfuscated and users' ability to opt out is limited (2024, 43-44).

Tzuo's model doesn't acknowledge the complexity of the data economy that underlies the subscription economy and the issues associated with it, as they have been powerfully documented by Shoshana Zuboff (2019) and plenty of other authors, who have pointed out the connection between transactional data enabled by electronic payment systems and surveillance (e.g. Fourcade and Healy 2013; Lauer 2020). As a result of this connection, the relationship between finance and information has been reversed (Fourcade and Healy 2024, 136): transactions are now less about money and more about the data trails they leave (Maurer 2014) and the monetary record becomes "a proxy for the intimate secrets and desires of its users" (O'Dwyer 2019, 12).

The asymmetrical relationship that exists between an individual customer and digital service is often tainted with suspicion regarding personal data that is being collected and how it might be used. At the same time, customers find themselves captivated and engaged by the mechanisms that are designed to meet their needs across many channels and moments.

The Agencement Model of the Market

The mechanisms that produce ongoing attachments and dependencies are of a particular interest to Michel Callon. Callon's work *Markets in the Making: Rethinking Competition, Goods, and Innovation* challenges the fundamental distinction between agents and goods that underlies the established thinking about markets as encounters between agents with needs and desires, and goods with specific properties that address (or don't) these needs and desires (Callon 2021, 46).

Callon introduces the agencement model of the market which sees goods as processes that go through different trajectories throughout their "careers" (c.f. Appadurai 1986) as they are designed, produced, commercialised, attached to customers, and as they become singularities—the value of which can be established and experienced by the customer. Throughout these "careers", value is created collectively and by multitudes of heterogeneous agents, tools and devices. Callon's neologism, agencement,[1] refers to assemblages that act as systems to accomplish various courses of action (Callon 2021, 140).

According to Callon's model, agents and goods are constitutive and transformative of each other. This is particularly obvious in the case of digital

products that utilise data to offer more personalised experiences. For example, an Amazon book recommendation that appears after one has bought a book can shape the customer's understanding of themselves as it reflects their tastes back to them, telling them something about themselves they didn't know before. Subscriptions are great examples of how goods and agents come together and transform each other. The attachments that are formed between customers and goods in these relationships are often long-term and therefore shaped by various factors that impact not only their perceived value but also their respective trajectories. For example, a music fan feels transformed by music they've discovered through Spotify, but they've also transformed their Spotify algorithm and provided data that informs the algorithmic recommendations for others.

The capacity of a good to create an attachment is what makes Callon's agencement model distinct and novel as it acknowledges that goods play a role in the structuring and dynamics of market activities (Callon 2021, 255). Callon points out that "if agents get attached, or are captured by goods, it is at least in part because the goods are made to be … captivating" (Callon 2021, 255). Anyone working in business understands that goods are, indeed, made to be captivating. However, Callon's point is that this power of goods to be captivating comes together with the customer's agency and ability to be captivated. Through market encounters, both goods and agencies (customers) are transformed as they become attached.

The idea that digital subscriptions are utilising various devices to foster long-term attachments might not be surprising to the reader. However, Callon's analysis of agencements reveals the shortcomings of our approaches to willingness-to-pay studies that assume that an agent is able to make an informed evaluation of how much they're willing to pay for their subscriptions and even a prediction about how much they'd be willing to pay should some features of the subscription change.

The Problem of (E)Valuation: Framing and Overflowing

As subscriptions are so embedded in people's lives, relationships, finances, routines and other subscriptions, they are difficult to disentangle, compare, evaluate or review. Even in corporate or institutional contexts where formal processes are established and quantitative rankings or usage statistics are available, decisions around subscriptions are embedded in relationships and social contexts (Ivanov et al. 2020).

To establish the value of such attachments requires complex informational devices that are both quantitative and qualitative in nature. While price is a quantitative metric, as Callon notes, it also carries qualitative and symbolic meaning. For example, the price of a wine bottle indicates its place in the wine hierarchy, just as subscription names like "Premium" or "Plus" reflect similar hierarchies.

From a business perspective, subscription pricing is complex and ambiguous. As Tzuo explains, "at the end of the day, you're not pricing an object, you're pricing an outcome" (Tzuo 2018, 152). Outcomes are important for customers as well but how do they anticipate whether an outcome will be worth the costs over time?

Evaluation doesn't happen inside a customer's mind.[2] The key to Callon's understanding of market dynamics and evaluations is the role of framings.[3] Market agencements, explains Callon, "are structured by framings that format the courses of action" (Callon 2021, 362). Callon uses the example of a dating app: there are many qualitative and quantitative points of information that a user evaluates when they see another user's profile. How they frame that encounter and what they value formats their action. However, "[a]ny framing produces overflowing," (Callon 1998, 38) meaning that over the course of evaluation and engagement, the user reframes the encounter—perhaps they might frame the potential match based on their pictures and immediately reframe it based on their education or religion. Enabled by tools and information, this dynamic of framing and reframing is the condition for any attachment to form.

Now consider a person who is trying to evaluate a dating app subscription that provides them with access to many more profiles—how does one decide on the value of such a subscription? How might getting that subscription impact how they evaluate individual profiles and how they interact with the app over time?

Subscriber Fatigue or a New Form of Customer Agency?

In their book *The Ordinal Society*, Fourcade and Healy analyse the intersection of digitally enabled products and services and a broader context of surveillance, and how their coming together impacts society and individuals. The abundance of data enables the proliferation of rankings and ratings, which impact all aspects of our lives, everyday decisions and future prospects. These rating and ranking systems support better services, for example by ranking Uber drivers or potential matches on dating apps, but they also enable more addictive experiences (from video games to social media scrolling), they impact one's ability to use services such as insurance or

credit, and they lead to more surveillance (Fourcade and Healy 2024, 261-270). This ordinal society of data, metrics and rankings, impacts how customers are seen as agents, as well as how they see themselves.

Engagement is often understood as a key signal of a customer getting some kind of value out of a service, and vice versa, engagement generates valuable data about tastes, preferences, spending as well as everyday behaviours. To initiate and boost engagement, various strategies informed by the knowledge of human behaviour are used, such as free trials, social experiences, or nudging.[4] In the effort to get as many people to subscribe to a service or prevent them from unsubscribing, many businesses have used "dark patterns", such as difficult cancellation processes, confusing presentation and communication of different tiers, prices and billing options (Sinders 2023).

At the heart of these strategies is an understanding of customer engagement with technology as both active and passive at the same time—active in the sense of being affected by a product and willingly interacting with it while actively producing data, and passive in the sense of being captivated and made dependent on it. Nudging, for example, is intended to incentivise agents to adopt and maintain certain behaviours without questioning them (Callon 2021, 155).

However, customers are often at least vaguely aware that with every transaction and interaction, data is collected. Sometimes, this data is fed back to them directly, in the form of various statistics that help them learn something new about themselves. They are able to reflect on their own passivity and take action to change their (passive) behaviours.

We go back to Callon and his reminder that goods and agents transform each other. The complexity of subscriptions and the difficulty of establishing their value shapes how customers make decisions, it shapes the framings that format their attachments, and the factors they consider when establishing the value of their subscriptions.

We examine how this new type of agency that considers its passive attachments and even addictions, evaluates subscriptions and we suggest how to reframe the question of willingness to pay. To understand what Tzuo calls "subscriber fatigue", we need to examine how the market encounters with subscriptions transform agencies as well as markets, creating new alignments and misalignments between customers and businesses. Addressing them can help us identify new opportunities.

Methodology

Our research was unusual in that it was not commissioned by a client but instead, it was initiated by our data science team. Our objective was to explore various factors that impacted how people valued subscriptions and how these findings could improve Stripe Partners' approach to willingness-to-pay studies and provide a deeper and broader perspective on the subscription economy and how it might evolve.

We conducted 2-hour remote interviews with 6 participants in the US and 6 in the UK. Our sample included 5 women, 6 men and 1 non-binary person, aged between 21 and 67, skewing younger. We recruited participants with varying levels of income and various numbers of paid subscriptions (from 1 to 10+).

In our interviews, we explored attitudes towards subscriptions and their perceived value, without any strict definition of what a subscription was. We probed participants around different types of subscriptions they were using for free or were paying for, including shared subscriptions and subscriptions paid for by others.

The Problem with Willingness-to-Pay

Willingness-to-pay (WTP) exercises are typically used in two situations: 1. to decide which of a set of possible directions to pursue (because demand for that direction is greater than for alternatives); or 2. to help set prices that maximise profits. For WTP exercises to be informative they need to provide results that are as reflective as possible of what real world behaviour might look like.

The premise of WTP studies is the idea that a customer is able to identify their needs and desires (demand), evaluate their options (supply) and establish how valuable something is to them and how much they're willing to spend to acquire a particular good or service. A qualitative perspective can then help understand relevant factors that are important to people when considering and evaluating the different options and deciding whether or not to pay for them.

The assumption behind WTP exercises related to subscriptions is that the qualities of an offered subscription would impact customers' demand, and vice versa, that understanding customers' unmet needs and desires could be used to improve the offer.

However, we want to show that the current assumptions that inform WTP approaches are not sufficient in helping us understand how people go about the calculations that inform their decision-making and evaluations regarding

subscriptions. The question of willingness to pay in the context of the subscription economy revealed fundamental problems of our understanding of markets and market agencies.

Subscriptions as Complex Attachments

Subscriptions are complex products that offer various types of content, services, features and entitlements, many of which participants didn't use or weren't aware of. While a single subscription can offer multiple benefits, participants would frame its value through their own individual use cases and goals. For example, some users of Amazon Prime saw it as a shopping and delivery subscription while others valued it primarily for entertainment.

Subscriptions were also recognized as complex because they were often bundled together or came as benefits to other subscriptions. Frank (67, US), had about 10 TV subscriptions, many of which were bundled together with other services, such as broadband or Walmart.

> We have them separately, but they're as a result of a promotion. So for an example, Paramount Plus was given to us for free as a result of joining Walmart Plus…. And Peacock was the result of a similar promotion. I don't remember exactly what we joined or what was offered to us, but we pay 99 cents a month for the next year and a half or something for it, so why not have it?

Participants often had multiple subscriptions with overlapping features and benefits, which impacted their awareness of what they were paying for, were entitled to and how they could manage the service. Yolanda, a 48-year-old single mother from New York City, was so anxious when it came to her spending that she did not even want to pay her bills automatically and checked her bank accounts every day. She had recently "broken up" with a lot of her subscriptions, the hardest of these breakups being Amazon Prime. However, during our conversation, she remembered she had subscribed to a toilet paper subscription on Amazon Prime during the pandemic. During the interview, she went on Amazon and realised she still had the subscription:

> Oh, I missed that. Missed this breakup. Um, yeah. So I guess I did do that. Wow.

Yolanda's resistance towards subscriptions was a result of her recognising that she had formed attachments with her subscriptions, which she wasn't able to stay on top of. She understood that companies could exploit people's forgetfulness. This was also important to explain to her son, who was very keen to get subscriptions but would often forget to manage them.

> It's the same thing that my son gets caught up in. … Not thinking about it three months from now. And then they're billing me. So it's their banking on you forgetting and them being able to bill you.

Subscriptions Are Recognised as Business Models That Leverage Customers' Passivity

While subscriptions are difficult to define, participants recognised them as business models that were designed in ways that leveraged their vulnerabilities, especially their tendency to develop dependencies, their forgetfulness and their lack of awareness over the personal information they shared. Participants thought that subscriptions exploited their passivity and therefore had to be actively, or pro-actively managed.

The youngest participant, Archana (21, US) was aware of the difference in attitudes towards subscriptions between herself and her parents who were a lot more "distrustful" than her because they were not familiar with how subscriptions worked. In contrast, she felt she understood the business model and was able to benefit from it.

> I don't think my parents would ever consider a meal kit service. They would think that's dumb. And they're also very distrustful of subscriptions. Every time I'm like, oh, I'm gonna have four months of this, they're like, 'they're gonna charge you for the rest of your life.' I'm like, 'it's not that hard to cancel'. So there's an attitude of, like, why would I buy a subscription when I can just buy the item? And an attitude of, be very careful what you subscribe to, because it's very easy to get tricked… Their concerns are valid, but I think I trust myself to navigate subscriptions more than they trust themselves. Like, I trust myself to be able to catch that and act accordingly.

For Archana, there were both pros and cons with regards to the subscription business model. On the one hand it reduced the cognitive burden of decision-making or problem-solving and could help habitualise certain behaviours, but on the

other, it might mean having a suboptimal pairing of goods with needs as one's circumstances changed.

> The pro is that something that you want, something that you like, you don't have to make the conscious effort to get it again and again. It'll just always be there for you. But the con is that it's hard to evaluate whether you still need something or, like, you end up paying for things you may not be using or need as much as you think you do. Because once you make that decision to start having the service, you're not reevaluating every month—do I still need this? Do I still want this?

Many participants recognised the different mechanisms that were used to get them "hooked" on subscriptions. For example, many felt like specific TV shows, exclusive sports content or free trials were used to get people to get a new subscription. Participants had to develop strategies for working around these mechanisms, such as setting reminders for cancelling subscriptions, re-evaluating their expenses or using gift cards to control their spending.

> You know when you do the free trial and you forget about it. I did that with Now TV and I was really angry because there was literally nothing else to watch on it. I didn't enjoy it at all and now I always set an alarm if I do that to make sure that I unsubscribe and don't disappoint myself (Aaron, 42, UK).

All our participants saw subscriptions as potentially risky attachments[5]—they were risky because of the different mechanisms that were orienting them to behave in certain ways they weren't necessarily aware of. The mechanisms were an important characteristic of this type of commodity but if managed carefully and used in a way that would generate value, they were seen as convenient and even beneficial.

With his 10+ subscriptions, Frank (67, US) was conscious about managing the subscriptions he was paying for. He had seen a YouTube video that offered advice on how to be smart about subscriptions by pausing or cancelling them:

> In order for you to save money with all these subscriptions, realise that you're really only on a month-to-month basis. So if you have a particular show that you like and it's ended and you really don't want to see anything else, just drop the subscription …Suspend the subscription until the next season comes on.

In order to get the desired outcomes, participants considered how a subscription business model worked and how it might align with their active and passive behaviours and decisions over time. Within this dynamic and fluid picture of subscription evaluations, we identified the following factors that facilitate such an evaluation: 1. Desired effects, 2. Temporality, 3. Affordability and 4. Engagement.

The 4 Factors

1. Desired effects

Our participants expected subscriptions to have effects on their lives. Subscriptions played an important role in supporting their lifestyles, relationships, everyday habits and practices, as well as their goals and aspirations.

In general, the most valued subscriptions were for entertainment. Many participants thought it was essential to have something to watch as this provided them with talking points but also was something they often shared with others. Many of Frank's subscriptions were for his extended family, including his brother who was in Japan, his grandchildren who'd come and visit and he even shared subscriptions with friends. He thought it was a "nice little thing to be able to give access to someone."

Entertainment provided important effects for many participants. It helped them decompress or escape and provided relief from busy lives. Nana considered TV show subscriptions to be essential:

> Especially, in a place like New York, I feel like a lot of people do have to work, so when you want to decompress, you don't want to be limited.

Participants expected to be attached to their subscriptions—a subscription that has had an effect on their lives would be considered most valuable. For Iman (37, UK), her gym membership was "invaluable" to her. With a baby daughter and a busy schedule, her gym membership not only provided her with the opportunity to go swimming, but it also incentivised her to take the time for herself:

> Pure Gym is something that I know that if I get rid of it, then there's nothing holding me accountable to be able to go to the gym and because I enjoy swimming, it's just something that I do for myself. So that's why I do it. It doesn't really hold value, really because, maybe it's gonna sound really cheesy, but it's, like, invaluable.

Another young mother, Isoke (34, UK) used the Calm app because it had effects on her mind and calmed her down:

> Especially when you have managed the kids for the whole day, you just really need some peace. So, you know, it brings that peace, that calmness, really. Like it revives me, it re-energises me.

Participants valued when subscriptions became incorporated into their lives, daily habits, rhythms and practices. They valued personalised experiences because they helped them achieve their desired effects without effort, allowing them to be passive. Nana (26, US) defaulted to Netflix when he didn't know what to watch. Netflix provided the desired effects while also enabling Nana to offload the cognitive effort needed to make decisions, which, Nana thought, has changed the way people watch TV more broadly:

> So when I sit down to eat my lunch and I don't have any TV show to watch, I automatically go to Netflix because they're the goat. They're, like, dope. They're the ones that got people hooked on this way of watching stuff, streaming.

Paid subscriptions were expected to deliver the desired effects in a more predictable and straightforward way compared to free alternatives or alternatives with advertising. This is why Archana (21, US) picked streaming services over social media (see quote above) and why she paid for her Spotify subscriptions as opposed to having to listen to ads:

> On Spotify, the ads are disruptive. If I'm about to fall asleep and I get a loud, raucous ad, then I'm awake.

Participants recognised that not all subscriptions became part of their daily habits, but they still were valuable to them because they represented their aspirations and provided motivation. Jane (26, UK) talked about how Headspace has helped with her anxiety and she still kept paying for it to keep up the habit and also help others:

> I do feel like I should just continue to have it and to use it, especially if it's something that I'm recommending to people.

Although Headspace was not something she was actively using, she felt like it was part of who she was—a mental health professional. In this way, Headspace was part of Jane's identity and habitus[6]: the choice to keep it wasn't determined by how

often she was using it to regulate her mental health. By recommending Headspace, Jane was able to impact other people's mental health, and, in return, gain authority as an expert.

As participants adopted new subscriptions, these would impact the value of existing ones. Thus, they were often negotiating and reprioritising the different desired effects they wanted to achieve and the subscriptions that help them achieve them. Iman (37, UK) had been subscribed to The Good Housekeeping Magazine since she was young and felt like it was "part of her". She kept the subscription because reading the magazine provided a detox from her phone:

> Like, sometimes, you know, when I don't want to spend time on my phone and things like that, I open the magazine and I spend some time in the evening where I'm not actually on the phone and I'm actually just reading the magazine.

As people's lives change, this impacts how they get the effects they desire from subscriptions. Some subscriptions remain incorporated in their lives, while others may be forgotten, paused or cancelled, or incorporated into new relationships.

2. Temporality

Temporality plays a crucial role in the effectiveness and perceived value of subscription services, as the benefits derived from subscriptions are meant to accrue incrementally over time rather than instantaneously. When asked to generate a formula for what made a subscription valuable, most participants mentioned frequency and recency of use.

For Archana (21, US), Spotify was her most useful subscription because she was using it every day.

> I've had my Spotify subscription for years without cancelling. And it might be one of my longest running subscriptions. That is something that I don't re-evaluate because that's how I go to sleep every night as I listen to my Spotify.

Some subscriptions provided better personalisation over time, or better deals for loyal customers in order to provide continuous value to customers. Participants also found greater value in some long-term subscriptions. For example, Archana appreciated the fact that over 10 years that she'd been using Spotify, it reminded her of various different periods in her life. The prospect of future personalised

experiences was an important benefit to her and played an important role in her original intention to form a long-term attachment to Spotify:

> I think their algorithm does a very good job for me personally. That's what persuaded me to buy it in the first place.

The right opportunity was another important temporal factor. A lot of participants understood that their engagement with subscriptions was dependent on their life circumstances. Jane (26, UK) understood that her use of entertainment subscriptions fluctuated:

> I think it fluctuates how much I watch Disney. Like, I'll go through phases. I've maybe not been watching it as much recently, so I'd maybe think, oh, maybe I should just cancel it for a few months and see how it goes.

Certain life stages were particularly significant triggers for adopting new subscriptions, especially moving house, becoming financially independent, starting work, looking after others, especially children, or retiring. Similarly, specific events or occasions inspired participants to get a subscription. For example, when Nana (26, US) and his girlfriend got a puppy, and because they lived on the 4th floor of an apartment block and didn't want to have to go outside, they bought a subscription to a patch of grass called Doggy Lawn.

Many participants considered subscriptions to be a long-term commitment. Ideally, they would prove to be rewarding over time, however, they could also produce diminishing returns and become forgotten. This was a risk many recognized, and some were so worried about being stuck with a subscription they couldn't get out of, forgot to use, or manage, that they'd intentionally try to pause them or avoid them altogether. The fact that the value of subscriptions was relative to time also meant that participants struggled to establish the value of a particular subscription. Thus, the value of subscriptions was never stable and could not be easily predicted.

3. Affordability

Being able to afford subscriptions was an important factor that impacted people's willingness to adopt one. However, while price was an important point of reference, participants were also taking other factors into consideration in order to determine whether they could afford a subscription or not. Participants needed to consider whether they'd be able to stay on top of their bill along with other bills, for

how long and whether it might be difficult to pause or cancel their subscription, and they also considered any alternatives.

For example, during the time of our research, many participants were aware that Netflix had changed the rules around sharing subscriptions.[7] This triggered many different considerations. Archana (21, US) was sharing Netflix with a friend who was also sharing it with someone who wasn't living with them. Archana thought that if she could no longer use this account, she wouldn't get her own Netflix subscription. Instead, she'd continue using her Dropout subscription over Netflix. Dropout offered enough comedy to provide the important escape from social media, which was the main reason Archana paid for a streaming service.

> Dropout is so cheap comparatively. Like, it's a dollar a month when it works out between all my friends. It feels like bad value to pay $6.99 a month for Paramount, for example, and then pay $1 for Dropout and watch Dropout more.

Importantly, she thought Dropout was paying their creators fairly and she wanted to support their business model:

> It's not necessarily just supporting a creator, any creator. It's supporting specific creators that reflect my values. … It baffles me that Netflix and Disney and Paramount and Hulu cannot afford to pay their actors and writers livable wages. And that's why there's a strike going on.

For young students, such as Archana and her friends, the change of Netflix's conditions felt destabilising. Part of the value of Netflix was the ability to share shows that others were watching too.

> The benefit it brings to my life is also that now that I share a subscription with all of my friends, that means we have something in common. Automatically, we all watch the same stuff.

Sharing his Netflix account was important to Frank (67, US) as well. He was sharing it with his brother in Japan. For Frank, paying more was not a problem but the change in conditions was disappointing and made Frank critical of Netflix's business model:

> If you're [Netflix] so hard up for money, maybe consider not making so many television shows. Part of the reason why you're the number one in this regard is that you make it available for so many other people. And if you want to, you know, rather than trying to crack down, if you want to

> charge a couple of dollars more a month, then okay, do so. But allow for us to be able to share it with whoever we want, and it doesn't matter where.

For many participants, there was eventually a ceiling where budgets constrained the number of subscriptions they could have. This was more important to help them determine what they'd be willing to pay than the price of an individual subscription. Nana (26, US) told us how he wanted to learn how to DJ and make music. He had already been paying for Apple Music, Spotify Premium and YouTube Premium but the DJ tool he wanted to use required that he used Tidal so he got Tidal as well. However, he then cancelled both Tidal and the DJ tool, when he realised that, while his expenses were increasing, his income was remaining the same.

> I have gotten to a place where I just cancelled the ones that don't necessarily work for me in this moment, just because I feel like I don't have that luxury anymore of just, like, having this go on for however long.

The objective value of a subscription is impossible to determine. While price was an important piece of information, it wasn't enough for participants to determine value. Callon reminds us that price is not about the value of a good, it is a quality of the good (Callon 2021, 288). Price, says Callon, "is neither the consequence nor the measure of the value of a good. It participates in the formulation of this value…. Price is a cause of value, but only one of its numerous causes. The consumer buys the price…" (Callon 2021, 300.) Price helped participants compare similar or different products and experiences. Iman (37, UK) felt in general she wouldn't pay over £30 for a subscription but wouldn't mind paying £7.99 or £5.99 because that's "like a cup of coffee".

Price could also be negotiated or lowered. Many participants actively looked for ways to lower the costs of their subscriptions—for example by sharing it with others, by registering for free trials with new email addresses or by paying with gift cards to avoid unwanted renewals. Importantly, the affordability of a subscription was also relative to how they engaged with it. For example, Archana (21, US) paid for Jack's Flight Club because she knew she was going to book flight tickets. She then intentionally maximised its benefits so that it would have turned out to be good value for money.

4. Engagement

As mentioned above, participants wanted to use subscriptions to achieve various desired effects. But they understood that in order for subscriptions to be good value for money, they'd have to use them frequently. Thus, getting a new subscription required forming new habits. Yolanda (48, US), who had recently started dating, was considering getting a subscription that would enable her to discover beauty products. She saw this as an investment into herself, as well as potentially doing something different. This could even impact how she'd re-evaluate the importance of entertainment.

> I think the whole makeup thing is fun for me, and it's doing self care. … I'm not sure I want to have more fun. And I guess it's something new and different for me. Netflix and Hulu, you kind of know what you're getting.

Subscriptions were important for providing access, convenience and flexibility. Being able to access content or service as one needed was valuable and participants registered these occasions as engagement. Often, they would also consider it valuable to have the possibility of this engagement, should their needs arise. Jane (26, UK) was paying for her Audible subscription even though she wasn't using it, in case she would like to use it in a car or at the gym.

Sometimes, just having a subscription, without using it, was important for their identity, values and relationships. Iman (37, UK) wasn't using many of her streaming subscriptions for herself but shared them with others. For example, she was paying for her Pret a Manger subscription even after she stopped using it because her brother and nieces could still use it. Often, even as participants' lives changed and they could disengage from their subscriptions, they were still sharing them with others, creating further dependencies and attachments.

Nana (26, US) also talked about keeping subscriptions despite not using them because they represented who he wanted to become. In his instance this could be becoming a DJ. Having a dedicated subscription represented the aspirational identity and the possibility of achieving his goal. He compared it to how others get a gym membership to motivate themselves to get fit:

> There's some things you just want to try out or, like, you aspire to try out.

Many participants recognised that their engagement with subscriptions was a result of the different mechanisms that were used to get them "hooked" so that they would become part of their routines. Archana (21, US) recognised that for

subscriptions to become more valuable, they were aiming to be embedded in people's lives.

> Subscriptions that are successful entice you to use it more often, so it becomes part of your routine. Like, for example, Netflix does this with their two-part releases. So, like, I was watching the Witcher, season 3, and they released the first four episodes one day, and then they released the second four episodes two weeks later, even though everything was already done. The reason for this is they make you use it more in terms of, like, you're using it today and you're using it two weeks from now. And that makes it part of your routine, that makes it something that you see as more valuable.

Participants often mentioned that the value they were getting from subscriptions was related to their ability to be aware of their behaviours, change their behaviours and navigate their activity and passivity. Iman (37, UK) had forgotten she was paying for her Elevate subscription. This realisation made her reactivate her interest and start using it more to achieve the desired effects that motivated her to get the subscription in the first place. For her, the cost was less important than the chance to re-engage with her aspiration:

> I forgot that I was paying for it. Honestly, this is how bad I am with money. So for me, I think I probably will start using it more often. Instead of forgetting that it's even there.

Most participants considered frequency and recency of use to be important indicators of how much value they were getting from a subscription. However, most also realised that their actual behaviours did not support this formula. Having the possibility to use a subscription was also valued because it was seen as motivational and aspirational.

On the other hand, pausing a subscription could be seen as a motivation to pursue other things in life while still valuing the subscription and intending to resume the attachment at a later time. Getting a paid subscription could be considered a way to "hack" one's behaviours and motivate oneself to create new habits.

How people engage with a subscription plays an important role in how they attribute value to it. However, engagement was just one of the indicators of how a customer might attach to their subscription and to what effect.

The 4 factors we identified surface how important it was for participants to navigate their activity and passivity to get value out of their subscriptions. We found

that participants were aware of their attachments, dependencies, and their day-to-day weakness of will and correlated passivity. In other words, they were aware of their lack of awareness around their habitual behaviours and around the broader consequences of forming attachments to subscriptions. This lack of awareness was exacerbated by the fact that participants were never quite sure what they were paying for, were entitled to and how they could manage the services.

This complexity and subjective understanding of subscriptions means that evaluating them is a relatively difficult task. Callon argues that goods lie on a continuum from those that are more easily defined by their qualities, to those that are more uncertain because their qualities are less clear, and customers can ascribe their own meanings to them (Callon 2021, 260). Subscriptions are very much at the uncertain end of this spectrum.

Evaluating Subscriptions Requires Calculative Agencies and Calculative Tools

Evaluating individual subscriptions is difficult, it's even more difficult to curate one's portfolio of subscriptions. While participants often had subscriptions that were overlapping in terms of the entitlements they offered, they also often felt like they were missing some subscriptions. For example, Nana (26, US) had several music subscriptions but wouldn't recognize that his Amazon Prime was bundled with the Amazon Music subscription. Jack (65, UK) had multiple subscriptions for watching sports but also had to go to the pub to watch his favourite team play because none of his subscriptions provided access to all the content he wanted to watch.

The valuation and evaluation of complex products, often bundled and highly personalised, such as subscriptions, is an understudied subject. In our analysis, we go back to Callon whose model stresses the importance of calculative devices as well as calculative agencies.[8]

Often, the way subscription options are presented to people is through different tiers—this helps to create choice for people, but it also provides references that people can use to compare options. This information provides cognitive tools for people to make calculations. Price, as discussed above, is just another quality of a subscription. For example, the price of a Netflix subscription places it within the market of other entertainment subscriptions offered by competitors and the possibility to choose from 3 different tiers of Netflix subscriptions also creates an internal hierarchy within which each tier is positioned against the other.

The way various subscriptions and bundles are presented to potential customers help them compare them with other options. For example, the resolution of 1080p means very little, without being positioned next to the resolution of "4K+HDR" and a higher monthly price. The price and other specs, thus, play the role of information devices and serve as cognitive tools that can help customers make decisions. Without these tools, customers are not able to determine the value of a subscription, not to mention the price they're willing to pay.

Callon (1998) doesn't buy into the idea of people having an innate ability to make rational decisions. In line with the symmetrical approach of Actor Network Theory, he argues that it is only thanks to various calculative devices that people enter into market transitions, and that any markets can exist, for that matter. Calculative devices include all sorts of measuring tools, currencies (especially money), accounting tools and management processes that enable pricing, comparisons and calculations. "Calculativeness," he says, "couldn't exist without calculative tools" (Callon 1998, 23). The dynamics of markets, as well as asymmetries between agents, can be explained by agents performing various types of calculations and using various types of calculative devices.[9]

According to Callon and Muniesa, calculation

> …starts by establishing distinctions between things or states of the world, and by imagining and estimating courses of action associated with things or with those states as well as their consequences (2005, 1231).

Calculation is neither purely quantitative, nor qualitative, and therefore cannot be reduced to a mathematical calculation or an intuitive judgement. In the market of subscriptions, customers are also taking into consideration their potential long-term engagement with a subscription. While they cannot know exactly how successful this will be, they can interpret various signals and predict how they might be activated or passivated by a particular subscription and form an attachment.

Looking at different Netflix subscription tiers, Archana, (21, US) reflected on how she would interpret them:

> I'm not looking at the video quality or the resolution when I'm deciding whether or not to buy it. I need to see monthly price. These bottom two check marks are telling me about how much I'm going to get it used. Downloads are another way for me to use it. Watching on my TV or my phone are different ways for me to use it. And then what I would need to

> see is, what kind of stuff do they have on there? And that's about, does it spark joy or not?

Archana was able to combine the information presented with her predicted behaviours. She could then compare this to other options, especially the Dropout subscription. The fact that the Dropout content was something she and her friends talked about, helped her prioritise the Dropout subscription.

The importance of information devices and calculative devices that help people frame an offer is important in helping them determine their willingness to pay for a subscription. However, Archana's response also reflects the existing information asymmetry, which left her unsure whether a subscription would spark joy, and how affected she might be by the subscription and to what effects and outcomes.

Pick Your Poison

According to Callon, markets are animated by what he calls *affectio mercatus*, passions, tastes, habits and energies that enable attachments and dependencies. Callon's agencement model stresses the importance of the material component of market activities. To experience a desire for a product requires that one can be affected by it, that they are moved or changed by it (c.f. Latour 2004). Subscriptions too, have a material component. Their effects are felt as they impact people's daily habits, rhythms, moods, energy levels, attention spans, tastes, as they generate new desires and new actions. The effects of ongoing encounters facilitated by subscriptions or experiences can come in virtually infinite forms. For example, one subscription can be used to change one's mood, be inspired or to fall asleep. However, for an attachment to be ongoing, it needs to also create a form of dependency and leverage people's dispositions to form dependencies.

Markets leverage these dispositions to captivate customers by appealing to people's learning systems and the underlying two different decision-making circuits, one responsible for impulse decisions and another corresponding to longer deliberations (Callon 2021, 254). The attachments that are formed between agents and goods can be of various intensity and length, with some utilising human hedonistic drives to create addictive attachments.

Archana (21, US) felt like games were most powerful at leveraging human dispositions to form dependencies but thought that subscriptions too, were using rewards to become addictive:

> Games get you to spend money by, like, you know, they're hitting that dopamine reward center, too. … For example, Clash of Clans has this formula where every day you log in, whether you're not, you have a subscription, you get some bonus thing. So they make it part of your routine by rewarding you for making it part of your routine. Subscriptions that are successful, I feel like, do much of the same where they entice you to use it more often, so it becomes part of your routine.

Callon predicts that "[e]ngineering addiction and addictive practices makes explicit and weaves together on a whole new scale the stuff that contemporary man is made from" (Callon 2021, 283). The insights from our research show how agents respond to these new attachments, how they are shaped by them and how they adapt to the market that is dependent on utilising them.

Getting "hooked" or becoming addicted to hits of dopamine go contrary to our understanding of the will. As our research shows, participants were aware of their lack of awareness, of their passivity, dependencies and weaknesses of the will or weaknesses of their memory. Many participants felt that this was part of being an agent in the markets of subscriptions. In this context, it was impossible to determine their willingness to pay. Nana (26, US), who felt dependent on Netflix expressed this uncertainty through a pause when he said:

> If Netflix decided, oh, yeah, well, we're so incorporated into your life, we're charging $50 now, I'm not paying $50 for Netflix [pause] probably.

How do businesses respond to the market that is shaping and being shaped by agents who feel like their best choice is to "pick their poison"?

Adopt and Adapt: Reframing Willingness to Pay

> I'm open to new subscriptions if it makes life easier. Cost doesn't really come into it, up until a point. Then you're paying a lot just to be lazy. (Nana, 26, US)

Subscriptions are agencements, they should be thought of as dynamic processes or systems that produce outcomes by incentivising attachments. To understand how value is attributed to subscriptions, multiple factors need to be considered. Our research has identified factors relating to desired effects, temporality, affordability and engagement. Importantly, we also explained how, as a result of these

attachments, a new type of customer agency emerges: one that navigates their own passivity and activity and recognizes addiction as part of its condition.

For this agency, price and other information devices that are used to present subscriptions—such as streaming quality, number of downloads, etc.—serve the purpose of helping them understand how a subscription might leverage their conscious and unconscious behaviours, their weaknesses or cognitive biases. In the context of information asymmetry, exacerbated by the use of data and algorithmic technologies, customers experience the need to, above all, assess their current and anticipated behaviours and circumstances.

In essence, when determining the value of subscriptions, customers need to establish what might be the outcomes of the adoption of a subscription and how they might adapt to it. Attachment, Callon argues, is a structured process that can be organised. This process involves experimentation, investigations and trials through which "affects are produced, identities transformed and connected, durable connections made, and solid associations between goods and agents established" (Callon 2021, 265).

Without attachments, markets cannot exist. Equally, markets cannot exist without detachments. "Subscriber fatigue," referred to by Tzuo, needs to be understood not in the context of singular subscriptions, but in the context of the broader subscription economy, in which the sum of various subscriptions people are paying for is affecting how they are evaluating, using, pausing or cancelling various individual subscriptions.

Last but not least, the role of personal data needs to be considered when thinking of the value of subscriptions. Consumption practices are key to how modern consumers craft their lifestyles and their selves. In 1999, Nikolas Rose wrote about the condition of the consumer: "we are obliged to make our lives meaningful by selecting our personal lifestyle from those offered to us in advertising, soap operas, and films, to make sense of our existence by exercising our freedom to choose in a market in which one simultaneously purchases products and service, and assembles, manages, and markets oneself" (103). Digital technologies that leverage personal data have enabled a more detailed understanding of one's decisions and habits. Perhaps we are less impacted by what we see in soap operas, and more impacted by the reflection of our behaviours as they appear in our screen times, step count or number of accomplished tasks. The data that is available to us impacts how we organise our lives and further feeds into the pressure to be responsible for and

educated about ourselves (Rose 2004, 88). Coupled with predictive technologies, subscriptions raise new questions about how modern agents are governed through their consumer practices.

The modern subscriber understands their dependency as well as their need to be responsible for their lives and their decisions. Where businesses are profiting from predictable cash flows and recurring revenues, customers experience challenges and powerlessness associated with their inability to manage the subscriptions they are using or not using. We would like to invite the community of ethnographers to provide more understanding and solutions that would support customers while also helping grow sustainable businesses.

About the Author

Iveta Hajdakova, PhD, is an Associate Director at Stripe Partners. She holds a PhD in Anthropology from Charles University in Prague and was a 2012–13 Fulbright-Masaryk Visiting Scholar at Columbia University in New York. Iveta specializes in economic anthropology and the anthropology of experience. Her approach draws inspiration from feminist thought, philosophy of technology, and Science and Technology Studies.

Morgan Williams is a Senior Consultant at Stripe Partners. He holds an MA in Sound for the Moving Image from the Glasgow School of Art as well as an MA in Philosophy from the University of Glasgow. Morgan specializes in AI and AI ethics and has a deep understanding of the media industry, having previously worked in sound for film and music.

Notes

Acknowledgements: We would like to thank our curator Claudia Ruge and our colleague Cath Richardson for their feedback and suggestions.

1. The translator of Callon's book, Olivia Custer, explains that agencement might be translated as "arrangement" or "layout". "It can refer to the way elements are laid out in space or the way they combine to create an effective whole. In order to preserve the insistence on the agency arrangements, the choice was made to import it into English" (Translator's note, Callon 2021, 428). Callon explains that the term is compatible with the definition of an assemblage. However, he doesn't want to use the word assemblage because it creates the sense that the various entities associated with an assemblage are just put together. Callon wants to stress the collective action that structures these assemblages. At the heart of the market agencement model are the various interdependent framings that, together, make it possible for a good to be attached to an agency (Callon 2021, 361-362).

2. For more on the various tools and processes used in evaluation, see Lamont (2012). According to Lamont (2012), it is a cultural process. Lamont proposes that to understand valuation and evaluation, we need to understand various other practices and factors: 1. How does categorisation and legitimation happen? 2. What are the different criteria and "grammars of worth"? 3. How is knowledge produced and what non-human agents and devices are used to produce that knowledge? Karpik (2010) has also provided a useful perspective on valuation, which has inspired a case study by Hajdakova, Karol and McDonald presented at EPIC 2020.

3. Framings and the tension between framing and overflowing is important to Callon's explanation of market dynamics, pricing, marketing and market competition.

4. For more on nudging, see Callon (2021, 155) and Fourcade and Healy (2024, 145).

5. We noticed that participants felt angry at themselves, guilty or ashamed if they've failed to stay on top of their subscriptions. These feelings point towards a moral economy associated with subscriptions. Fourcade and Healy also point out the connection between ordinalization (the use of data, metrics, rankings, etc.) and morality (2024, 245-247). For the purposes of this paper, we did not focus on the moral aspect of the subscription economy albeit we think that morality plays a crucial role in shaping people's attitudes towards subscriptions and informs how they navigate their activity and passivity, and should be closely examined.

6. Building on Pierre Bourdie's concept of the habitus, Callon argues that how people calculate and evaluate goods is part of their habitus (Callon 2021, 195).

7. Netflix. 2023. 'Update on Sharing'. May 23, 2023. https://about.netflix.com/en/news/update-on-sharing-may-us

8. In his earlier work (Callon 1998), Callon talks about "calculative agencies" whereas in his later work (Callon 2021), he talks about "qualculations" to highlight the qualitative aspect of calculations. For the purposes of this paper, we found the concept of "calculative agencies" clearer and more sufficient.

9. There are plenty of asymmetries between agents involved in a transaction and calculative agencies, do not use the same calculative tools. The more complicated and complex the network of relations and entities is, the larger the number of potential asymmetries (Callon 1998, 45).

References Cited

Appadurai, Arjun. 1986. "Introduction: Commodities and the Politics of Value." In The Social Life of Things: Commodities in Cultural Perspective. Edited by Arjun Appadurai. Cambridge University Press. 2–57.

Adelman, George. 2022. “Foreword by FT Strategies.” In Subscription Economy: A Transformed World. FT Strategies. https://www.ftstrategies.com/en-gb/insights/the-latest-report-on-the-subscription-economy-a-transformed-world

Callon, Michel. 1998. “Introduction: The Embeddedness of Economic Markets in Economics.” In The Laws of the Markets. Edited by Michel Callon. Oxford and Malden: Blackwell Publishers.

Callon, Michel and Fabian Muniesa. 2005. “Peripheral Vision: Economic Markets as Calculative Collective Devices.” Organization Studies 26(8): 1229–1250.

Callon, Michel. 2021. Markets in the Making: Rethinking Competition, Goods, and Innovation. New York: Zone Books.

Fourcade Marion and Kieran Healy. 2013. “Classification Situations: Life-Chances in the Neoliberal Era.” Accounting, Organizations and Society 38(8): 559–572.

Fourcade, Marion and Kieran Healy. 2024. The Ordinal Society. Cambridge: Harvard University Press.

FT Strategies. 2022. Subscription Economy: A Transformed World. Accessed August 5, 2024. https://www.ftstrategies.com/_/view_pdf/file/906/b26b64cb001/

Goldfield, Hannah. 2024. ‘Why New York Restaurants Are Going Members-Only’. The New Yorker, March 18, 2024. https://www.newyorker.com/magazine/2024/03/25/why-new-york-restaurants-are-going-members-only

Hajdakova, Iveta, Sohit Karol and Deb McDonald. 2020. “Beyond User Needs: A Meaning-Oriented Approach to Recommender Systems.” EPIC Proceedings pp 191–202.

Ivanov, Asen., Catherine Anne Johnson, and Samuel Cassady. 2020. “Unbundling Practice: The Unbundling of Big Deal Journal Packages as an Information Practice”, Journal of Documentation. https://atrium.lib.uoguelph.ca/items/703c9e04-0bf9-4d68-868c-5113bbd892a7

Karpik, Lucien. 2010. Valuing the Unique: The Economics of Singularities. Princeton University Press.

Lamont, Michèle. 2012. Toward a Comparative Sociology of Valuation and Evaluation. Annual Review of Sociology 38(1): 201-221.

Lauer, Josh. 2020. “Plastic Surveillance: Payment Cards and the History of Transactional Data 1888 to Present.” Big Data & Society 7(1). https://doi.org/10.1177/2053951720907632

Latour, Bruno. 2004. “How to Talk About the Body? The Normative Dimension of Science Studies.” Body & Society 10(2-3), 205-229.

Maurer, Bill. 2014. "Postscript: Is There Money in Credit?" Consumption Markets & Culture 17(5): 512–518.

Netflix. 2023. 'Update on Sharing'. May 23, 2023. https://about.netflix.com/en/news/update-on-sharing-may-us

O'Dwyer, Rachel. 2019. "Cache Society: Transactional Records, Electronic Money, and Cultural Resistance. Journal of Cultural Economy 12(1): 133–153. Accessed August 5, 2024, from https://www.academia.edu/30478963/Cache_Society_Electronic_Money_Transactional_Data_and_Cultural_Memory

Rose, Nikolas. 1999. Governing the Soul: The Shaping of the Private Self. London: Free Association Books.

Rose, Nikolas. 2004. Powers of Freedom: Reframing Political Thought. Cambridge: Cambridge University Press.

Sinders, Caroline. 2023. "How Companies Make It Difficult to Unsubscribe." The Pudding. Accessed August 4, 2024. https://pudding.cool/2023/05/dark-patterns/

Tzuo, Tien and Gabe Weisert. 2018. Subscribed: Why the Subscription Model Will Be Your Company's Future – and What to Do About It. New York: Portfolio/Penguin.

Tzuo, Tien. 2015. "The Subscription Economy: A Business Transformation." Medium, March 10, 2015. Accessed August 4, 2024. https://ientzuo.medium.com/the-subscription-economy-a-business-ransformation-83d6fb24a2f9

Tzuo, Tien. 2024. "Is the Subscription Economy Over?" LinkedIn, March 9, 2024. Accessed August 4, 2024. https://www.linkedin.com/pulse/subscription-economy-over-tien-tzuo-w6vwe?trk=news-guest_share-article

Ughade, Nupura. 2024. "All You Need to Know About the Subscription Business Model." Chargebee Blog, July 15. Accessed August 5, 2024.

Zuboff, Shoshana. 2019. The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. New York, NY: PublicAffairs.

Quit Playing Language Games with My Heart; or, Conversational AI and Knowledge Sharing

LARRY S. McGRATH

This paper shares guidance for designing conversational AI based on findings from linguistic and social analysis of core shifts introduced by large language models. When Duolingo and Babbel added AI features to their language learning apps, the games that users play on the apps was transformed, giving us a window into the structure of AI-driven user interactions in general. Dynamic, turn-taking dialogues resembling natural conversation now appear alongside multiple-choice and drag-and-drop vocabulary games. The transformation is not just technological: It also lays bare the shifting language games that we have come to play more broadly in our interactions with LLMs such as Chat-GPT, Claude, and Gemini. People engage with them in back-and-forth conversations in which different kinds of speech acts are at play. By contrast, traditional conversational agents based on decision-tree learning rely on discrete, finite, and close-ended commands. The paper argues that the future of AI conversation is one in which rule-based interactions co-exist with intention-based interactions powered by generative AI. Design elements on Duolingo, in particular, draw on diverse styles of communication and offer guidance to ensure that conversational AI platforms in general remain open to new forms of dialogue that enrich both artificial and natural conversation alike.

> "Here the term 'language-game' is meant to bring into prominence the fact that the speaking of language is part of an activity, or of a form of life."
>
> – Ludwig Wittgenstein, Philosophical Investigations §23

1. Introduction

Speaking a language is like playing a game. For Wittgenstein, learning a language – that is, grasping how to describe, question, command, and explain – involves a process of mastering the rules of the games that shape human interactions. In this paper, I take Wittgenstein's words literally and examine language games in the learning apps, Duolingo and Babbel, which transformed in 2023-2024 from rule-based to AI-powered conversation platforms (Freeman 2023). The transformation is not just technological; it also lays bare shifting modes of interaction with chatbots. Unlike decision-tree chatbots, intention-based AI interfaces support communication through back-and forth dialogues in a dynamic manner not wholly unlike natural conversation. I suggest that studying language games in everyday conversation can enrich the design of language learning apps and ensure, more broadly, that

McGrath, Larry. "Quit Playing Language Games with My Heart; or, Conversational AI and Knowledge Sharing." *EPIC Proceedings* (2024): 254–275. https://www.epicpeople.org/conversational-ai-and-knowledge-sharing. ISSN 1559-8918.

conversational AI platforms remain open to the diverse forms of interaction that already animate dialogue among humans.

Language games serve to demonstrate that words' meaning is better understood as part of a conversation among humans rather than as a representation of the world. For Wittgenstein, speakers of a language can understand each other because they partake in a shared activity. What speakers achieve in mutual understanding, however, comes at the risk of incomprehension, failure, and rejection. Much can go wrong in conversation. No facts – whether in the world or in our minds – guarantee the successful transmission of meaning. Such is also the case, as I show in this paper, when playing games with language learning apps. As Large Language Models (LLMs) grow in size and complexity, expanding conversational AI platforms' emulation of natural dialogue, so too do they replicate the fragility of their fallible creators. An important characteristic of language learning apps' sense of realism, I'm suggesting, is to be found in their potential for confusion. Duolingo and Babbel can bewilder. AI integrations that make the apps approximate human-like conversation also make them succumb to what Stanley Cavell describes as "the deceptions and temptations and dissatisfactions of the ordinary human effort to make oneself intelligible (to others and to oneself)" (Cavell 2004). As thinkers in Wittgentsein's wake have shown, language games can always run aground and land in misunderstanding. Our techniques for navigating confusion in quotidian conversation offer a guide to steering artificial conversation back on course.

Duolingo and Babbel exemplify the transformative impact of AI-powered conversation platforms. The apps use Natural Language Processing (NLP) to create personalized and adaptive learning experiences. Duolingo uses AI to tailor lessons to individual users' performance, providing each learner with appropriate levels of challenge and support (Duolingo 2023). Similarly, Babbel employs AI to recognize users' speech and provide real-time feedback on pronunciation and grammar (Babbel 2023). AI's integration not only enhances user engagement through realistic dialogue and instant feedback; a window also opens onto broader interaction patterns between humans and conversational user interfaces. Indeed, conversational AI is a partner in dialogue.

Conversational AI platforms are chatbots that people interact with by typing natural language, often over the course of back-and-forth interactions. There are two kinds of platforms. First, universal AI platforms such as Chat-GPT, Claude, and Gemini have a general scope of knowledge; their intended use is as standalone products. Second, product-specific AI chatbots (including Duolingo and Babbel) have a scope of knowledge limited to an organization with a specific service; their

intended use cases are as embedded features designed to address the organization's needs. I'm suggesting in this paper that language learning apps' AI features open a window onto the deeper structure of conversational interaction with the former platforms.

My own conversation is with scholars in Science and Technology Studies as well as Human-Computer Interaction, especially those who have moved beyond debates about AI platforms' epistemic capacities. An outdated current of thought remains captivated by AI's ability to mimic intelligence, reasoning, creativity, or consciousness akin to that of humans. On the one hand, technology enthusiasts contend that conversational AI realizes a theory of mind (Kosinski 2023; Summers-Stay 2023). Infamously, Google fired Blake Lemoine, an engineer, after he claimed that the Language Model for Dialogue Applications (LaMDA) had reached a level of human-like consciousness (Wakabayashi 2022). On the other hand, Noam Chomsky argues that nothing of the sort is possible. "The human mind is not, like Chat-GPT and its ilk, a lumbering statistical engine for pattern matching, gorging on hundreds of terabytes of data and extrapolating the most likely conversational response or most probable answer to a scientific question" (Chomsky 2023). Both sides of the debate, however, frame epistemic capacities as *possessions* of humans or AI systems. I want to suggest the ethnographers are well positioned, instead, to trace the behaviors and patterns that shape interactions *between* humans and AI platforms. Autonomy is bound up with automaticity. Neither are located on one side of the human-machine divide. A fresh scholarly focus on relations (rather than identities) guides recent inquiry into the natural conversation styles that precede and exceed conversational agents (Li 2023; Packer (2023). My article builds on that literature. As conversations with AI platforms approximate conversations among humans, I show how we can improve the former by drawing lessons from the latter.

In the pages that follow, I offer an armchair anthropology primarily of Duolingo (and secondarily of Babbel) guided by the philosophy of language. Wittengstein's insight that meaning is a form of action in a community orients my study of the forms of interaction between human and computer made possible by AI-powered language games. In the second section, I present AI features in Duolingo and Babbel and explain their leap beyond rule-based language games. In the third section, I show how two kinds of language games – traditional rule-based games and dynamic conversation games – engender distinct varieties of learning. Learners commit errors of knowledge when playing traditional games. With dynamic conversations, however, errors of acknowledgement take place when conversations follow unintended directions. In the fourth section, I elaborate on the theme of acknowledgment and

argue that it's pivotal not just in language learning apps but also in emergent interaction patterns with universal AI platforms. As technological progress brings those platforms' models closer to natural conversation, they also bear the risks of uncertainty and skepticism that haunt our all-too-human efforts to interpret others' intentions in everyday language games. In the final section, I show how language learning apps come to terms with these risks by combining rule-based interactions with open-ended conversations. With Duolingo in particular, the combination offers an elegant solution to the challenges that conversational AI platforms confront as they emulate natural conversation.

2. From Rigid Rules to Dynamic Conversations

Babbel was founded in 2007 with the mission of "creating mutual understanding through language" (Babbel 2024). Duolingo followed in 2011 the method as its centerpiece: "learn by doing" (Freeman 2023). Both mobile apps still feature language games in 2024 whose purpose is to learn languages and facilitate human conversation via speaking, writing, reading, and listening. Duolingo employs brief 2-5 minute exercises such as matching paired words, translating sentences, speaking sentences, and multiple-choice questions. Babbel employs longer scenarios in the form of dialogues, fill-in-the-blank exercises, and pronunciation practice. What makes the exercises *games* is both the progressive structure of the exercises, which advance by levels, as well as engaging features such as points, badges, rewards, and challenges (Saleem 2022). Although Duolingo explicitly gamifies exercises with playful designs, a cartoon bird mascot, and a streak system that rewards daily usage, games are central to both apps in the form of progress tracking and certificates for finishing courses. Like a board game, users can take out their mobile device, play briefly with Babbel or Duolingo, and return the apps to their pocket.

The apps initially relied on decision-tree algorithms to personalize language games for users' language skills. Decision trees are used in machine learning to make predictions based on a finite data set with predefined categories. The apps used Simple Adaptive Learning to match the difficulty of a game with the user's ability. Like a tree, the model begins with a root node (the user's initial language assessment after signing up and subsequent exercise levels while using the app); decision nodes split the dataset on the basis of defined conditions (such as the user's performance on an exercise and the time passed since the prior exercise); leaf nodes generate predictions (such as correcting the user's response and presenting feedback). Predetermined rules guide the decision nodes of the tree. They include, among

others, progression rules (if the user completes a lesson with at least 80% accuracy then the next level unlocks), exercise rules (if the user struggles with writing then more typing games appear), and error rules (if the user makes the same error three times then detailed explanations follow). Although both apps use supervised decision trees to personalize the games presented to users, Duolingo came to deploy more sophisticated machine learning models. By 2023, a recurrent neural network internally called Birdbrain supported game selection on the basis of a user's entire history on the app (Freeman 2023). A long short-term memory model translates users' performance into 40 vectors. Completing a game updates the vectors. Individual games are tailored to each user's learning level. Yet, the internal elements of any game remain fixed. A finite number of routes through the decision tree circumscribe the bounds of possible interactions with the traditional language games on Duolingo and Babbel.

Rule-based algorithms underlie traditional, rule-based language games on the apps. The user is introduced to the meaning of new words by using the words – that is, by playing games. The app takes turns by displaying a prompt; the user responds by answering. The user learns by doing in the sense that rules are embedded in the games; he or she does not grasp a rule apart from the words' use. "For a large class of cases," Wittgenstein wrote, "in which we employ the word 'meaning' it can be defined thus: the meaning of a word is its use in the language" (Wittgenstein 1953, §43). When a user performs well, and correctly applies a rule, the algorithm advances to more advanced topics. When the user fails to apply the rule, the algorithm revisits weaker areas. The rule-based games reflect the contractual terms of language. As John McDowell puts it, "to learn the meaning of a word is to acquire an understanding that obliges us subsequently – if we have occasion to deploy the concept in question – to judge and speak in certain determinate ways, on pain of failure to obey the dictates of the meaning we have grasped" (McDowell 1984, 325). The linguistic rules used to evaluate users' turns depend on specific vocabularies and grammatical structures, the categorical values that function as decision nodes in the apps' decision trees. The limited range of input prevents overfitting (when a machine learning model predicts accurately from training data but not from new data). To-and-fro, the app and user take turns playing language games.

Turns are finite, discrete, and close-ended. In the apps' traditional language games, a finite set of clickable individual words appear. For pronunciation and free-form response games, the apps recognize only predefined answers. Each of the user's turns constitutes an individual unit in the game. The units are close-ended in the sense that they follow a binary logic of validation. Answers are correct or

incorrect. Each individual turn in the game either applies or violates grammatical rules and, as a result, triggers the algorithmic rules in the decision tree to issue the next turn.

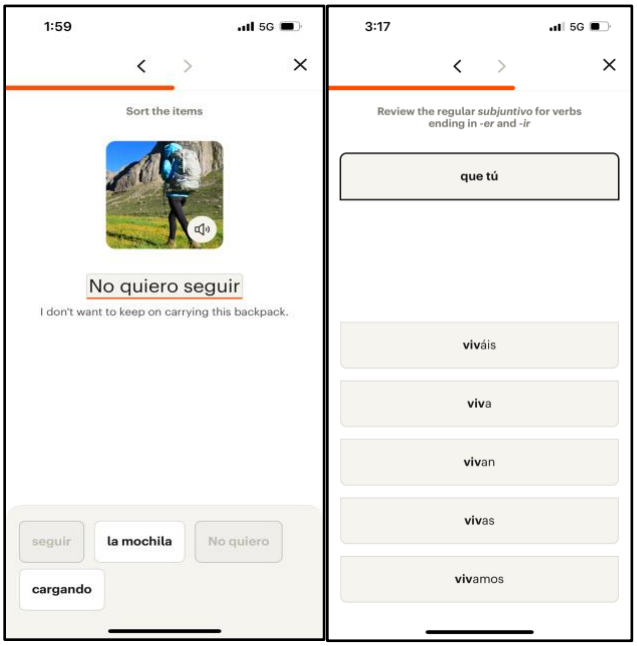

Figure 1a–b. Babbel – Rule-based language games

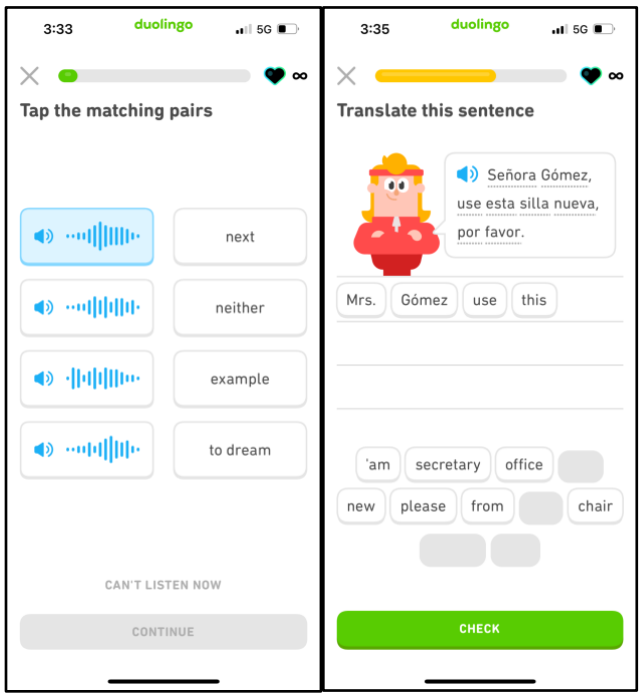

Figure 2 a–b. Duolingo – Rule-based language games

Yet, the traditional language games in Duolingo and Babbel lack the continuous character of natural conversation that unfolds between speakers of a shared language. That is why the apps have demonstrated mixed results in studies. When it comes to effective rule learning, Duolingo and Babbel help reinforce knowledge of vocabulary and grammar. In one study, 54 college students improved their Spanish scores by .7 of 5 levels on the standardized exam of the American Council on the Teaching of Foreign Languages after using Duolingo for 15 minutes per day over 12 hours (Jiang 2021). A review found that app-based vocabulary learning can be as effective as traditional instruction methods for vocabulary. However, students have had to go elsewhere to learn conversation skills (Tommerdahl 2022). Fluid turn-based conversation, as students would enjoy in classrooms, has been a shortcoming of traditional app-based language games.

Babbel and Duolingo added AI-powered features in 2023 to make their language games feel more like natural interactions with humans. Babbel added AI-enhanced speech recognition to "Everyday Conversations" for French, German, Italian, and Spanish. Duolingo's "Roleplay" surpasses Babbel's AI integration in its sense of realism and fluidity (currently available only for French and Spanish). Whereas Babbel's users read aloud predetermined lines, Duolingo's users can express anything that they'd like in open-ended interactions. "Roleplay" comes far closer to the back-and-forth dynamics of natural conversation.

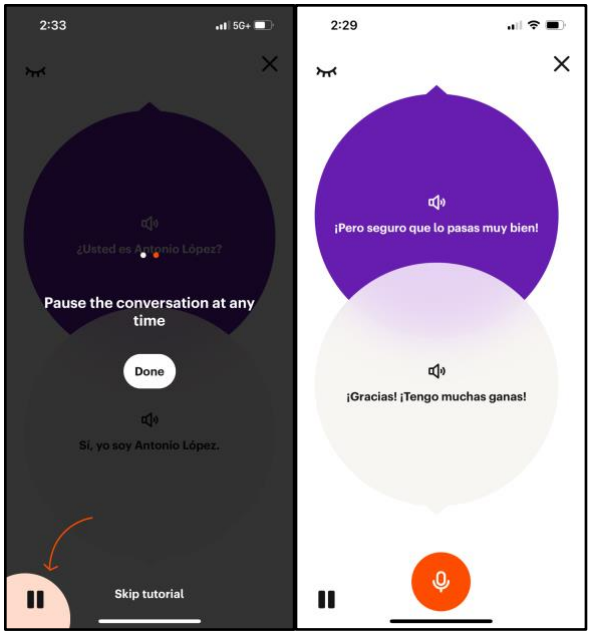

Figure 3a–b. Duolingo – "Roleplay" AI feature

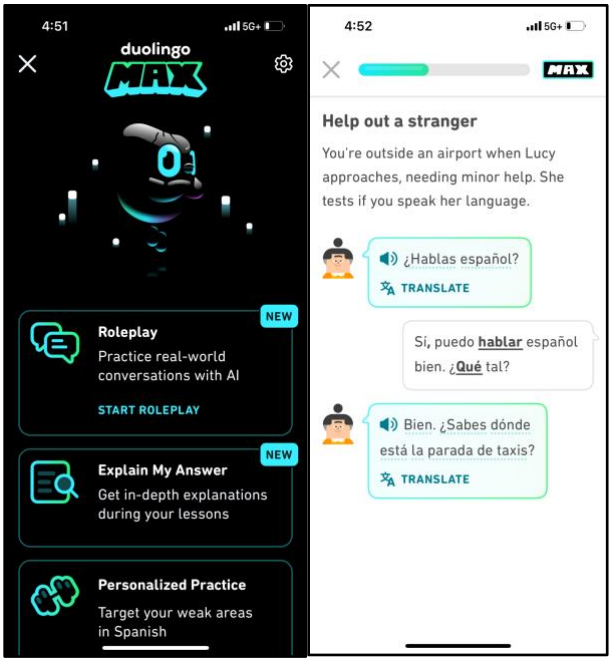

Figure 4 a–b. Babbel – "Everyday Conversations" AI feature

Continuous conversation comes to the fore in Duolingo's "Roleplay." Users do not carry out finite, discrete, and close-ended exercises. Rather, dialogue is open-ended. The uncertainty of the chatbot's response hangs on every utterance. Users can engage literally or add commentary; they can elaborate by injecting color, making revisions, or taking exceptions. Points are rewarded for turns with more words. Akin to a fluid conversation between humans, words are not set in stone but confront the pushback of a fellow conversant. Although Duolingo is unmistakably non-human, "Roleplay" depends on an interaction pattern that stands out from those of the app's traditional language games. Back-and-forth dialogue endures through time as each turn between chatbot and user spills into the unforeseen utterances of the next turn.

Neural networks facilitate the continuous character of dialogue in "Roleplay" by registering the various ways that users articulate the same utterance. Instead of parsing individual entries word by word, the networks deploy hidden layers between a user's input and Duolingo's output. Each node in the layers applies a set of attention mechanisms to the input and passes the result through an activation function before sending the output to the next layer. The attention mechanisms function as weights and biases; they assess all prior states of text according to learned measures of context and relevance. "Roleplay" integrates a Generative Pre-Trained Transformer (GPT), which applies weighted mean reductions to interpret input in the light of the collective input of users' interactions with the app's language games

(Duolingo 2023b). (Babbel has not disclosed which model facilitates "Everyday Conversations.") The result in Duolingo is a conversation whose continuity extends across alternating turns between chatbot and user.

Users of "Roleplay" play a turn-taking language game. They learn not just the grammatical rules of meaning but also partake in the social rules that structure a form of life shared with others. Conversation is the conduit of community. The conclusion of the chatbot's turn solicits the user's to begin. It's up to him or her to form a complete thought and bring it to a close in a manner that completes the turn and invites the chatbot's response. Conversations with "Roleplay" last seven to 10 turns. In their study of turn-taking, Harvey Sacks, Emanuel Schegloff, and Gail Jefferson wrote, "For socially organized activities, the presence of 'turns' suggests an economy, with turns for something being valued and with means for allocating them, which affect their relative distribution" (Sacks 1974, 696). Each turn is not a wholly isolated unit (although Duolingo does deploy feedback in response to incorrect utterances). The chatbot and user articulate turns in an organized sequence, what the pioneers of conversation analysis call an "economy" whose distribution of social roles has at once "general abstractness and local particularization potential" (Sacks 1974, 700). Users' words serve both to respond to the chatbot's individual utterances and to advance the next turn in a structured conversation.

The transition from language learning apps' rule-based decision trees to dynamic conversation supported by neural networks brought about a shift in interaction patterns from close-ended and finite actions to open-ended and continuous dialogue. Much like a natural conversation, the prior turn in "Roleplay" weaves with the subsequent turn. The chatbot's warp crosses the user's weft such that neither are entirely separate from their dialogical tapestry.

3. Learning a Language Via Two Varieties of Knowledge

Learning a language depends on making mistakes. Indeed, opening one's mouth, putting pen to paper, typing on a keyboard – all entail risks when sending new words into the world. The mistakes of schoolchildren who baffle their teacher as they stumble about trying to apply seemingly obvious rules appear again and again in Wittgenstein's writings. "Let us now examine the following kind of language-game," he writes, "when A gives an order B has to write down a series of signs according to a certain formulation rule… At first perhaps we guide his hand in writing out the series 0 to 9; but then the *possibility of getting him to understand* will depend on his going on to write down independently" (Wittgenstein, 1953, §143). On his own, B writes,

0, 1, 2, 3, 4… and suddenly, the series goes awry. B writes … 5, 7, 6, 9. The scenario illustrates how rules are communal phenomena that bond us to their future application. B's mistakes do not arise because he fails to share with A a mental picture – perhaps a formula – of ordinal numbers. The student's understanding is, ultimately, to be found in applying the rule over time in a shared community. Wittgenstein concludes, "And hence also 'obeying a rule' is a practice. And to think one is obeying a rule is not to obey a rule. Hence it is not possible to obey a rule 'privately': otherwise thinking one was obeying a rule would be the same thing as obeying it" (Wittgenstein, 1953, §202). To speak a language is a public act. And coming to understand the rules of a language involves a long and frustrating process of unexpected errors in the face of errors. That's why errors are key elements of Duolingo and Babbel.

Error messages appear in the apps when users make mistakes. For the traditional language games, decision trees register errors across various aspects of language – e.g., errors of grammar, vocabulary, pronunciation. Error messages take a few forms: incorrect answer messages display the correct answer and explain words' meaning; grammar messages display the correct tense, word order, or article; typos qualify as correct responses accompanied by clarifications of words' spelling; missing diacritics (in the case of Latin languages) also qualify as correct responses with messages that display correct accents.

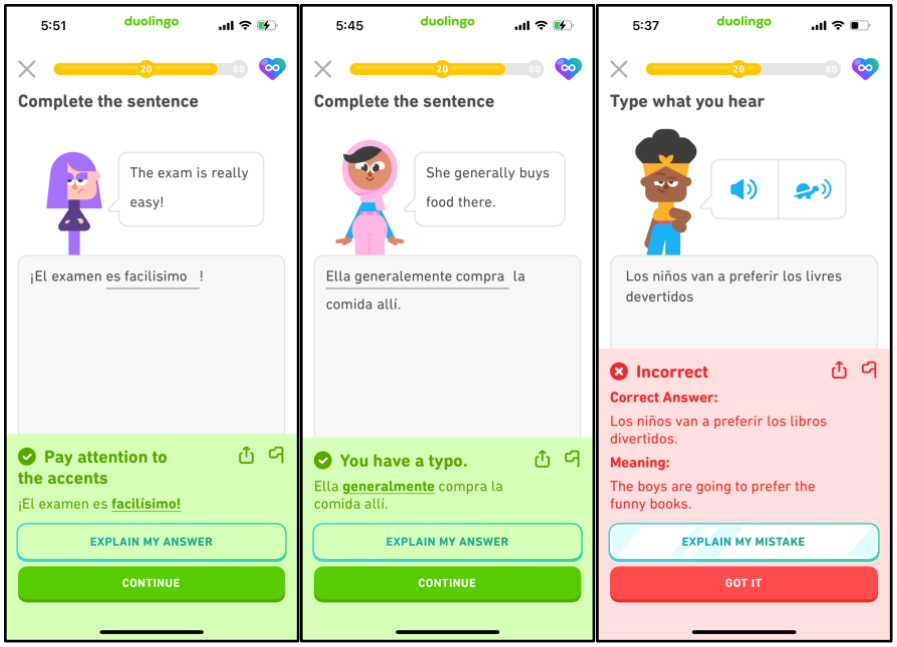

Figure 5 a–c. Error messages in Duolingo

The error messages respond to users' *errors of knowledge*. A user fails to correctly apply a rule of language. The messages are designed to guide users back to the rule's correct application. We're reminded of Wittgenstein's students, whose misapplications of a rule evince that they have yet to understand it. The error messages' goal is to restate the rule, encourage its correct applications in future instances, and thereby grow users' knowledge of a language.

Errors of knowledge also afflict LLMs. For its public debut, Google's Bard (now known as Gemini) incorrectly stated that the James Webb Space Telescope took the first image of a planet outside our solar system. In fact, the European Southern Observatory's Very Large Telescope already did so in 2004. Such "hallucinations" are false predictions that take the form of factually incorrect output (Ji 2023). Google's LaMDA model trained Bard on the basis of more than 1.56 trillion words and deployed 137 billion parameters to predict erroneously that the James Webb Space Telescope took the first image of an exoplanet.

Interactions with LLMs can also bring about varieties of error other than incorrect facts. *Errors of acknowledgment* occur when the course of conversation goes astray, takes unforeseeable turns, and results in misunderstandings. Words fail to act on each other with the appropriate force. And the resultant turns in a back-and-forth dialogue prove dissatisfying. Divergent interactions or shifting contexts in Duolingo's "Roleplay" are examples of acknowledgement errors that can arise artificial and natural conversation alike, even when conversants might utter factually correct sentences.

Take as an example when conversation runs aground. Duolingo's "Roleplay" organizes dialogues according to topics, such as "Help out a stranger" and "Mingle at a party." The topics serve as points of departure but – like a meandering dialogue – they need not be the final destination. Users can stray from the topic on account of confusion or an effort to steer the conversation toward their own ends. Because the AI–powered conversation unfolds continuously through turns, each turn has the potential for either conversant to lead the next down errant paths, even when chatbot and user contribute empirically correct content. Below, a user of French "Roleplay" responds to a conversation about parks in Paris by instead suggesting that they go to a park in Montréal. The conversation takes an errant turn to discuss a North American country – not a European country. Chatbot and user are not on the same wavelength. Although the sentences are meaningful, the conversation exemplifies an error of acknowledgment due to misaligned intentions.

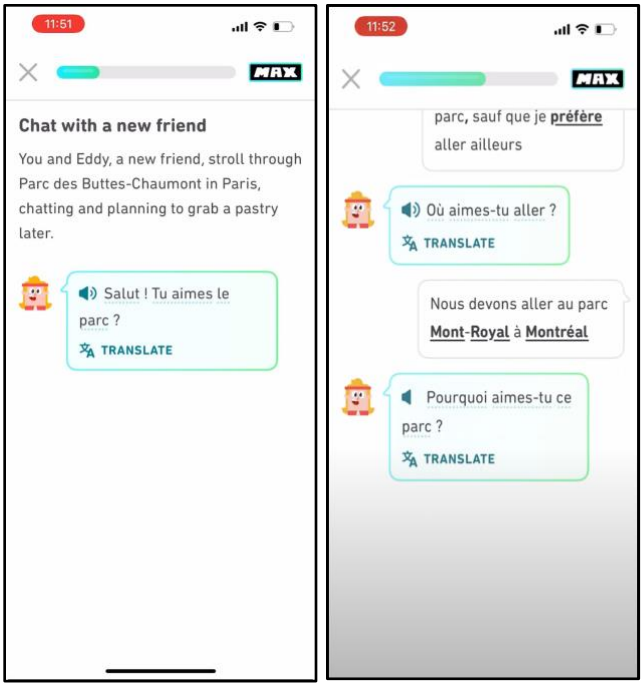

Figure 6 a–b. "Roleplay" in Duolingo

In the second example, new contexts lead to ambiguity and misunderstanding when words' meaning shifts. Learners of a language grasp new meanings of words by using them in different situations. Creativity and flexibility are key. Consider a user of Spanish "Roleplay" who meets a friend in Barcelona's Parque Güell. The user mentions that he gets lost in the designs ("Siempre me pierdo en sus diseños"). "Pierdo" can mean getting lost physically or metaphorically, as if to be captivated by the designs. For a Spanish language learner familiar with just the word's literal meaning, the idea that his friend might be lost could feel unsettling; it would certainly fail to convey his sense of amazement stirred by Gaudi's architecture. The error of acknowledgment reflects the same words being used differently by chatbot and user.

Acknowledgment is a peculiar species of knowledge. I acknowledge another when my utterances convey not just facts but also an appeal to respond. The word stems from 15th- century Middle English, *aknow*, meaning an admission of one's knowledge. I confess my words, as it were, and make claims on another. As Cavell put it, "acknowledgment is more than knowledge, it includes an invitation to action" (Cavell 1969, 259. In the least, that action enjoins another to respond in kind and continue the conversation. At the most, my words' action makes an ethical appeal. When I feel pain in front of a friend, for instance, I might wince, groan, and grab the hurt part of my body. My speech does not convey a desire that my pain be known.

(This is the sense in which Wittgenstein remarked, "It can't be said of me at all (except perhaps as a joke) that I *know* I am in pain" (Wittgenstein, 1953, §246). I seek my friend's acknowledgement: that he responds attentively, asks how I feel, and gestures to help.

The passage from rule-based to AI-powered language games, viewed from the vantage of linguistics, hinges on the different speech acts at play. Games of knowledge assess users' constative utterances: words whose description of a state of affairs is true or false. These are generally statements ("He attends the wedding") with meaning and lexical qualities that are assessed according to their truth. With AI-powered conversations, Duolingo empowers users also to make performative utterances, which, as J.L. Austin showed, do not describe but bring about a state of affairs (Austin 1962). These are statements (to use Austin's favorite example, "I take thee to be my lawfully wedded wife") whose force implicates another. They enjoin the other's acknowledgment in order to create a new truth. In the case of "Roleplay," performative utterances impel the chatbot to continue the conversation in an intended direction.

As AI-powered conversation approximates natural conversation, conversations between users and chatbots fall victim to the infelicities that already haunt our struggles to acknowledge each other. In natural conversation, my pronouncement of marriage might fall apart when spoken in the wrong context (perhaps I lack the authority of a rabbi, priest, or civil servant) or when the other leaves my words unacknowledged (and never says "I do"). Even when the utterances spoken might be empirically valid, unforeseen turns and shifting contexts can equally lead conversation with a chatbot to errors of acknowledgement.

Errors of acknowledgment could also be the result of willful manipulation. Take as an example Air Canada's customer service chatbot, which invented a bereavement refund policy thanks to a conniving user's back-and-forth conversations with the bot (Matsakis 2023). He claimed that he was entitled to a free ticket because he was traveling to attend a grandparent's funeral. After Air Canada rescinded the reimbursement, a Canadian tribunal ruled that the company had to enforce the policy. The customer kept a screenshot and demanded the refund that had been promised. We find humans interacting with AI platforms as if they were humans – testing or coaxing chatbots. The history of computing is rife with malfeasance. Humans hack systems to steal information or falsify documents using text and image processors. But with AI platforms, we might imagine that users no longer break into a bank; a manipulative user instead persuades the guards to unlock the safe. In this

case of Air Canada, the bereavement policy was not lying in wait. It was the product of perlocutionary utterances between chatbot and user.

The errors of knowledge that attract much attention in popular and academic accounts of AI-powered conversation should not occlude the errors of acknowledgement that arise not simply due to mistakes of engineering. Acknowledgment remains a delicate struggle in everyday interactions. Neither human nor machine are entirely to blame. If solutions to such errors of acknowledgement have proved elusive, whether by modifications of interface design or back-end engineering, that is in large part because the infelicities of human interaction magnify in artificial interaction. Intentions can misalign. Either the chatbot misinterprets the user's intentions or the user confounds those of the chatbot. Perhaps, still, a mischievous user flippantly beguiles an AI-powered conversation. The potential for acknowledgement to go awry knows no bounds.

4. Language Games at The Heart of Our Intentions

Acknowledging another's intention is rarely instantaneous. There is a back-and-forth interaction by which we express our thoughts and solicit the other's response. To be sure, an AI-powered chatbot does not acknowledge a user's intentions *tout court*. A machine cannot re-*cognize* a user's words, tone, or facial expression with the emotional immediacy that humans of a common culture perceive in each other. AI models make predictions about users' intentions on the basis of learned associations found in training data. Duolingo's "Roleplay" nonetheless registers the linguistic signs of users' intentions thanks to Natural Language Understanding (NLU). Deployed by the transformer model, NLU concatenates patterns of intentional statements by tokenizing inputs across training data and embedding the tokens in vectors representing their relationships to other words. Context management finalizes the prediction. NLU tracks each turn in a dialogue to maintain conversations' coherence over time. Although AI-powered conversation relies on prediction, it cannot be said that predicting the intentions in users' natural language input is a small feat.

The intentional interface in Duolingo's "Roleplay" marks a shift from the command-based interfaces of modern computing (Norman 2023). Command-based interaction began in 1963 with the Teletype Model 33 ASR teleprinter, which registered users' lines of text via keyboard and punch tape. Originally designed for the US Navy, the teleprinter was soon used to send and receive messages in company offices. Command-based interactions became digital thanks to Graphic

User Interfaces (GUIs). Alto by Xerox PARC enabled users in 1973 to issue commands by clicking visual elements. Apple's Lisa (1983) and Macintosh (1984) inaugurated a world in which users click a mouse to engage windows, icons, and menus. AI changed everything. Instead of executing tasks via discrete commands, LLMs facilitate interactions by which users input what they want in the form of natural language. Duolingo's "Roleplay" deploys the technology to weigh mean reductions and interpret individual users' input against the collective input of user interactions. As a result, the app's underlying GPT predicts users' intentions.

Turn-taking with an AI-powered intentional interface such as Duolingo unfolds through time. Much the same, natural conversations advance through time when conversants speak intentionally. That is, we share more than bare facts with each other. Hardly innocent, the nature of intentionality is the subject of longstanding debates in the philosophy of language. If we think of users' back-and-forth conversations with chatbots as anything like natural conversations – rife with the hazards of knowledge and acknowledgment – then everything hinges on how we think of intentions. That is especially the case if natural conversation is to yield lessons that can improve the design of intention-based interfaces, as I have suggested. After all, intentionality establishes the conditions of a conversation's success or failure.

Consider three approaches to intentionality. For Aristotle, intention involves purposeful reasoning. All animals experience desire (appetency) but humans stand apart for the capacity to rationally deliberate (βούλησῃ) about the means to fulfill our desires (Aristotle 1999). With prohairesis (προαίρεσις) we act intentionally and make choices that lead to a desired end. This line of thinking established in Western thought the principle that intention consists of rational action.

A second approach considers intention to be a mental state. When psychology took shape in the late nineteenth century as an experimental and clinical science, Franz Brentano contended that intentionality is the "mark of the mental" (Brentano 1874). Beliefs and desires direct the mind toward the world in a way that objects do not; they lack the intentional relation to believe or desire a state of affairs. Advancements in brain science in the twentieth century drove John Searle to suggest that intentionality is the mind's causal power to direct actions (Searle 1980). Although the second approach shares much with the folk concept that intentions motivate decision-making, it's difficult to apply to user interactions with conversational AI because computers lack the ability to access the thoughts and desires anterior to the words that users type in an interface. All that is available for natural language processing is the user's input.

A third approach is that intention depends on the order of reasoning. Elizabeth Anscombe suggested that intentions are not mental states that precede our actions (Anscombe 1957). We act intentionally when we are prepared to answer the question, "why?" Anscombe took as an example a man moving his arm up and down while holding a pump. We might ask, "why is he moving his arm?" We arrive at his intentions by explaining that the man acts *because* (and not just that) he is operating the pump, which he's doing because (and not just that) he's pumping water to the house. The order of the arm's actions up and down really exists in the world; it is not the sum of merely physical events nor anterior thoughts or desires. *This* order of actions, connected by giving and taking reasons, constitutes the reasoning involved in intention.

So too does the order of reasons between chatbot and user make it possible to achieve a mutual acknowledgment of each other's intentions. That the chatbot's output follows my input, and in turn, that I extend the interaction by building subsequent input on prior output, makes a successful conversation possible. For H.P. Grice, we implicate each other. Conversations consist of more than "a succession of disconnected remarks…They are characteristically, to some degree at least, cooperative efforts; and each participant recognizes in them, to some extent, a common purpose or set of purposes, at least a mutually accepted direction" (Grice 1975, 45). The order of reasons, as Anscombe considered them, can be understood in Grice's remark to take the shape of a "mutually accepted direction." Errors of acknowledgment arise when intentions misalign; conversants express what might be otherwise factual claims without the shared connections expected in an intentional exchange.

Errors of acknowledgment are not a bug but a feature of everyday conversation. When conversants fail to connect – when we fall short of cooperative efforts and our words crumble into disconnected remarks – the urge ensues to flee or fasten the conversation. On the one hand, I might seek an escape from the frustration of elusive intentions (perhaps yours appear opaque or my own prove difficult to convey). On the other hand, I might seek iron-clad knowledge of the other's intentions; to peer behind my interlocutor's words, as it were, and grasp the underlying mental picture. Yet, the wish for a perfect communication, a complete transparency between minds, is a fantasy; it masks the reality that language makes our community possible as much as it divides us. Deprived of acknowledgment in a shared space, we seek knowledge as consolation in the recesses of another's mind. What we seek are linguistic rules: answers to our doubts that are independent, finite, and discrete. Confronted with a deficit of acknowledgement, we might wish that our

broken conversation could be fixed if what stood between us were merely a deficit of knowledge. For Anscombe, Austin, Cavell and other thinkers who wrote in the wake of Wittgenstein's theory of language games, the precarity of natural conversation drives in us the ineluctable desire to step beyond an exchange of reasons, to attain rigid – yet illusory – rules and thereby regain our bearings in the face of uncertainty.

Users can feel a similar urge to regain their bearings when they encounter errors of acknowledgment in AI-powered language games. Conversations in "Roleplay" can stray from the intended topic or expose new contexts that bring words' meaning into question. As technological progress brings AI-powered conversation platforms into closer proximity with natural conversation, they confront the uncertainties that cause human interaction to slip and leave us desiring words' traction. In Duolingo, we find a mix of language games – games of knowledge and acknowledgment – that evince the possibilities and pitfalls of natural conversation.

5. Designing AI Chatbots to Deepen Conversation

Until 2023, users played only rule-based games on Duolingo and Babbel. Close-ended exercises would transport users from the dynamism of everyday dialogue in order to make moves that are finite and discrete. With the integration of GPT-4 on Duolingo, continuous conversation launched "Roleplay" into the orbit of natural turn-taking communication. Yet, desire for close-ended features with finite and discrete designs persisted. Clickable individual words as well as drag-and-drop vocabulary offer relief from the uncertainties and infelicities that can arise when users feel lost in continuous conversation. "Roleplay" includes design features that provide footing when the grounds of the game slip away. The language game's interface offers a mix of rigid rules and dynamic interactions.

Facets in "Roleplay" appear beneath the dialogue frame and provide filtered navigation when a user is unsure how to proceed. The feature displays a finite set of words to select and add to the dialogue. The facets are spatially positioned outside the conversation, providing rule-based relief, as it were, when users are unsure what to say. When errors of acknowledgment arise in natural conversation, we might desire a script to know which moves to make. The urge, although unattainable in natural conversation, remains a real and insistent impulsion that Duolingo visualizes in the form of faceted conversation. Roleplay thereby offers knowledge when users face errors of acknowledgment.

Figure 7. Facets in Duolingo's "Roleplay"

A second design feature is translation overlays, which appear when users hover over foreign words in the dialogue frame. At times when the user is unsure of words' meaning, especially when multiple meanings might be at play, he or she can make the English translation appear. It opens an escape hatch, an opportunity to flee the uncertainty of translation, and momentarily return to a language (game) whose rules are known. Unsure of another's intentions in natural conversation, we might desire telepathic powers to know what he or she *really means* behind the veil of language. "Roleplay" accommodates the desire with translation overlays.

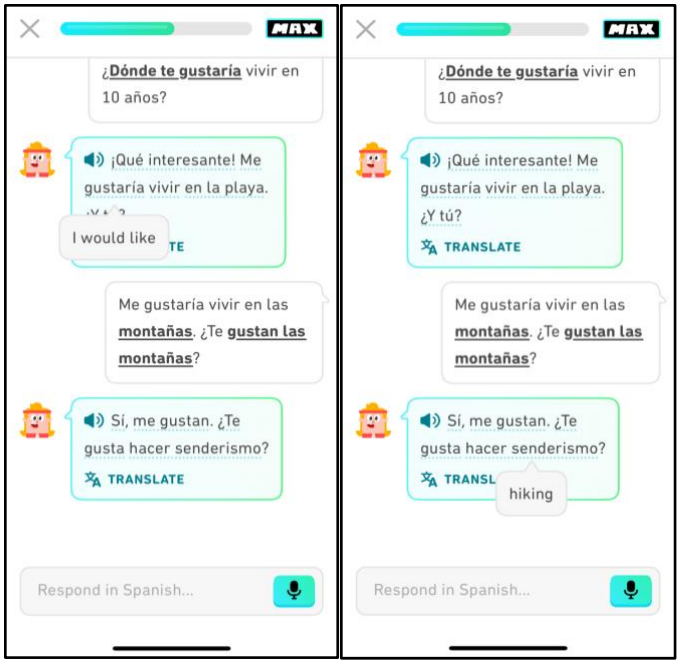

Figure 8 a–b. Translation overlays in Duolingo's "Roleplay"

Both design features make visible what language hides. They respond to the temptations of a transparent communication liberated from the opacities and illusions of language. Whereas misunderstanding and incomprehension can leave us groping to secure acknowledgment in day-to-day interactions, AI interactions secure traction for understanding and comprehension. Facets and translation overlays make text clickable. The errors of acknowledgment implicit to natural conversation appear as knowledge explicit in artificial conversation. As a result, the features fulfill the illusory – yet inescapable – urge for a transparent conservation devoid of precarity. Language games with chatbots rise above the games we always already play in language.

We might quit playing language games with our heart, and start playing language games with chatbots, when AI-powered conversation combines rule-based, clickable options with the dynamic realism of back-and-forth conversation. The future of AI conversation, I'm suggesting, is one in which the rule-based interactions of traditional decision trees co-exist with intention-based interactions powered by generative AI. We are just beginning to witness AI platforms' capacity to mimic the continuity of natural conversation; progress is to be made by learning from the pitfalls of human communication as well. With AI platforms as our interlocutors, companionship takes many forms. Users might learn a skill with Claude and treat the

platform as an educator. Others might test a thesis with Chat-GPT and use it as a debate coach. Perhaps we explore a niche topic with Gemini and interact with it as we would with a librarian. Across diverse language games, conversational AI should bring to light the ineluctable desires churning beneath the surface of natural conversation; to thereby suture artificially what our natural words, as the Backstreet Boys sang, can break asunder:

> I should have known from the start,
> You know you have got to stop,
> You are tearing us apart,
> Quit playing (language) games with my heart.

References Cited

Anscombe, G.E.M. 1957. Intention. Basil Blackwell.

Aristotle. 1999. Nicomachean Ethics, translated by Terence Irwin, 2nd ed. Hackett Publishing Company.

Austin, J.L. 1962. How to Do Things with Words, edited by J.O. Urmson. Oxford University Press.

Babbel. 2024. "About Us." Babbel, September 16. https://uk.babbel.com/about-us

Babbel. 2023. "Learn with Your Own Voice: Babbel Launches Two New Speech-Based Features." Babbel, May 5. https://www.babbel.com/press/en-us/releases/learn-with-your-own-voice-babbel-launches-two-new-speech-based-features-us#

Brentano, Franz. 1874. Psychology from an empirical standpoint, edited by A. C. Rancurello, D. B. Terrell, & L. L. McAlister (1995). Routledge.

Cavell, Stanley. 2004. Cities of Words: Pedagogical Letters on a Register of the Moral Life. Harvard University Press.

Cavell, Stanley. 1969. Must We Mean What We Say? Cambridge University Press.

Chomsky, Noam. 2023. "The False Promise of ChatGPT." The New York Times, March 8, 2023. https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html

Duolingo. 2023a. "Large Language Models + Duolingo = More Human-Like Lessons." Duolingo Blog, April 12, 2023. https://blog.duolingo.com/large-language-model-duolingo-lessons/

Duolingo. 2023b. "Introducing Duolingo Max, a Learning Experience Powered by GPT-4." Duolingo Blog, March 14, 2023. https://blog.duolingo.com/introducing-duolingo-max-powered-by-gpt-4

Freeman, Cassie, Audrey Kittredge, Hope Wilson, and Bozena Pajak. 2023. The Duolingo Method for App-Based Teaching and Learning. Duolingo Research. https://duolingo-papers.s3.amazonaws.com/reports/duolingo-method-whitepaper.pdf

Grice, H.P. 1975. "Logic and Conversation," Syntax and Semantics 3: Speech Acts. Academic Press.

Ji, Ziwei, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii. 2023. "Survey of Hallucination in Natural Language Generation." ACM Computing Surveys 55 (12): Article 248. https://doi.org/10.1145/3571730

Jiang, Xiangying, Jenna Rollinson, Luke Plonsky, Erik Gustafson, and Bożena Pająk. 2024. "Evaluating the Reading and Listening Outcomes of Beginning-Level Duolingo Courses." Foreign Language Annals 54 (4): 974-1002. https://doi.org/10.1111/flan.12600

Kosinski, Michal. 2023. "Theory of Mind May Have Spontaneously Emerged in Large Language Models." arXiv Preprint arXiv:2302.02083. https://doi.org/10.48550/arXiv.2302.02083

Li, Xinyuan. 2023. "'There's No Data Like More Data': Automatic Speech Recognition and the Making of Algorithmic Culture." Osiris 38 (1): 165–82

Matsakis, Louise. 2023. "Air Canada's Chatbot Soars Out of Control, Invents Refund Policy," Wired, February 6. https://www.wired.com/story/air-canada-chatbot-refund-policy/

McDowell, John. 1984. "Wittgenstein on Following a Rule," Synthese 58 (3): 325-363.

Norman, Donald. 2023. "AI: First New UI Paradigm in 60 Years," Nielsen Norman Group, June 18. https://www.nngroup.com/articles/ai-paradigm/

Packer, J.D., and James Reeves, eds. 2023. Abstractions and Embodiments: New Histories of Computing and Society. Baltimore: Johns Hopkins University Press

Sacks, Harvey, Emanuel A. Schegloff, and Gail Jefferson. 1974. "A Simplest Systematics for the Organization of Turn-Taking for Conversation." Language 50, (4): 696-735

Saleem, Awaz, Narwin Noori, and Fezile Ozdamli. 2022. "Gamification applications in E-learning: A literature review." Technology, Knowledge, and Learning 27: 139–159. https://doi.org/10.1007/s10758-020-09487-x

Searle, John. 1980. "Intention and action," Philosophical Perspectives 4: 79-105.

Summers-Stay, David, Christopher R. Voss, and Scott M. Lukin. 2023. "Brainstorm, Then Select: A Generative Language Model Improves Its Creativity Score." In The AAAI-23 Workshop on Creative AI Across Modalities.

Tommerdahl, Jodi, A. Olsen, and C. Dragonflame. 2022. "A Systematic Review Examining the Efficacy of Commercially Available Foreign Language Learning Mobile Apps." Computer Assisted Language Learning 37 (1): 1-30. https://doi.org/10.1080/09588221.2022.2035401

Wakabayashi, Daisuke. 2022. "Google Engineer Says AI Has 'Consciousness' in Remarkable Claim." The New York Times, July 23, 2022. https://www.nytimes.com/2022/07/23/technology/google-engineer-artificial-intelligence.html

Wittgenstein, Ludwig. 1953. Philosophical Investigations. Translated by G.E.M. Anscombe (2001). Blackwell.

Seeing It from Other Eyes: How First-Person Data Reshapes the Role of the Applied Ethnographer

MARIA CURY, *ReD Associates*
ERYN WHITWORTH, *Meta Reality Labs Research*
ARIEL ABONIZIO, *ReD Associates*
GABRIEL COREN, *ReD Associates*
MAYA POTTER, *ReD Associates*
TAMARA MOELLENBERG, *ReD Associates*
MIKKEL KRENCHEL, *ReD Associates*

This paper argues that first-person data—video, audio, and other data that is recorded from the research participants' point of view through wearable technology—affords new possibilities for anthropology and the social sciences to capture and work with longitudinal, immersive, behavioral data at scale. We position first-person data capture as an ethnographic method, complimentary and adjacent to others. Collecting detailed data on what people see, hear and experience through wearable technology opens a new way of understanding people, communities, and societies that has the potential to fundamentally change the social sciences as researchers gain access to nuanced, real-time interactions and behaviors. We also show the value of ethnography in informing the development of wearable technology, as well as engineering and research to develop next-generation technology across areas like computer vision and foundational AI. The paper draws on a series of projects the authors conducted for Meta Reality Labs Research during 2018–2024, which resulted in 1000+ hours of first-person footage that were analyzed alongside participant observation, semi-structured interviews, and diary studies.

Introduction: The Potential of First-Person Data

Ethnography, the backbone of anthropological research, has traditionally relied on methods such as participant observation and interviews to delve deep into the lived experiences of individuals and communities. In the context of applied research, these methods have provided rich, nuanced understandings of human behavior and social dynamics. Over the past decade, the advent of wearable technology offers new opportunities to enhance and expand these traditional methods. In this paper, we argue that first-person data—video, audio, and other sensing input recorded from the participants' point of view using wearable devices—can be a powerful ethnographic tool that not only complements but also transforms and elevates the practice of applied ethnography.

Cury, Maria, Eryn Whitworth, Ariel Abonizio, Gabriel Coren, Maya Potter, Tamara Moellenberg, and Mikkel Krenchel. "Seeing It from Other Eyes: How First-Person Data Reshapes the Role of the Applied Ethnographer." *EPIC Proceedings* (2024): 276–315. https://www.epicpeople.org/seeing-it-from-other-eyes-first-person-data. ISSN 1559-8918.

First-person data, sometimes also termed egocentric data, refers to information captured from the perspective of the individual experiencing it. This type of data is collected through wearable devices such as action cameras, body cameras, smart glasses, and audio recorders, providing an immersive and situated viewpoint of the wearer's environment and interactions. First-person data can take the form of images and videos that capture visual aspects of the wearer's surroundings from their viewpoint, audio recordings that document conversations and ambient sounds, and self-narration where individuals verbally recount their thoughts, feelings, and experiences as they occur, from their point-of-view.

Figure 1. Three still-images of first-person video data captured using a wearable device during one of our studies; all are from the embodied vantage point of the person wearing the camera, simulating what they are seeing. The left image shows their hand holding a book as they are seated, the center image shows their hands at a kitchen sink washing a pan, and the right image shows them drawing on a tablet at a desk covered in art supplies.

Although wearable recording devices have existed for years (e.g., GoPro, Meta Ray Bans), the use of these kinds of devices for research has largely remained contained to fields like cognitive science, computer vision, and ergonomics, the data analyzed mainly quantitatively and often for the purposes of improving the underlying technology. However, recent developments in wearable technology—from lighter form factor to longer battery life to AI-enabled software for processing and analyzing the captured footage—have made first-person data capture methods more accessible and practical for applied ethnographers.

This paper examines the potential of first-person data and how it might transform and expand the role of applied ethnography in the future. We argue that first-person data affords new possibilities for anthropology and for the social

sciences more broadly, to capture and work with longitudinal, immersive, behavioral data at scale. This paper has the dual purpose of reframing first-person data as an ethnographic method and showing the value of ethnography in informing the development of nascent wearable technologies. Collecting detailed data on what people see, hear and experience through wearable technology opens a new way of understanding people, communities, and societies that has the potential to fundamentally change the social sciences as researchers gain access to nuanced, real-time interactions and behaviors previously beyond reach. The rise of this kind of first-person point-of-view data, viewed through the ethnographic lens of understanding people, culture, interpersonal dynamics, meaning, and human experience, also invites applied ethnographers to engage with and influence the core engineering and research sciences emerging from fields like computer vision and foundational AI—and we'll argue, to help shape how and in what directions next generation technologies like these develop.

We draw on methodological learnings accumulated from projects the authors conducted for Meta Reality Labs Research between 2018-2024, which used first-person data capture as a method to inform the development of next-generation wearable devices. These projects resulted in 1000+ hours of first-person footage, which were analyzed alongside participant observation, semi-structured interviews, and diary studies. Instead of detailing the findings from these individual studies, our focus in this paper is to position first-person data capture as a method that is adjacent and complementary to traditional ethnographic methods. Additionally, we show how both traditional and egocentrically augmented ethnographic methods have a critical role to play in the development of nascent wearable technologies.

There are Four Objectives to This Paper

First, we aim to lay a provisional foundation for ethnographers to experiment with first-person data capture as a technologically mediated method. By integrating first-person data with methods like participant observation, contextual inquiry, surveys, interviews and diary studies, our paper argues that first-person data collection has the potential to become a critical tool in any social scientists' toolbelt—and invites other ethnographers to experiment with and further define this nascent method. We frame first-person data capture as a new and complementary method to traditional ethnographic methods that have had a long lineage of experimentation with different kinds of media and points of view in anthropology, and as a method that allows great depth and breadth of data collection and elicits new forms of ethnographic analysis. We do this by providing an overview of the use

of wearable technology and first-person data in research within other disciplines and within anthropology specifically, and by outlining the approaches we took to collect first-person data across our studies.

Second, we aim to discuss the methodological benefits surrounding the use of first-person data capture devices in ethnographic research. Whereas, at first glance, wearable devices might appear "participant-less" and voyeuristic, we argue that researcher participation is taking a different form—the researcher is not absent, but their presence is transposed into the wearable glasses, carrying its own positionality. The displacement of the researcher into wearable recording devices affords new modes of engagement. For instance, we argue that researchers gain the ability to participate in moments of participants' everyday lives that would have been unavailable to them had they been physically present—intimate conversations with loved ones, heated family discussions, and late-night insomnia bouts. Moreover, we argue that first-person data enables researchers to gain a new temporally grounded perspective of the lives of their participants by watching their context and lives unfold from their perspective. Unlike traditional ethnographic methods, which center the research participant as the subject of study, first-person data centers the participants' context, as seen by both the researcher and the participant. This approach challenges the conventional dichotomy of observer and observed, suggesting a hybrid model where observation has the potential to become a shared, co-constructed experience, especially when the data is analyzed in co-constructed ways as we will outline as a core tenet of working with this kind of data. Some of the highlighted benefits include the ability to capture longitudinal data of different kinds (audio, visual, including objects, environments, people), the scalability of the data collection, and the emphasis on showing what people do, not just what they say, and more intimate reflections in what they say. We argue that this can unlock insights with a greater scale of validation, with the ability to inform what drives behavior and behavior change, and with the ability to shed light on the gaps in understanding that might exist between, e.g., institutions and individuals.

Third, we lay out some of the risks of this approach. We highlight how the frame of a video or still image risks inflicting data with a dangerous aura of objectivity, reifying research participants' lives. Additionally, there is a risk that, in focusing on a first-person perspective, the research findings default to an individualistic view of the world. There is also the risk of privacy and confidentiality, given the immediacy and intimacy that first-person data affords. To address these risks, we posit that, like other ethnographic methods such as participant observation, first-person data capture is a method that is still historically situated and contextual. First-person and

ethnographic data are not just captured but rather constructed in the act of observation and sensemaking. First-person data, in that way, continues to be a form of situated and relational knowledge: the viewpoint, however technologically mediated, continues to be from somewhere and it carries its own biases and limitations. First-person data still requires analysis, and indeed even a new analytical skillset. We outline some of the emerging elements of this analysis.

Fourth, we aim to explore the implications of this new method for applied ethnographers – for the kinds of skillsets ethnographers would benefit from having (e.g., technical literacy, analysis triangulating different kinds of data) and for the kinds of impact ethnographers are poised to have (e.g., new ways to work at scale, new ways to address behavior change). We also argue for the presence of ethnographers in early-stage wearable technology development: We show that ethnographers have a role to play informing the very kinds of sensors and capabilities that next-generation wearables should have to collect and surface data that is meaningful to people. The ethnographer's first-person data collection centers participants' experiences, social interactions, real-world context, and inner states, and applies anthropological frameworks to these, in ways that may be novel and beneficial to the engineers and research scientists who are often already collecting forms of first-person data (e.g., in labs, for dedicated tasks) to shape how these technologies develop. This collaboration allows ethnographers to become critical nodes in the design of high-quality, specialized data collection circuits, rooting innovation in real-world human experiences, and potentially leading to more empathetic, useful, and socially meaningful technologies. Doing so requires greater literacy in how AI systems work, new partnerships with engineers and developers, and new types of outputs that are legible and relevant to early-stage product development.

We conclude by reflecting on what it might take for first-person data collection to become a mainstream research practice for applied ethnographers, not only in terms of technological readiness, but also the skills, costs, ethical frameworks, and partnerships that will have to be developed by the community. As we navigate an era where wearable technology increasingly intertwines with daily life, this exploration hopes to serve as a call to action for more discussion about, and experimentation with, how first-person data might both expand and transform the role of applied ethnographers.

Overview of Wearables and First-Person Data in Research

The use of first-person data in research has a history that spans several decades, evolving in tandem with advancements in wearable technology. First-person point-

of-view footage has been used extensively in scientific research, predominantly within medical research, with GoPros capturing surgery footage as one prominent example (Ganry et al. 2019; Graves at al. 2015; Huang et al. 2024). In tech-forward domains such as computer vision, robotics, and ergonomics, first-person data has been extensively explored to enhance machine learning algorithms, improve hardware interfaces, and develop AI systems; these studies have often focused on informing the design of wearable devices.

Beyond the scope of technology development research, first-person data has been used to observe naturalistic behaviors and contexts in situ since the 1960s and 1970s. In ergonomics, for example, wearable cameras have been used to analyze workers' movements and interactions with their environment, leading to improvements in workplace design and safety protocols (Charbonneau and Doberstein 2020; Rane et al. 2023). Similarly, in health sciences, wearables have facilitated the monitoring of patients' daily activities and behaviors, contributing to better personalized healthcare solutions (Tyler et al. 2020). In psychology and cognitive science, first-person data provided an unprecedented window into human behavior outside of laboratory settings (Yoshida and Smith 2008; Tangen et al. 2022). And in the environmental sciences, teams of animal behaviorists have employed egocentric video and audio capture along with wearables and sensors to document real-time naturalistic behavioral patterns of difficult to access—and often endangered—species, such as sea turtles (Seminoff et al. 2006). For these fields, the integration of wearable technology into research methodologies expanded the possibilities for capturing rich, real-time data, enhancing the depth and breadth of behavioral studies.

Recent technological advancements have made collecting first-person data in research increasingly possible and feasible, as wearables become smaller, their battery lives become longer, and the quality of the recording increases. Numerous companies are heavily investing in wearable technologies that capture first-person data, including devices like the Tab, Humane AI Pin, Limitless Pendant, Rabbit, Meta Ray Ban, the Tobii Pro Glasses and Meta's Project Aria – the latter two which are used for research purposes. These wearables promise often continuous, real-time data from the user's perspective, enabling detailed documentations of daily activities and interactions.

Combined with AI capabilities to process the vast amounts of data collected (e.g., transcript and summary generation, object and scene recognition, labelling), first-person data opens new kinds of analyses. Major corporations are developing AI models capable of visual analytics, promising capabilities such as object detection,

action classification, and contextual understanding of user intent. Notable examples include IBM Watson Visual Recognition, Google Video AI, Microsoft Azure Video Analyzer, and Amazon Rekognition Video. Whereas wearable devices facilitate new forms of data collection, emerging software with AI tools promises to analyze this data in ways previously unattainable through traditional methods.

Despite the growing body of research using first-person data capture in various tech-forward fields, and the growing possibilities enabled by technological advancements, a significant gap remains in its application within the social sciences, particularly in ethnography. The social sciences have been slower to adopt these technologies, and there is a paucity of literature that examines the methodological, theoretical, and practical implications of first-person data capture for ethnographic research. We explore some precursors and parallels to this kind of data, in the following section.

This raises some key questions: what does this new kind of data mean for applied ethnographers? What kinds of analyses become possible through ethnographic approaches to first-person data? This exploration challenges ethnographers to think beyond conventional methods and consider how first-person data can be harnessed to provide richer, more nuanced understandings of human behavior and social interactions.

In the sections below, we recontextualize first-person data as a valuable tool for applied ethnography, rooting it within existing discussions in anthropology and exemplifying how first-person data was collected and analyzed in recent applied ethnographic studies we conducted. Then, we highlight some of the methodological benefits and limitations of first-person data as a method. We hope to show how this approach not only broadens the scope of questions that ethnographic research can address but also positions ethnographers to contribute more significantly to the design and development of emerging technologies.

Precursors and Parallels to First-Person Data as Ethnographic Data: Experimentation with Positioning, Point of View, and Media Across Anthropological Practice

First-person data presents a profound shift in ethnographic research methods and goals. Yet although it seems radically new, anthropologists have long experimented with emerging technologies for data capture in the field, with positioning themselves relative to their subjects, and with exploring diverse recording methods and media. We posit that first-person data possesses precursors and parallels to historical approaches of capturing ethnographic data. We here examine

examples of fieldwork, autoethnography, visual anthropology, and sensory ethnography to position first-person data within a lineage of ethnographic practices, and to show that experimentation with media technologies to capture human experience is a staple of anthropology.

Fieldwork's Inauguration Was about 'Being There' to Collect a New Form of Data about Human Experience—and the Role of the Ethnographer was Partly to Provide a First-Person Point of View

When in the late 19th century ethnologists began fieldwork to study non-European cultures, neither the data they captured nor their methods for collecting it had been considered adequate for exploring ethnographic questions. Only in the early 20th century, Franz Boas and Bronislaw Malinowski formalized fieldwork as a core anthropological method. While Boas codified the culture concept that shaped academic anthropology, Malinowski, marooned in Melanesia, emphasized immersive fieldwork—being there with one's subjects, immersed in their worlds, observing their practices firsthand—as a foundation for ethnography (Boas 1911, Introduction; Malinowski 1922). Malinowski highlighted the distinction between raw data and its interpretation, positioning the ethnographer as the principal medium for recording and rendering ethnographic data.

By 1930, the ethnologist and surrealist intellectual Michel Leiris also argued that seeing firsthand is the way of making sense of strange surrounds, and that the ethnographer's eye, their first-hand witnessing, is like a skein that mediates between familiar and strange, self and other, universal and particular, self and world (Leiris 1930).

This approach made the ethnographer's experience integral to the aptness—and the authority—of their data, introducing reflexivity into ethnographic research. While traditional ethnographic monographs remained dominant, post-structuralist literary theory and decolonial contexts led to a reevaluation of ethnographic representation among a post-war generation of anthropologists. Clifford Geertz's *The Interpretation of Cultures* (1973) elevated ethnographic data to layer up into a "thick description" of the symbols and meaning that play out in cultural contexts – establishing a new standard.

Geertz' students, such as Paul Rabinow and Keith Dwyer, further pushed ethnography by bringing their subjects' voices directly into ethnography itself, triggering a wave of criticism about 'raw' ethnographic data and its purposes (e.g., Clifford and Marcus 1986; Crapanzano 1986; Dwyer 1982; Fabian, 1983; Nader 1972; Ortner 1974; Rabinow 1977). A hard-won disciplinary unity about what

ethnographic data was—as well as how it was to be collected and given form—dissolved. But the possibilities for what ethnographic research and representation could ultimately be and become radically opened as well. Since then, within academic anthropology but increasingly in applied anthropological contexts as well, what ethnographic data could teach us and others about our changing world, the variety of sites, spaces, and subjects it could be drawn from, and the kinds of questions and problems it could be relevantly applied to, has expanded greatly. This is especially true when it comes to emerging phenomena at the intersection of science, technology, and society.

The use of personal experience as a primary source of data for making sense of cultural phenomena has been part of the ethnographic enterprise at least since Malinowski (and through and beyond his "children"). Carolyn Ellis' work, "The Ethnographic I: A Methodological Novel about Autoethnography" (2004), however, gave it a name. The practice has been in many ways perennial: for example, Franz Boas' student, Zora Neale Hurston, was the first anthropology graduate student to return to the community in which she was reared in order to conduct a form of self-study, leveraging one's identity and experience to gain access to—and to apply one's insider/outsider positioning among—a peoples and a place one intimately knows (Hurston 1935, 1938).

Working with first-person data has a kinship with self-study, in that it also relies on a first-person point of view to generate rich, contextual data. However, unlike autoethnography where the researcher's first-person perspective is primary, first-person data as we explore in this paper allows for the centering of participants' perspectives, which expands the scope of what is possible ethnographically but shifts the role of the ethnographer with respect both to that data and to the subjects whose lives it derives from as well.

Despite Text's Predominance for Presenting Ethnographic Findings, Emerging Media Technologies Have Been Vital Tools for Capturing Ethnographic Data and Experimenting with Perspective

From the start, ethnologists carried cumbersome media technologies into the field as instruments of documentation, integrating visual media (film, photography) into their processes of collecting and rendering ethnographic data. Of this overlap of the ethnographer's eye with the camera's lens, Michaela Schaüble observes that it "not only indicates that the two forms of data and knowledge acquisition are highly compatible but also shows that the pioneers of the discipline recognized the

potential of the film and photo camera as a technical extension of the ethnographic eye very early on" (Schäuble 2018, 1).

Documenting ethnographic data visually was about more than equipping ethnographers with technical protheses. By the 1930s, photographic technologies were serving at least one of two additional functions for the ethnographic fieldworker, which we continue to find elements of in first-person data capture technologies. First, as a method for visual fieldnote taking— a kind of visual duplicate to verbally dictated ethnographic data, rendered in a different form. Second, as a critical authority-establishing device, intended for the publication of ethnographic findings following fieldwork. Photos from the field possessed an 'I-was-really-there' aura which wooed audiences—academic and popular—from the onset. Margaret Mead and Gregory Bateson exemplified the visual fieldnote method by carrying film and cameras to Bali and New Guinea from 1938 until the late 1940s, rigorously recording all the while developing an analytical armature for integrating visual data into ethnographic research and reporting (as outlined in, for example, their introduction to Balinese Character: A Photographic Analysis (1942)). Claude Levi-Strauss helped popularize the 'I-was-there aura' when he carried his own photographic equipment into the Amerindian interiors of South America, the documentary artifacts for which were exhibited to much fanfare (and critique) at the Musee L'Homme in Paris.

Ethnographic film, however, had from the start less academically acceptable purposes than photography. Unlike still images, moving pictures could not quite be cast as back-up note-taking nor quite establish an 'I-was-there' authority for the ethnographer, since unlike photographs, film couldn't be integrated into the textual media format of ethnographic monographs. Perhaps ethnographic film's resemblance to the media of popular, mass entertainment—rather than to contemporary artistic media—had something to do with its casting aside by the bulk of academic anthropology's establishment.

Nonetheless, many imaginative ethnographers recognized film as technology for ethnographic research of a uniquely valuable kind. David MacDougall (2006) has glossed video ethnographic data as uniquely suited for examining "temporal," "corporeal," and "personal" aspects of human experience – aspects which we position first-person data (often in video form) as uniquely suited to examine as well today. Marcel Griaule's filmic studies of the Dogon in Mali Griaule showed what was possible along these lines and inspired many others. One exemplary legatee of Griaule's, from a generation later, was Jean Rouch, who mobilized ethnographic filmmaking to explore genre-defying ways of capturing, engaging with, and

presenting ethnographic data while also instantiating what he called "ethnofictions," or what many ethnographic filmmakers today call "the participatory approach" (see Taylor 1996). Of note is also the enduring power of Robert Gardner's work, especially Dead Birds (1964), and the breathtaking, if controversial, collaborations between film maker Tim Asch and anthropologist Napoleon Chagnon, taken of and among the Yanomami throughout the 1960s and 1970s (Asch 1961, 1974; Asch and Chagnon 1971, 1974, 1975; Chagnon 1974; Harper 2003).

While still photographs are artifacts of an encounter, when audiences encounter photographs — embedded and printed with captions alongside ethnographic texts or hung on a museum's walls — they tend to experience these as if they were raw, unmediated renderings of culture captured in action. Despite being by-products of an individual's experience, still photographs nonetheless can easily feel like remnants of an objective third-person observation. Ethnographic films, in contrast, cannot conceal that they are made with media consciously introduced into an ethnographic encounter and carefully reconstructed afterwards in editing. The process is visible in the product: that is, the raw data and its interpretation can't be untangled. Although academic anthropology has tended toward discomfort with this fact and often honored that discomfort by disavowing the appropriateness of film as a medium for ethnographic outputs, artists, audiences, and many anthropologists have recognized the distinctive advantages that data in film form can afford. As Schäuble observes, "visual ethnographic practices comprise reflexive, interactive, participatory techniques that provide access to specific local and/or embodied forms of knowledge and perception that would otherwise be difficult to access or that [would] remain invisible" (2018, 2). As will become clear in our discussion of some of the mixed methods that we have developed for working ethnographically with first-person data—the raw files of which overwhelmingly anchored in digital film recordings—the access Schäuble suggests traditional video recording techniques allow are further intensified by always-on wearable technology. And yet at the same time, as we will later outline, the perceived objectivity of the outputs are a risk (and illusion) with first-person data as with photographs.

From the vantage of our current ethnographic work with first-person data, the film medium continues to invite innovations around reflexive, interactive, and participatory methods. But even before wearable tech and ubiquitous cameras on personal devices, ethnographers imagined and were playing with the shifts in perspective that film as a medium for ethnographic data production could conceivably afford. In some instances, they were rendering visible not the ethnographer's observations of their interlocutor's worlds but the traditional ethnographic subject's first-person field of view instead. We note the "Navajo Film

Themselves" projects that Sol Worth, John Adair, and Richard Chalfen contrived in 1973, handing the cameras over to the former subjects of an ethnographic gaze, to see how they saw themselves and the world—and what they captured of it from their first-person points of view. And indeed, by empowering people to make ethnographic media that depicts their experiences from a direct and personal, engagement is likelier to extend *beyond* the ethnographic encounter, since media can be consumed or re-experienced alongside audiences or in contexts that continue to be meaningful for the former subject's everyday lives. This kind of first-person data production and continued engagement is observable in indigenous media like Maori TV and National Indigenous TV in Australia (see Ginsburg 1994), indigenous radio (see Fisher 2016), and Nanavut Independent TV.

Anthropologist and experimental film-maker Lucien Castaign Taylor, who beginning in the 1990s founded the Sensory Ethnography Lab at Harvard University, set about using film to render multisensory experience in sight and sound. As Castaign Taylor puts it: if "anthropology is to create a space for the visual" it "would entail a shift from the attempt to convey 'anthropological knowledge' on film…to the idea that ethnography can itself be conducted filmically" (1996, 86). Taylor poses the questions: "What if film not only constitutes discourse about the world but also (re)presents experience of it? What if film does not say but show? What if film does not just describe, but depict? What then, if it offers not only 'thin descriptions' but also 'thick depictions'?" (1996, 86). The Lab's filmic collaborations, from Sweetgrass (2009) to Leviathan with Verena Paravel (2012), model how a sensory medium can indeed "embody" ethnographic knowledge, blurring the boundary between subjects, technology, viewers, affect, and experience. The way we see it, first-person data can similarly enrich ethnographic research, providing a unique embodied viewpoint that traditional third-person observations cannot achieve.

From Ethnomusicology to Anthropology Composed of Sounds, Sensory Data Beyond the Visual Has Been a Powerful Medium for Tapping into and Capturing Emotions, Affect, and Memory

Lastly, there is sound recording—sensory ethnographic data beyond the visual—which less depicts than evokes lives and worlds, experiences and perceptions. The first-person data that we will describe in the next sections is just as much oriented towards audio as towards visual input from a person's daily life, so we provide a brief review of sound's role in ethnography over time here. As with cameras, phonography was experimented with in the field by early anthropologists as soon as the technology emerged. From Frank Hamilton Cushing in the nineteenth century, to Benjamin Ives

Gilman and Walter Fewke's recordings of Zuni Melodies (1891), to Alice C. Fletcher's and J.C. Fillmore's "Study of Omaha Music" (1893), American ethnologists sought to capture the sounds of people's languages, music, and song: impressions in wax, aural grooves of social life and human expression. Herzog's studies of plains ghost dance rituals and the sound worlds of the great basin were especially instructive in this regard (1935, 1936).

Folklore and phonography formed the basis for the emergence of ethnomusicology, which sought to develop archives of cultural production that in some ways (unintentionally) parallel and precede the development of datasets of various kinds used for training artificial intelligence today. Carl Stumpf in Germany founded the Berlin Phonograph Archive not long before Alan Lomax (sometimes joined by Zora Neale Hurston) set out in emerging radio days to record the sounds of ethnic and of folk peoples across America, producing an archive that became the basis for the National Archives and Library of Congress in D.C.

More recent anthropological studies using sound have taken a more dynamic and participatory approach to data collection – echoes of which are in the methodology we will outline below. For example, the work of anthropologist, ethnomusicologist, linguist, sound engineer, and composer Steve Feld shows us how sensory data, especially sound, offer unique possibilities for making sense of people and their worlds—and the role of audio in amassing first-person data sets of promise and possibility. In his 1982 monograph of about the sound-worlds and Kaluli language of the Bosavi people of Papua New Guinea, Feld captures vocalizations alongside ethnographic audio recordings of the rainforest worlds—the soundscapes—that compose their sense of being in and of their place among their worlds: waterfalls, birds, soulful lamentations, non-linguistic utterances (Feld 2012a). The kind of first-person data through sound that Feld invites the Bosavi to record in audio (and remix and edit and compose further) helps constitute what Feld calls "acoustemology," an orientation and a type of sensory data rife for "investigating the primacy of sound as a modality of knowing and being in the world (Feld 2012a, 2012b, 2015).

From their first forays into fieldwork to traditions in visual ethnography, sensory ethnography, and anthropology in and through sound, we see anthropologists experimenting with first-person data, exploring the avenues of access into human experience and perception its different varietals afford. We situate our work with first-person data as issuing from this tradition of anthropologists innovating with data forms and formats and experimenting with new media and new technologies as they emerge. In what follows, we outline the direction that we have taken in our own recent ethnographic work with first-person data and its current generation of devices for self-capture.

Seeing the World from Someone Else's Perspective: Overview of First-Person Data Collection and Analysis

The aim of this section is two-fold: to illustrate the ethnographer's experience of working with first-person data and to offer core tenets for engaging with it.

This paper draws on methodological insights from a series of studies conducted in the United States between 2018 and 2024 for Meta Reality Labs Research. Throughout these studies, which broadly sought to inform the development of next-generation wearable and personal assistive devices, our teams collected over 1000 hours of multimodal first-person point-of-view data using a range of technology to capture footage. Study participants represented a diverse range of adult ages, ethnicities, genders, occupations, and living arrangements. They were compensated for their contributions, which included capturing first-person data during both targeted exercises and normal daily life. They were made aware of the purposes of the studies, the identities of the organizations conducting the research, and could withdraw from the studies at any time. Our teams obtained consent from the research participants and other individuals in their social ecologies who participated (e.g., those in view during the recordings or whom we interviewed).

We have published on aspects of some of these studies in previous EPIC Proceedings (Cury and Whitworth et al. 2019; Cury and Kim 2021), and here outline two of the main studies:

Study: Internal States during Goal-Oriented Tasks

This study set out to map relationships between what is observable in context and what is subjectively experienced by an individual, in particular their experience of mental effort. Providing our research participants with a pair of eye-tracking wearable glasses used in academic and commercial research at the time, we studied and subsequently patterned the diverse individual experiences of goal-oriented domestic activities. We investigated how the same kinds of domestic tasks, like washing, chopping and sauteing vegetables, which may outwardly look the same across two individuals, can feel or be experienced totally differently across the individuals – one person might be happily in a groove as a they cook, while the other might be distractedly worrying about something else. What were some of the observable or machine-readable clues of the participants' subjective states and experiences of mental effort? How could recognition of how someone was experiencing their context inform the kinds of interventions or support they would find meaningful? To answer these questions, we conducted full-day ethnographies

with 18 people in the great New York and Seattle areas. Participants recorded their experience in twenty to thirty-minute intervals and were asked to describe their mental effort in-the-moment using experience-sampling techniques, scales and metrics common in psychology. The outcome was a framework informing how and when next-generation AR/AI assistants should surface information to users, and how a machine-learning model might detect these instances.

Study: Detecting and Encoding Personal History Information That People Find Most Meaningful

This study set out to help identify what kinds of sensors wearable devices would need to have to collect meaningful personal history information to help people better navigate everyday life. From the human perspective, what information from their pasts (whether that was yesterday or many weeks ago) do people need, resurface, and have questions about in-the-moment? From a technology perspective, what would future technologies need to sense in-the-moment to make sense of the situation and ultimately service those needs? Our role as ethnographers was to inform engineering decisions by observing people in their contexts. We employed traditional ethnographic methods like contextual inquiry, participant observation, and diary studies alongside first-person data capture using the Aria device, a non-commercially available wearable developed by Reality Labs for the purposes of research (Engel et al. 2023). Our 13 participants recorded with the Aria devices for 4-6 hours per day over approximately two weeks, indoors in their homes with the consent of others with whom they shared their home environment, with the option to self-narrate their experiences with their description of what they were doing, why they were doing it, and how they felt about what they were doing. Participants were also given specific 'missions' to capture data using Aria on areas of interest, such as memorable moments, significant objects, and daily challenges, tailored to the study's goals. This work resulted in a framework for the information most relevant to people throughout their everyday lives and a minimum viable list of sensory capabilities that a hardware device would need in order to sense that information. This framework informed how wearable-enabled AI systems sense and encode observable reality based on what information is relevant to people.

What First-Person Data Looks and Feels Like

Across our studies wearables captured a range of moments such as:

- Social interactions, e.g., super bowl parties, family dinners, heart-to-hearts, heated debates, family milestones, pillow talk, divorce announcements…

- Religious rituals, e.g., Havdalah service, Bible study, Muslim prayer…
- Leisure activities, e.g., dog training, sewing, painting nails, crafting, board game nights, Oscar watch parties, yoga sessions, baking, cooking, dance party, pancake night, writing projects…
- Household activities, e.g., dishes, childcare, gardening, meal preparation…
- Professional activities, e.g., remote work, finalizing book drafts, composing lectures, networking calls, planning a workshop, discussing career changes…
- Personal hardships, e.g., crying, arguments with family, saying goodbye to a foster dog, insomnia nights, dealing with the loss of a parent…

Here are video stills of first-person data collected from a wearable device to exemplify the kinds of footage our research participants recorded:

Figure 2. Two still-images of first-person video data captured using a wearable device during one of our studies; both are from the embodied vantage point of the person wearing the camera, showing what they are seeing. The left image shows their hands as they handwash clothes at the kitchen sink. The right image shows one hand holding a vase they are filling with water and one hand holding the plant that will go in the vase. Not shown is how in these recordings the individual was also narrating actions, thoughts, and motivations.

Here is an excerpt of one of our ethnographic fieldnotes as we review the first-person video data collected by research participant Ryan in one of our studies – note that this is written from the point of view of the researcher, not the participant who recorded the footage, but includes quotes of some of the audio of what was said in the recorded footage:

> It's day 3 of the study and I'm watching Ryan playing with his 6-month-old baby Leo. The father and son are sitting on the living room carpet with an assortment of colored ducks. Leo puts the ducks in his mouth, then drops them. Ryan picks them up and gives them back. He seems to be focused on Leo, who is examining the red duck. Although I can't see Ryan face, his tone of voice makes me think he is smiling.
>
> Although I've interviewed Ryan in-person, it is strange to watch his pov footage. It's as if I'm sitting on his shoulder. Ryan and his wife Amy are new parents and just moved into their first home. Their days are packed. They take turns working and caring for their son in shifts. They balance his feedings, diaper changings and naps with host of household chores: cooking their own meals, clean up, dishes, laundry.
>
> Ryan exhales sharply as he gets up. Ryan had mentioned his accident and his nerve issue during the interview—but I had not realized it was this bad until I began to watch his footage. Now, I see how unrelenting and the agony he endures day-to-day as he writes emails, cooks dinner, even plays with his son. It feels intimate to watch him in physical pain—in the moments when he is alone.
>
> I now know his 'tells' –signs that he indicate he is experiencing some kind of smarting, tightness or numbness. I know the sounds: his sharp exhales, whimpers, grunts. I know what he does—the stretching, massaging, shaking, or knocking—to find some relief. I know his triggers too: when he bends over or puts weight on his arm. While I have learned these patterns, I imagine that a machine would pick up on many more.
>
> Later that night, Ryan stands in the doorway. I think he may be using the doorframe to stretch. He exhales. His pain seems acute. His wife, Amy, who is in the other room looks up. 'Babe, you ok?' She knows too. My arm is numb as shit.' he says.

Moments like this one in the fieldwork helped inform the kinds of personal information that people find relevant and have questions about in their day-to-day – in this case, what might be causing a recurring event and how that might be changing over time.

Core Tenets for an Ethnographic Approach to First-Person Data

To provide methodological guidance for ethnographers engaging with first-person data, we offer three tenets that have proven to be critical to our own studies' data collection and analyses. The first tenet of engaging with first-person data is to contextualize it using additional ethnographic methods to gain a richer, more

rigorous understanding. The second is to encourage participant self-reflection at different timescales to capture the evolving or layered nature of experience. The third tenet is to co-interpret first-person data alongside participants to ensure analysis remains deeply rooted in participant perspectives.

Tenet 1: Collect and Triangulate First-Person Data with Traditional Ethnographic Methods

For the applied ethnographer, the value of first-person data emerges when it can be triangulated with other ethnographic methods during collection and analysis. Across our studies, we collected first-person data alongside several other 'kinds' of ethnographic data to contextualize participant's experience that they had captured with wearable technology.

Here is an excerpt of our ethnographic fieldnotes as we review the first-person video data collected by research participant Amy in one of our studies. Moments like this one brought together a range of ethnographic data alongside the first-person recordings, and informed our understanding of the motivations behind the kind of personal history data people wish to collect, and the implications this might have for how people might wish to access and engage with it:

> Amy is standing at the kitchen counter cutting an orange. Her husband Ryan (also a participant) is in the living room holding her son, Leo. As I review the couple's respective footage, I watch them navigate jam-packed days taking turns working full-time jobs and caring for Leo. They also balance his feedings, diaper changings and naps with host of household chores: cooking, dishes, laundry.
>
> Amy holds out a piece to her husband. "Should we give him a taste?" she asks. As Leo begins sucking on the orange piece, Amy exclaims gleefully: "This is Leo's first experience with a citrus fruit."
>
> The next day, in the digital diary, Amy wrote: "Leo explored so many different foods today! He loved the banana pancakes Ryan made for him. But did NOT like avocado. Nope. And he drank water (first time) from his baby cup—more Leo Firsts. His baby tooth just popped through his gums! YAY!"
>
> Amy, like many parents, clearly values her son's milestones. From my initial interview with Amy, I know these recording milestones have a deeper significance. Amy is the daughter of refugees. Her parents fled Vietnam before she was born. The trauma that plagued her family's past meant that

her parents and elder siblings rarely spoke about 'the old days'. Amy has no stories or photos from her own infancy and very few from her childhood.

Amy takes me on a tour of her world--her house and her phone. She shows me the family group chat as well as her treasures: a photo of her parents on the altar, who both passed away last year. They never met their grandson. Amy tells me: "I don't know what my first word was. Or what foods I liked or didn't like as an infant. So I guess I want to record these for him. It's his story, so I want him to have it."

There are many ways to combine first-person data collection with traditional ethnographic methods. Here is an example of the fieldwork structure closely mirroring one of our studies, which interspersed first-person data collection with other ethnographic approaches to data collection and interim analyses:

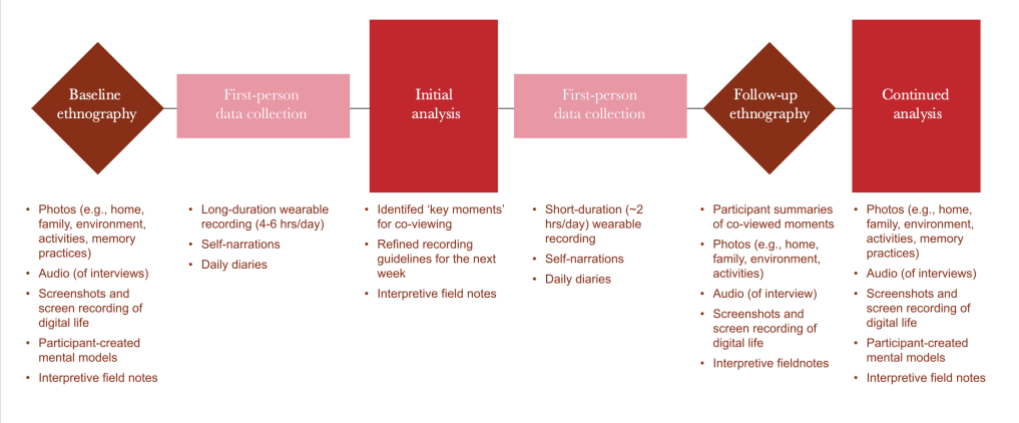

Figure 3. A schematic of the sequence of fieldwork activities in one of our studies – from baseline ethnography to first-person data collection, initial analysis, more first-person data collection, follow-up ethnography, and continued analysis. Under each of these sequenced activities we include bullet points for the kinds of data captured (e.g., in baseline ethnography we list photos, audio, screenshots of digital life, participant-drawn mental models, and fieldnotes). These bullet points closely follow what is already in the description of the methods and activities in this section of the paper.

Here are example ethnographic methods we triangulated with first-person data:

Foundational interview

Across studies, we conducted an initial ethnographic interview which provided baseline insights into the participants' lives, including their social circles, daily routines, home environment, profession, and values. A foundational interview helped us make sense of and interpret the first-person data that participants subsequently captured with the wearable research devices we equipped them with.

Rapport-Building Embedded with Tech Support

As with any traditional ethnographic engagement, building rapport with participants is critical. Atypical to traditional ethnographic practice, yet necessary to ethnography predicated on the collection of first-person data, was the need to onboard participants to the wearable research devices they'd be donning for the bulk of the study as part of the rapport-building. This also included helping to 'socialize' the devices in the context of their households and family life and providing ad hoc technical support to participants. Based on our own piloting of the wearables, we drafted user manuals for the devices and troubleshooting guides addressing the kinds of challenges one might encounter using these devices in home environments rather than in controlled lab settings, such as internet bandwidth issues or where to safekeep the devices while charging. In the initial onboarding, we described our own experiences with the technology (having tested it ourselves), acknowledging how different it might feel, at least at the outset, to wear the equipment while going about one's day-to-day activities.

Digital & Physical Tours and Artifacts

To ensure that the ethnographer would have enough background to situate participant-collected first-person data in the context of participants' everyday life, we gathered extensive contextual data, largely in the form of personal artifacts, both physical and digital. This included detailed tours of their homes, favorite belongings, and physical environs, with the participant describing origin stories and layers of meaning in an archeology of their possessions and living spaces—all documented photographically. We toured their digital environs: the spaces they navigate and inhabit online, their self-representation, the information ecologies and info-storage practices that constitute the digital layer of their lives and which shaped the everyday experiences of each participant. We compared the relationship between their physical surroundings and their digital surroundings. We asked participants to only share with us what they felt comfortable sharing.

Fieldnotes

The 'fieldnotes' composed in the wake of the initial ethnographic encounter with each participant—intended to forge a baseline context for the participant's lived experience—were by necessity multi-media and multi-modal, comprising photographic images, screen shots, video clips, lists, quotes, transcripts, drawings, and the ethnographer's own interpretive text. Fieldnotes drawn from the aggregation

of these sources resembled 'data packets' more than the traditional ethnographer's scribbles—informational archives that could be analyzed many ways.

Tenet 2: Encourage participant self-reflection at different timescales

Participant self-reflection is critical to attuning the first-person data to what matters most in the scene, from the point of view of the participant, otherwise it can be difficult to assess what is happening (beyond the obvious) and what might be most meaningful.

Here is an excerpt of our ethnographic fieldnotes as we review the first-person video data collected by research participant Mark in one of our studies – the self-narration from Day 2 and Day 3 helps the researcher understand the meaning behind what was happening in the scene in the recording from Day 1:

> Day 1
>
> Mark is an avid photographer and I watch as he reviews past photo albums.
>
> Day 2
>
> (morning)
>
> I watch Mark down sit down in his bedroom armchair with his laptop and iPhone. He narrates: "I'll be starting to draft that letter to my ex. Here will be the last communications that we'll ever have. The phone call I got from my friend this morning—he said to wait. My parents gave me different advice. At least I can get some of my thoughts down. This all stems from looking at those pictures yesterday."
>
> (evening)
>
> I read Mark's end-of-day digital diary. He writes: "This day was mostly at home…let me list some notable accomplishments from the day:
>
> I drafted a deep and heartfelt letter to my soon-to-be-ex-wife
>
> I went grocery shopping twice
>
> I cooked an enchilada casserole
>
> I played a game of chess with my daughter

Two forms of participant-generated self-reporting at different time scales anchored our analysis of first-person data: in-the-moment 'self-narration' of what they were doing and why, and what they were thinking and feeling during their recordings, and daily end-of-day diary entries for reflections.

Self-Narration

'Self-narration' served two primary methodological purposes: first, to add a contextual layer to first-person data, while also offering an emotional and affective layer to the first-person data as it arose in situ for participants; second, to provide the ethnographer with a contextual foothold into the first-person recordings they'd be analyzing – researchers could easily scan through the audio and transcripts to find key moments to analyze in more detail.

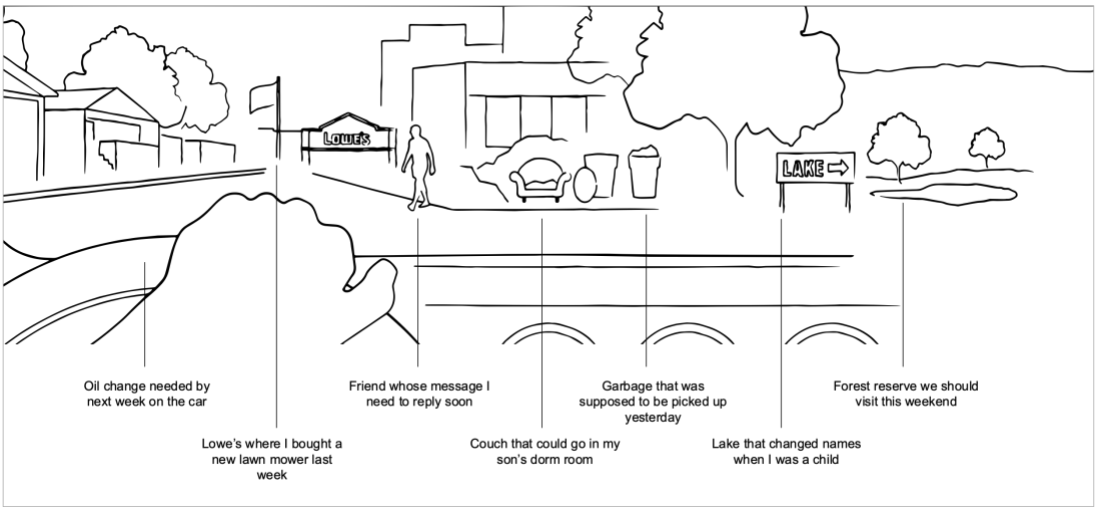

Figure 4. A line-sketch developed by ReD Associates, illustrating a scene someone would see from the windshield while driving: their hand at the wheel, and in front of them tree-lined streets, a hardware store, a pedestrian, a couch left out by the garbage for donation, and a sign pointing to a nearby lake and a forest. Lines from each of these elements lead to a short inlayed text description of what might be meaningful about these elements for the person whose point-of-view is embodied in this sketch. E.g., the steering wheel might make them think about the oil change they need on the car next week, the pedestrian crossing the street might remind them of the friend they need to message, and the couch might be an item they could take for their son. This sketch is meant to show the kinds of meaning and relevance that self-narration can reveal as part of first-person data collection.

Diaries

The daily diaries, a reflection on what was and was not captured in the first-person footage that day as well as the participant's point of view on what was memorable for or notable to them, directed the researcher to 'moments' of first-person data recording that were potentially rich with personal meaning for the participant, and which warranted in-depth analysis.

Tenet 3: Co-interpret first-person data alongside the participant

First-person data requires a filter to digest large amounts of 'raw' data (video footage, audio recordings) — to add weight to specific events or periods of time and to draw connections. While the self-reflections do part of the work in this, also

revisiting the recorded first-person data with participants during analysis offers an opportunity for the ethnographer to go beyond what was done, said, and reflected on in-the-moment, and for the research participant to exert more agency and control over how their footage is being interpreted, thereby giving them more of a voice in the outcomes of the study.

Here is an excerpt of our ethnographic fieldnotes as we review the first-person video data collected by research participant Aisha in one of our studies – the co-interpretation with Aisha was critical for understanding what happened in the moment, which could have otherwise gone misinterpreted:

> I meet with Aisha for a follow up interview and to review some of the first-person footage she recorded. I tell her that I have pulled out some 'specific moments' from the past week—and that we are going to watch them together. I say that I want to hear about her experience in that moment. Aisha pipes up, "Oh, are we going to watch the fight?"
>
> "The fight?" I don't remember watching a 'fight'.
>
> "Yeah, the fight. You know, the argument in the kitchen." I did not see an argument in the kitchen. However, one of the clips I pulled was a moment where Aisha's mother seems to be scolding her in Fula. Aisha is quiet, washing dishes as her stepmother employs a stern tone. Half of what I witnessed was in Fula.
>
> While I thought she was being reprimanded, to Aisha, this was a dispute.
>
> We rewatch the scene together and she explains the ins-and-out of their dynamic—its cultural dimensions and their relationship.

Co-Summarization

Across our studies, we summarized relevant moments from the first-person data to begin to draw out patterns and insights. We took notes about what the participants were doing, experiencing, feeling, and how that might matter to the participant in the context of the moment they had captured. In addition, we drew out what meaning it might hold within the broader context of the participant's lived experience beyond the frame of the first-person data captured in that moment, and what learnings and implications the moment had for our focus areas of study. To pressure-test ethnographers' initial accounts with participants, we conducted focused ethnographic follow-up sessions with each participant. A central activity for these revisits was co-viewing and co-summarizing selected 'moments' or recordings with the participant. We cued a clip or string of clips and asked: "what's going on?" and recorded what the participant said. We asked what we were missing, gathering more

context and artifacts to support the participant's summary. We shared our own summaries with them, and then co-summarized the clips upon reviewing once more, ensuring that together we reached the best possible summaries of first-person data moments. We outline more of this method in Cury and Whitworth et al. 2019: 262-3.

The three tenets outlined above are not meant to be exhaustive, and we imagine more will surface as first-person data continues to be used in applied ethnographic research.

Benefits of First-Person Data Capture

We contend that by viewing the world from the participants' perspectives, researchers can access aspects of their lives that are typically hidden, gain longitudinal insights, rethink the observer-observed dichotomy, remain more attuned to the participants' contexts, and combine qualitative and quantitative data collection at scale. These benefits collectively enhance the depth and breadth of ethnographic research, making it more robust and reflective of real-world human experiences.

Benefit 1: Access to Aspects of Participants' Lives Unobservable through Traditional Methods

One of the primary benefits of first-person data capture is the ability to access aspects of participants' lives that would have been unobservable through traditional ethnographic methods alone. The presence of a researcher often influences participant behavior, as they may alter their actions when being observed. While the researcher's skills in building rapport and triangulating observations address this inherent aspect of ethnographic fieldwork, the addition of first-person data capture with a wearable device while the researcher is not present provides more of what participants do alongside what they say and provides instances that might be difficult to be present for otherwise, like morning wake-up routines or moments of boredom. At first the participants from our studies found the wearable technology novel and the equipment captured their attention. However, we spent considerable time ensuring tailored fitting of the devices for comfort, providing clear and careful instruction in relatable terms and in different formats (e.g., videos, written guidance, live demonstration and testing), demystifying how the devices work by explaining the technical details that might not be readily apparent (e.g., the extent to which the participant could view and interact with their recordings), and emphasizing that the participants were in control of deciding when to record and when not to record. After a few tries or over a few days, the wearable devices integrated into daily

routines, allowing participants to act more naturally over time. This normalization enabled the capture of spontaneous and genuine interactions, such as intimate family moments, conversations, and personal routines. For instance, one participant, Linda, talked about her interest in crafts and making furniture when we first met her. In her first-person recordings, we observed there was more to this: Linda coped with insomnia by engaging in crafting and furniture-making throughout the night. In her in-the-moment self-reflection, Linda described how she developed this practice by noting patterns in which activities made her sleepy. This, alongside other examples of participants describing how they discovered hidden patterns in how seemingly unrelated activities shaped their wellbeing, helped form one of the key findings in one of our projects, related to the kinds of information people find meaningful.

Benefit 2: Gaining Longitudinal Insights by Seeing Lives Unfold in Real Time

First-person data capture allows researchers to gain longitudinal insights by observing how participants' lives unfold in real time, over time. Unlike traditional methods that often rely on sporadic and isolated data points, first-person data provides a continuous stream of contextual information. This continuous observation helps researchers understand the temporal dynamics of participants' lives, such as what ruptures and enables consistent daily routines, how behaviors become habits, and how experience changes over time. By recording daily activities repeating across several instances over time, researchers can identify patterns and shifts in behavior, offering a more comprehensive view of how people navigate their environments, build their perceptions, and interact with others. For example, one participant Raquel initially described playing backgammon on her phone as a minor and unimportant part of her day-to-day. We could not have easily observed her playing backgammon because she played in interstitial moments of boredom. In her first-person recordings that she shared with us, we observed that she played the game much more frequently than she had initially described. When we revisited Raquel for a follow-up interview, this recurring footage over time allowed us to talk more openly about her habits and routines, and she described the ways in which she felt stuck, lonely, and uncertain about her future. This kind of longitudinal data can reveal what people are unaware of as they build habits slowly over time and enables us to explore the ways actual experiences might not match needs or aspirations. We can also more closely examine how specific actions might contribute to general feelings, e.g., of belonging or loneliness.

Benefit 3: More Participatory Research

First-person data capture challenges the traditional observer-observed dichotomy by creating a more immersive and participatory form of observation. The wearable devices act as an extension of the researcher, allowing them to 'be present' without physically intruding on the participants' space. When combined with methods of co-interpretation of the footage alongside the research participant, as described above in the core tenets, this transforms observation into a shared, co-constructed experience, where the boundary between the observer and the observed becomes fluid. (We note that first-person footage alone, without some element of co-interpretation, would likely not provide this.) This hybrid model fosters a deeper connection between researchers and participants, allowing for greater rapport-building and the ability to ask deeper follow-up questions (as with the example immediately above with Raquel). And it creates space for the participant to interpret and articulate their experiences in applied studies that may ultimately shape products or services they might one day use. As we describe in a prior paper, "[p]roviding research participants more opportunities to articulate their internal states, including what they need and what they don't need, rather than assuming or inferring from observations alone, seems particularly important for determining the relevance, helpfulness, and boundaries of an assistive technology in everyday contexts" (Cury and Whitworth, et al, 2019:264). In a study in which we asked participants to record first-person footage focused on their hands to inform future haptic technology development, we asked participants to replay the footage with us and describe, moment-by-moment, why they made certain actions and how they learned to make those actions; this allowed participants to surface aspects of skill-acquisition and the value they placed on their hands, in ways that they might not have articulated before (Cury and Kim 2021).

Benefit 4: Centering and Attuning Researchers to Participants' Contexts

By centering the data collection on the participants' point of view, first-person data capture makes researchers more attuned to the participants' contexts. This approach mitigates the researcher bias to psychologize participants by focusing on their subjective experiences (e.g., stated preference, beliefs, and judgments). With wearable data, researchers can observe how participants navigate their daily lives, interact with their surroundings, and behave in different contexts. This contextual sensitivity allows researchers to develop a holistic understanding of the participants, grounded in their lived realities rather than abstract self-perceptions alone. In one prior study, analyzing the contextual elements (e.g., objects, environment, social

dynamics) in the first-person data allowed us to meaningfully group aspects of context that seemed to correlate with participants' needs, moods, and what kind of task assistance they might be open to versus might reject (Cury and Whitworth et al. 2019; Zax and Cury 2020).

Benefit 5: Combining Qualitative and Quantitative Data at Scale

First-person data capture provides a novel way for researchers to collect qualitative data at scale while also incorporating quantitative elements. Most basically, a single researcher can have several participants recording first-person data at once, so that considerably more data is collected than if the footage came primarily from the one researcher. Moreover, the continuous and comprehensive nature of first-person data allows for the collection of large volumes of rich, detailed information per participant. Certain wearable devices can capture quantitative metrics, such as movement patterns, eye direction, and environmental data, which can be integrated with qualitative insights, and which will likely become more easily accessible and analyzable as AI assistance via software improves. This combination of data types enables a more robust analysis, bridging the gap between qualitative and quantitative research methods and enhancing the overall rigor and depth of ethnographic studies. Particularly in situations in which volume of findings or translating of insights into 'hard' and comparable metrics (e.g., correlations to elements in the environment) matter for the validity of the insights or for teams' abilities to work with the insights, first-person data can lend more credibility, authority, and relevance to ethnographic data. For instance, we hypothesize that if first-person data captured exactly how many times a person experienced pain in their shoulder over a five-hour period, that could serve as part of compelling evidence in a healthcare study pertaining to what contributes to quality of life or where there are gaps between physician and patient understanding.

Summary of Benefits

The integration of first-person data capture into ethnographic research offers numerous benefits that enhance the depth and breadth of understanding human experiences. By providing access to previously unobservable aspects of participants' lives, enabling longitudinal insights, rethinking the observer-observed dichotomy, centering researchers on participants' contexts, and combining qualitative and quantitative data at scale, first-person data transforms the practice of ethnography. Moreover, through the specific examples provided for each benefit above, we also show how these advancements position applied ethnographers to contribute more

significantly to the development of emerging technologies, ensuring that these innovations are deeply rooted in real-world human experiences and ultimately leading to more empathetic and socially meaningful advancements. We hypothesize that the benefits are applicable not only to studies with a technology-oriented outcome, but also studies related to health and wellness, behavior change, group collaboration, to name a few.

Risks of First-Person Data Capture

While first-person data capture offers numerous benefits for applied ethnographic research, it also presents several challenges and ethical considerations that must be addressed. These include the potential for perceived objectivity and reification of participants' lives, the risk of defaulting to an individualistic view of the world, and privacy concerns inherent in capturing intimate, first-person data – each of which we will explore here, but which do not form an exhaustive list. As this form of data collection becomes more widely used, we anticipate more risks will arise.

We argue that understanding the historical and contextual situatedness of first-person data is crucial to integrate this method ethically and effectively into applied ethnographic practices. Concretely, we suggest that embedding reflexivity in the analysis process, encouraging participant-driven self-narration, providing clear protocols for data collection, and maintaining transparency about data use can help applied researchers can mitigate these risks and uphold the integrity of their work. Some of these concrete ways to mitigate the risks of first-person data capture are also part of the core tenets of the method, which we outlined above.

Risk 1: Potential for Perceived Objectivity and Reification of Participants' Lives

One of the primary risks associated with first-person data capture is the potential for perceived objectivity and reification of participants' lives. The immersive nature of first-person footage can create an illusion of objectivity, where the recorded data is seen as a complete and unbiased representation of reality. In that way, first-person data capture shares characteristics with adjacent methods such as video- and photo-ethnography, which also involve visual documentation of participants' lives. As with video- and photo-ethnography, first-person data requires a critical awareness of the power dynamics and ethical considerations involved in representing others' experiences, and of the wider context of the data collection beyond what is in view. By situating first-person data within this broader methodological context, researchers

can draw on (and hopefully also expand) established ethical frameworks and best practices to guide their work.

We had several ways of addressing this risk. We leaned into participant-driven self-narration (as described above) which can bring ethical integrity to the research by allowing participants to control the narrative and highlight aspects of their lives that they deem important and wish to foreground or background. Additionally, we played back key recordings of participants' lives back to them, bringing them in as co-researchers of their own experiences.

Overall, in both working with the data and presenting the data to other teams, researchers must remain critically aware of their interpretive role and avoid presenting first-person data as an unmediated reflection of reality. Instead, they should emphasize the co-constructed nature of the data and the ongoing dialogue between the researcher and participants.

Risk 2: Defaulting to an Individualistic View of the World

First-person data capture inherently focuses on the individual's perspective, which can lead to an overly individualistic view of the world. This focus risks neglecting the broader social and cultural contexts that shape participants' experiences.

To address this, we integrated first-person data with other ethnographic methods that capture collective and communal aspects of life, such as participant observation and contextual inquiry about participants' pasts, their future aspirations, and aspects of social life not immediately in view. Also, while our participants recorded in private indoor settings they often included consenting individuals with whom their shared their daily life and physical spaces, e.g., roommates and spouses, and we sought to include these and other individuals, e.g., friends, extended family, in the participant observation and contextual inquiry portions of the fieldwork. In one study, we also mitigated some of this risk by recruiting 'dyads,' meaning we had two people collect first-person data within a single social ecology so that we could compare their points of view and better observe social dynamics.

With mitigations like these in place, we argue that researchers can balance the individual point of view with a more comprehensive understanding of the social and contextual dynamics at play.

Risk 3: Privacy and Consent Concerns

The intimate nature of first-person data raises ethical considerations particularly regarding privacy and consent. To navigate these challenges, we provided guidelines

on what is permissible to record, to ensure that the data does not inadvertently include sensitive information like addresses, financial documents, health records, or nudity. We established clear protocols guiding research participants on how to obtain the consent of any individuals they might have recorded. We showed participants how they could review and delete any recordings that included sensitive information, or simply any moments they did not wish to share in retrospect. Participants could not record outdoors or in public spaces – we focused on indoor, private spaces, with limited social interactions. As mentioned above, we maintained transparency about the uses of the data, the identities of the organizations conducting the studies, and the goals of the research, to foster trust and accountability between researchers and participants. This meant going over consent forms in detail, multiple times, and in multiple formats (i.e., individual time to read, followed by a verbal walk-through of the content with time for questions), and also going beyond the consent forms to discuss what kinds of moments each participant anticipated they might encounter that they would not be able to record or would not want to record, beyond the basic guidelines, and providing them with tailored guidelines and strategies for these. It also meant continued checking-in with the participants throughout the recording period to ensure they understood, could comply with, and felt comfortable with, the protocols for data collection.

A New Role for the Applied Ethnographer: First-Person Data Allows Access to New Aspects of Human Experience, but Embracing It Means Ethnographic Practices and Roles Will Require Change

With access to new aspects of human experience as detailed above, first-person data can transform ethnographic practice and has the potential to increase the impact of applied ethnographers. We believe first-person data can impact the field of anthropology because it is a good tool for understanding people, particularly when combined with other traditional ethnographic methods. We also believe that the integration of first-person data into anthropology might shift the discipline's influence among the social sciences. And we believe first-person data can change the role of the applied ethnographer in engineering. But this new role for the applied ethnographer requires new skillsets.

Expanding Ethnographic Tool Kits to Get at Previously Inaccessible Dimensions of Human Experience

The emergence of wearable devices for capturing first-person data is watershed, at least from the point of view of anthropology's modern mandate to examine what being human is and means today and tomorrow. First-person data offers a different kind of access to (and angle into) human experience. If it remains, generally, that the organizing mission of anthropology is to get as rich, varied, and deep a glimpse into human experience and the conditions that structure its possibility today, then first-person data is a novel platform for that mission's continuance—a lens and a depth of view that wasn't possible before. And for anthropologists who are interested in the figure of the human taking shape in relationship with powerful net-new technologies like AI, first-person data offers an especially valuable vantage because it works with the same kinds of inputs that AI-enabled wearables in the future likely will have: sensors capturing and recording context from the user's point of view to serve up personal assistance in-the-moment or later on. This parallel allows these future inputs and their value in human-device interactions to be simulated and studied anthropologically with today's wearables. We believe that with the availability of this exceptional lens into human experience, ethnographic attention to and experimentation with these new tools and data will be a necessity, not an option, for ethnography's continued relevance – just as the integration of video and photography proved crucial to the discipline.

Reconfiguring Anthropology's Role Relative to the Social and Behavioral Sciences

If ethnographers do embrace first-person data, leveraging it as equipment for diving deeper and differently into contemporary human experience than previously possible, we believe this could drastically shift present power dynamics among the human, social, and behavioral sciences. Relative to the social and behavioral sciences, anthropology (both academic and applied) anchored in ethnographic research has remained a stubbornly empirical science. This has meant mounds of data based on direct observation, on one hand, and tentative interpretation of it, on the other (especially since the reflexive turn of the 1980s). Theory is borrowed; methodological and conceptual innovation is favored. Trends and institutional norms across sociology, psychology, economics, however, have tended towards appropriating the methods and mien of the experimental sciences, supplanting qualitative observation with variants of quantitative and statistical methods, on one hand, and with broad-based theorizing on the other. Seriously integrating first-person data into

ethnographic practice, however, may dissolve the distinction between qualitative and quantitative data; at the least it blurs the distinction between human observation and interpretation and non-human/technical modes of observation and interpretation. The type of technologies, computational abilities, and data sets that are emerging will enable practitioners of this new ethnographic approach to pose and answer questions that no discipline-specific researcher or research program could have previously: for example, accounting for behavioral change individually or socially and at scale, less restricted by the effects of being observed.

Elevating Applied Ethnographers' Role Relative to Data Science and Engineering

The implications for applied ethnographers of working with first-person data extend beyond impact within the social and behavioral sciences. We hypothesize that ethnographers taking up this kind of data will not just increasingly impact but also become integral to the (r)evolutions in data and computer sciences and engineering currently underway. While ethnographers using first-person data via wearable devices gain new kinds of social science insights, they at the same time gain an understanding of how people meaningfully use and interact with first-person data and wearable devices and become versed in how to translate sensor data into aspects of reality that are meaningful to people. This, in turn, could be helpful in shaping how wearables capturing first-person data (for purposes other than applied ethnographic research, e.g., commercial or for tech development) are built.

Whether working traditionally or experimentally, academically or applied, ethnographers make sense of the irreducibly different. And today, data scientists, technologists, and engineers in fields like computer vision and artificial intelligence would benefit from anthropologists and ethnographers as collaborators, fieldworkers, and experimental partners to engineer, train, apply, and design the systems (and sciences) they are dreaming into reality. Systems that are non-human, but which can both sense and make sense of human experience and culture can best be pursued with anthropologists mediating between machine intelligence and humans' experience, between experimental technology for capturing novel human datasets and the human beings participating in its production.

Furthermore, this reconfiguration between anthropology and engineering heralds a landmark repositioning of human sciences and ethnographic methods, for these would no longer be pursued or consulted downstream of engineering (e.g., in UX). Instead, ethnographers engaging first-person data can collaborate laterally with engineers, enabling while humanizing their ambitions. Collaborating with engineers

(rather than being downstream) means that ethnographers embedded in tech and engineering companies would not be just engaging with teams in product, marketing, and some strategy; working hand in glove with engineers to envision, develop, apply, refine, and humanize future technologies would allow anthropologists to reach new audiences of decisions-makers within technology companies.

This Requires Revamping the Applied Ethnographer's Skills, Learning New Tools, and Reorienting to Their Subjects of Study

Meaningfully integrating new technologies and kinds of data into ethnographic practice, however, means expanding the tool-kits ethnographers are used to. This includes:

- Working with new data formats and media (e.g., mixtures of quantitative and qualitative data, of observational and artifactual data, physical and digital data, and of raw and coded data)
- Learning about engineering infrastructure, practice, process, and language to collaborate with engineers
- Working with nascent wearable technology and AI-enabled software that helps to organize and process collected data

It also means that ethnographers must revamp their skills in ways that include reorienting to their subjects of study, the locus of the interventions, the venues in which they work, and the types of instruments and technologies they engage for doing so. This includes:

- Helping participants (i.e., device-wearers) take on a lot of the data collection and gathering capacities that an ethnographer would traditionally conduct
- Allowing computer vision and sensors to encode information from the environment and working with those formats and limitations
- Gathering through and working in dialogue with sensing technologies and AI systems to scale data collection, to amass and interpret large data sets
- Interfacing with multiple research participants at once, without sacrificing the rapport-building that matters and the interpretive responsibilities and sensitivities

While the ethnographer can extend observation and data collection through emerging technologies, we firmly believe ethnographers must remain wholly accountable for analysis and interpretation and to the experience of research participants who are not just participants but co-researchers as well. We argue that at its best, technology s never displacing human prowess (or judgement), analytical

acuity, empathy, or engagement, but rather augmenting what is possible and extending ethnographic capacity into depths and details previously inaccessible and could scale it to previously unimaginable reaches.

Embracing first-person-data means ethnographic practices and roles will require change – expanding tool kits and revamping skill sets, repositioning the role of the researcher in relation to the participants and the equipment, reconfiguring the role relative to the social sciences and to engineering, and extending through emerging technologies to create impact with new audiences.

Conclusion

First-person data capture can significantly impact applied ethnography by transforming and expanding its scope. Traditionally, ethnographers have relied on methods such as participant observation and interviews to gather data. While these methods define the field and provide valuable insights, they are often limited by the presence of the researcher and the sporadic nature of data collection. First-person data collection overcomes these limitations by enabling continuous and immersive data collection from the participant's perspective. This allows ethnographers to explore new dimensions of human experience, capturing nuanced, real-time interactions and behaviors that were previously beyond reach.

Redefining What Applied Ethnography Can Be

First-person data broadens the set of questions that applied ethnography can answer by providing researchers with access to meaningful data that would have otherwise been inaccessible. High-quality wearable technology and new means of extracting, processing, and analyzing hours of footage provide the ability to capture intimate, everyday moments for a deeper understanding of participants' lives. This access to rich, contextual data enables ethnographers to investigate complex social phenomena, such as group dynamics, personal routines, and interpersonal interactions, with a level of detail that was previously unattainable. By redefining the boundaries of what applied ethnography can explore, first-person data opens new avenues for research and discovery. We hypothesize that first-person data can be helpful for applied questions that require a detailed and longitudinal understanding of the day-to-day, such as, how can people better adhere to health goals and regimens? How might teams collaborate better? What gets in the way of sustainability efforts?

Expanding the Role of the Applied Ethnographer

First-person data also expands ethnography's role in the engineering of future technologies. Historically, applied ethnographers have been excluded from critical early conversations in hardware development and training dataset design, often relegated to post-hoc analysis rather than active participation in the design process. However, the insights gained from first-person data captured through wearable devices provide ethnographers with valuable information about the types of sensors and capabilities needed for future wearable devices to be socially meaningful in real contexts. This knowledge allows ethnographers to become critical nodes in the design of high-quality, specialized data collection circuits that are critical to technology development (e.g., for training machine learning algorithms on real world data, or testing the capabilities and value of sensors). By contributing to the foundational stages of technology development, ethnographers can ensure that devices are designed with a deep understanding of human behavior, needs, and social dynamics, ultimately creating more user-centered and effective technologies.

Broader Challenges and the Need for Critical Engagement

The integration of first-person data into ethnographic research and into engineering presents broader challenges that require ethnographers to critically engage with emerging technologies like wearables and AI. One of the primary challenges is the need for ethnographers to develop technical fluency and an understanding of the capabilities and limitations of these technologies. This knowledge is essential for effectively integrating ethnographic insights into the design process and for advocating for ethical and socially responsible technology development. Additionally, ethnographers must navigate the ethical considerations and privacy concerns associated with first-person data capture, ensuring that their research practices uphold the highest standards of integrity and respect for participants.

Call to Action

We hope this paper inspires applied ethnographers and social scientists more broadly, to continue to experiment with, and discuss, first-person data, and to find ways to shape the early development of wearable technologies that will continue to capture first-person data for a range of purposes. By transforming and expanding the role of ethnographers, providing access to previously inaccessible data, and positioning ethnographers as critical contributors to technology development, first-person data has the potential to reshape the field of applied ethnography. However,

these advancements also present new questions and challenges that necessitate a critical and informed engagement with emerging technologies. By navigating these challenges thoughtfully, ethnographers can harness the power of first-person data to deepen their understanding of human experiences and contribute to the creation of more empathetic and socially meaningful technologies.

About the Authors

Maria Cury is a partner at ReD Associates. She studies the role of new technology and media in daily life to advise on product development, visioning, and strategy. She has led projects on AR/VR, haptics, and autonomous vehicles, and has conducted foundational studies on topics such as the smart home, computing, the role of entertainment media and books in people's lives, and the human senses in tech. She has background in anthropology, the arts, and cultural institutions.
mcu@redassociates.com

Eryn Whitworth is a member of the Project Aria team at Meta's Reality Labs Research. Eryn leads UX/UI efforts across multiple surfaces and topics, including interaction design, hardware evaluation, and responsible innovation. As a practitioner, Eryn blends ethnographic and laboratory methods to create insights from qualitative and quantitative data triangulation.

Ariel Abonizio is an anthropologist, artist, and business strategist who specializes in advising global technology companies on product and corporate strategy from inception to implementation. At ReD Associates, Ariel informs product, engineering and business leaders on how to design for thorny human concepts like trust, misinformation, cultural representation, and intimacy.

Gabriel Coren is a senior consultant with ReD Associates. As an anthropologist of science interested in emerging relationships between living nature and technology, he has conducted ethnographic work with biologists and engineers around the world. He earned a PhD in cultural anthropology from the University of California, Berkeley.

Maya Potter has worked across the worlds of emerging technology and cultural fandoms for companies across industries. She holds a bachelor's degree in anthropology and a master's degree in World Heritage Studies. Prior to joining ReD, Maya worked with NGOs and development organizations in South Asia.

Tamara Moellenberg is a senior manager at ReD Associates, leading mixed methods projects to define the future of emerging technologies and services. She has a background in literary and cultural studies. Prior to joining ReD, she completed her doctorate on conceptions of childhood in West African literature as a Clarendon Scholar, focusing on themes of socio-cultural identity and agency.

Mikkel Krenchel is a partner at ReD Associates. For over a decade, Mikkel has helped build ReD's presence in North America, advising executives while building the firm and the practice of human science in business more broadly. He focuses on how ideas around disruptive technologies shape our lives, for better or worse. He also serves as an advisor and thought partner to leaders at some of the

world's largest social media companies, telecom providers and electronics manufacturers, as well as clients across the finance, energy, and industrials sectors. Mikkel holds a degree in Sociology and International Studies from Yale University and is a former Danish national team rower.
mkr@redassociates.com

Notes

Acknowledgements: The authors would like to thank Richard Newcombe, Mingfei Yan, Kevin Hannan, Lambert Mathias, Laird Malamed, Saransh Solanki, Wang-Chiew Tan, Zhaoyang Lv, Faraz Kamili, Maddy Bowen, and Yan Xu at Meta Reality Labs for their intellectual contributions and support in the work that shaped this paper. The authors would also like to thank Natasha Jessen-Petersen, Fred Guo, Tine Walravens, Matthew Janney, and Matt Kay at ReD Associates for their support in the development of this paper.

References Cited

Adair, John, and Worth, Sol. 1973. *Through Navajo Eyes: An Exploration in Film Communication and Ethnography*. Bloomington: Indiana University Press.

Asch, Tim. Dodoth Morning. Documentary Educational Resources, 1961. 20 minutes.

Asch, Tim. Arrow Game. Documentary Educational Resources, 1974. 10 minutes.

Asch, Tim, Neal, James, and Chagnon, Napolean. Yanomamö: A Multidisciplinary Study. Documentary Educational Resources, 1968. 45 minutes.

Asch, Tim, and Chagnon, Napolean. Climbing the Peach Palm. Documentary Educational Resources, 1974. 9 minutes.

Asch, Tim, and Chagnon, Napolean. A Man Called "Bee": Studying the Yanomamö. Documentary Educational Resources,1975. 40 minutes.

Boas, Franz. 1911. Handbook of American Indian Languages. *Bureau of American Ethnology Bulletin*, 40 (1), 1–1069.

Charbonneau, Étienne and Carey Doberstein. 2020. "An Empirical Assessment of the Intrusiveness and Reasonableness of Emerging Work Surveillance Technologies in the Public Sector." *Public Administration Review*, 80: 780-791. https://doi.org/10.1111/puar.13278

Cury, Maria, Eryn Whitworth, Sebastian Barfort, Séréna Bochereau, Jonathan Browder, Tanya R. Jonker, Kahyun Sophie Kim, et al. 2019. "Hybrid Methodology: Combining Ethnography, Cognitive Science, and Machine Learning to Inform the Development of Context-Aware Personal Computing and Assistive Technology." *Ethnographic Praxis in Industry Conference Proceedings*: 254-281. https://www.epicpeople.org/hybrid-methodology-ethnography-cognitive-science-machine-learning/

Cury, Maria and Kahyun Sophie Kim. 2021. "Hands Are People Too: Reflections on the Value of Hands (and How to Study Them)." *Ethnographic Praxis in Industry Conference Proceedings*: 342-343. https://www.epicpeople.org/hands-are-people-too/

Chagnon, Napolean. 1974. *Studying the Yanomamö*. New York: Holt, Reinhart, and Winston.

Clifford, James and Marcus, George (eds.). 1986. *Writing Culture: The Poetics and Politics of Ethnography*. Berkeley: University of California Press.

Crapanzano, Vincent. 1986. *Tuhami: Portrait of a Moroccan*. Chicago: University of Chicago Press.

Dwyer, Keith. 1982. *Moroccan Dialogues: Anthropology in Question*. Baltimore: Johns Hopkins University Press.

Ellis, Carolyn. 2004. *The Ethnographic I: A Methodological Novel about Autoethnography*. Walnut Creek: AltaMira Press.

Engel, Jakob, Kiran Somasundaram, Michael Goesele, Albert Sun, Alexander Gamino, Andrew Turner, Arjang Talattof, et al. 2023. "Project Aria: A New Tool for Egocentric Multi-Modal AI Research." arXiv: 2308.13561. https://doi.org/10.48550/arXiv.2308.13561

Fabian, Johannes. 1983. *Time and the Other: How Anthropology Makes its Object*. New York: Columbia University Press.

Feld, Steven. 2012a [1982]. *Sound and Sentiment: Birds, Weeping, Poetics and Song in Kaluli Expression*. Duke University Pres

Feld, S. 2012b. *Jazz cosmopolitanism in Accra: Five Musical Years in Ghana*. Duke University Press.

Feld, S. 2015. "Acoustemology". In *Keywords in Sound*. Novak and Sakakeeny, eds. Durham: Duke University Press (12-21).

Fisher, Daniel. 2016. *The Voice and Its Doubles: Media and Music in Northern Australia*. Durham: Duke University Press.

Fletcher, Alice, and Filmore, J.C. 1893. Study of Omaha Music. Archæological and Ethnological Papers, Peabody Museum, Harvard University, 1(5).

Ganry, Laurent, Nicolas Sigaux, Kyle S. Ettinger, Salam O. Salman, Rui P. Fernandes. 2019. "Modified GoPro Hero 6 and 7 for Intraoperative Surgical Recording-Transformation into a Surgeon-Perspective Professional Quality Recording System." *Journal of oral and maxillofacial surgery: official journal of the American Association of Oral and Maxillofacial Surgeons* 77 (8): 1703.e1-1703.e6. https://doi.org/10.1016/j.joms.2019.03.026

Gardner, Robert. Dead Birds. Film Study Center of Peabody Museum, Harvard University, 1964. 84 minutes.

Geertz, Clifford. 1973. *The Interpretation of Cultures: Selected Essays*. New York: Basic Books.

Gilman, B. I., & Fewkes, J. W. (1891). *Zuñi Melodies*. Houghton, Mifflin and Company

Ginsburg, Faye. 1994. "Embedded Aesthetics: Creating a Discursive Space for Indigenous Media." *Cultural Anthropology* 9 (3): 365-382.

Graves, Steven Nicholas, Deana Saleh Shenaq, Alexander J. Langerman, David H. Song. 2015. "Video Capture of Plastic Surgery Procedures Using the GoPro HERO 3+." *Plastic and Reconstructive Surgery. Global Open* 3 (2), 312. 6 March. https://doi.org/10.1097/GOX.0000000000000242

Harper, Douglas. 2003. "Scenes in the Anthropological Life of Tim Asch," in *Tim Asch and Ethnographic Film*. Ed. E.D. Lewis. Routledge.

Herzog, G. (1935). "Plains Ghost Dance and Great Basin Music." *American Anthropologist*, 37(3), 403-419.

Herzog, G. (1936). *Research in Primitive and Folk Music in the United States: A Survey* (No. 24). American Council of Learned Societies.

Huang, Xin-Yue, Zhe Shao, Nian-Nian Zhong, Yuan-Hao Wen, Tian-Fu Wu, Bing Liu, Si-Rui Ma, Lin Lin Bu. "Comparative analysis of Gopro and Digital Cameras in Head and Neck Flap Harvesting Surgery Video Documentation: An Innovative and Efficient Method for Surgical Education." *BMC Med Educ* 24, 531 (2024). https://doi.org/10.1186/s12909-024-05510-2

Hurston, Zora Neale. 1935. *Of Mules and Men*. New York: Perennial Library.

Hurston, Zora Neale. 1938. *Tell My Horse*. Philadelphia: J.B. Lippincott Co.

Leiris, Michel. 1930. "L'oeil de l'ethnographe (a propos de la Mission Dakar-Djibouti) [The Ethnographer's Eye (concerning the Dakar-Djibouti Mission)]." *Documents* 7: 404-14.

MacDougall, David. 2006. *The Corporeal Image: Film, Ethnography, and the Senses*. Princeton: Princeton University Press.

Malinowski, Bronislaw. 1922. *Argonauts of the Western Pacific*. London: Routledge and Sons.

Mead, Margaret, and Bateson, Gregory. 1942. *Balinese Character: A Photographic Analysis*. New York Academy of Sciences.

Nader, Laura. 1972. "Up the Anthropologist: Perspectives Gained from Studying Up." Washington D.C.: U.S. Department of Health, Education, and Welfare.

Ortner, Sheri. 1974. "Is Female to Male as Nature Is to Culture?" in *Feminist Studies* 1 (2): 5-31.

Rabinow, Paul. 1977. *Reflections on Fieldwork in Morocco*. Berkeley: University of California Press

Rane, Nitin, Saurabh Choudhary, and Jayesh Rane. "Leading-edge Wearable Technologies in Enhancing Personalized Safety on Construction Sites: A Review" (November 13, 2023). Available at SSRN: http://dx.doi.org/10.2139/ssrn.4641480

Schäuble, Michaela. 2018. "Introduction. Mining Imagination: Ethnographic Approaches Beyond the Written Word." *Anthrovision* 4.2 (2016). https://doi.org/10.4000/anthrovision.2407

Seminoff, Jeffrey A., T. Todd Jones, and Greg J. Marshall. 2006. "Underwater Behaviour of Green Turtles Monitored with Video-Time-Depth Recorders: What's Missing from Dive Profiles?" *Marine Ecology Progress Series* 322: 269-280. https://doi.org/10.3354/meps322269

Tangen, Steffen, Alexander Olsen, and Ellen Beate Hansen Sandseter. 2022. "A GoPro Look on How Children Aged 17–25 Months Assess and Manage Risk during Free Exploration in a Varied Natural Environment" *Education Sciences* 12, no. 5: 361. https://doi.org/10.3390/educsci12050361

Taylor, Lucien Castaign. 1996. "Iconophobia: How Anthropology Lost It at the Movies." *Transitions* (No. 69): 64-88.

Taylor, Lucien Castaign. Sweetgrass. Cinema Guild, 2009. 101 minutes.

Taylor, Lucien Castaign, and Paravel, Verena. Leviathan. Sensory Ethnography Lab and Arrête ton Cinéma, 2012. 87 minutes.

Tyler J, SW Choi, and M Tewari. Real-time, Personalized Medicine through Wearable Sensors and Dynamic Predictive Modeling: A New Paradigm for Clinical Medicine. *Current Opinion in Systems Biology* 2020 Apr (20):17-25. https://doi.org/10.1016/j.coisb.2020.07.001. Epub 2020 Jul 7. PMID: 32984661; PMCID: PMC7515448

Yoshida, Hanako and Linda B. Smith. 2008. "What's in View for Toddlers? Using a Head Camera to Study Visual Experience." *Infancy*, 13: 229-248. https://doi.org/10.1080/15250000802004437

Zax, David and Maria Cury. 2020. "The Moonshot Tech Challenge of Our Era is Context." *Red Associates Perspectives*. Accessed 4 August 2024. https://www.redassociates.com/all-perspectives/the-moonshot-tech-challenge-of-our-era-is-context

What's Going on with Strategic Research in Big Tech?

MORGAN RAMSEY-ELLIOT, *ReD Associates*
CHARLIE LOTTERMAN, *ReD Associates*
CAMERON WU, *ReD Associates*

For years big tech was one of the major employers of, and training grounds for, industry ethnographers focused on solving complex, longer-term strategic problems. But the wave of layoffs that hit white-collar workers beginning in 2022 prompted many in the research community to question their historical value to companies and their offerings to them in the future. Many participants in this discourse arrived at the same conclusion: Researchers need to do more strategic research—but what does that mean? And does the core of the issue lie with research, tech, strategy, or some combination of the three? After surveying shifts in the macro context that have upended the status of researchers within businesses, specifically large technology companies, we apply Roger Martin's strategy playbook to develop a strategy for ourselves. We argue that to meet the moment, researchers should grow the strategic muscle of their companies by shifting their focus from how users interact with products to how their businesses interact with a wider world. By repositioning our offering to companies in this way, we argue that researchers can apply ethnography's unique capabilities to the most pressing strategic questions facing businesses today.

Introduction

Over the last 20 years, many industry ethnographers have grown up and thrived in big tech companies. For the purposes of this paper, when we say "big tech" we mean Fortune 500 technology companies that have historically well employed and supported applied ethnographers—such as Intel, Microsoft, Meta, Amazon, and Alphabet[1]. Understanding how people use technology in context has proven essential for answering *what* product teams should build for people, as well as the *why* behind those offerings. And looking at broader contexts (e.g., among competing products, in different temporal or spatial situations, or in historical and future-facing contexts) has also helped applied ethnographers position their work as a key input for *strategic* product decisions. Good thing, too, since influencing such decisions is not only a core personal aspiration for many of us, it's historically also been a core dimension structuring our performance evaluations and career trajectories. The more senior one's level at many of these companies, the greater the expectation to drive broad strategy and inform "big bets."

Ramsey Elliot, Morgan, Charlie Lotterman, and Cameron Wu. "What's Going on with Strategic Research in Big Tech?" *EPIC Proceedings* (2024): 316–334. https://www.epicpeople.org/whats-going-on-with-strategic-research-big-tech. ISSN 1559-8918.

Over the past few years, however, something has changed, not only for applied ethnographers but across the broader research community. The halcyon days of large research teams dedicated to greenfield, generative research exploring foundational human phenomena have given way to a bleaker landscape littered with "open to work" announcements, the bruised remains of whole pathfinding teams, and reduced stakeholder appetites and financial budgets for exploratory research. From EPIC and LinkedIn to brief coffee chats at the office canteen, many within our community are asking questions like: *Why and what do we do now? How can I protect my job and advance in my career? How can I afford to take risks in my research that are exciting and potentially high reward, without exposing myself to too much uncertainty?* Many prominent voices from the user experience (UX) research community have aligned around some version of the same answer: Researchers of all stripes, including applied ethnographers, need to do better at connecting their insights and recommendations to business outcomes.

While such calls to action are valuable, we have noticed a gap in the discourse. Much of the conversation doesn't adequately account for the changed business landscape of big tech (like higher interest rates, the maturation of many product spaces, and flat growth curves), or internal shifts (like the rise of data science as a powerful tool for product decision making), and consequent implications for where and how ethnographic research can and should play a role. Moreover, the effective meaning and practice of strategy in big tech has changed in response to such shifts. And within the tech research community, while we see many calls to do more "strategic" research, there seems to be little consensus on what we mean when we talk about strategy. If we as applied ethnographers want to drive even more strategic decisions in this landscape, we may need to update our own ways of thinking about strategy and our role in it.

What might help ethnographers who practice (or are aspiring to practice) strategic research is a clearer *strategy* for our own work, to make our approach more possible, more impactful, and more valued. We will outline one possible strategy, informed by a literature review and interviews with research practitioners, people leaders, and cross-functional stakeholders from data science and product management at some of the biggest tech companies today. This paper will thus:

1. Map the current macroeconomic/sociocultural/political forces shaping big tech companies' corporate strategies today,
2. Situate the enactment of strategy in big tech companies, and

3. Propose some recommendations for ethnographers looking to drive strategic decisions in this landscape, using a powerful yet under-appreciated (at least, within tech) strategic framework.

Section 1: What Is Going on in the Market Landscape?

We see at least four major shifts over the last roughly ten years that are shaping the role of research within technology companies. The first is that many of the big tech companies have evolved from startups to incumbents, and the overall tech product landscape has become more mature. More recently, we have also been experiencing relatively high interest rates and persistently tense macroeconomic market conditions. As a result, companies are experiencing increased pressure from shareholders to demonstrate strong bottom line results. At the same time, data science has matured into a powerful insights function that has become quite efficient at answering many important product questions. This section explores how such shifts have impacted how these companies' approach strategy.

From Startups to Incumbents

Twenty years ago, many of today's big tech companies were effectively startups that lacked financial repercussions for taking risks. Their evolution into mature companies has triggered many downstream implications for strategy and research. As startups, the business's strategy involved coming up with a new product, market or business model idea. Once identified, the company would build a minimally viable version, then track whether enough people would use it to justify continued investment. In this nascent business landscape, the approach was a bit like throwing spaghetti at a wall to see what stuck, followed by rapid scaling if it did (Kushner 2011). The underlying logic of product strategy was one of innovation more so than monetization.

Through our interviews, we heard that this cultural emphasis on innovation persisted even as tech companies were becoming well established through the 2010s, a period that coincided with many of them building large research organizations. This combination created a relatively lush environment for many ethnographic researchers: A persistent cultural focus on innovation, supported by growing opportunities for research teams to do foundational, ethnographic work to inform those innovation bets. But as these companies have matured and become beholden to shareholders, there is even more at stake. Tiny changes to the number of rows of pixels on the most efficiently monetized surfaces can directly and nearly immediately

result in billion dollar revenue impacts (McGee 2016). As a result, for many core product offerings from mature tech companies today, the underlying logic of business strategy has become more about optimization than innovation.

At the same time, the overall tech product landscape has become more mature and competitive. The first 20 years of the 21st Century were a period of massive growth in the number of product offerings launched (Palandrani 2020). The number of venture capital firms tripled in the years following the 2008 financial crisis. In 2006, only 3,500 startups were funded with $31 billion USD, but by 2021 those numbers ballooned to over 30,000 companies funded with $669 billion USD, an increase over 2,000% (Teare 2022). While the first iPhone was released in 2007, by 2016 there were more smartphones in use than people on the planet (Richter 2023). Watching this and other lines shoot up at such a rate, the financial incentive structures rewarded big risk taking, and in this context, one of the key strategic goals was being first to market with new big ideas. Investors and technologists alike saw reasons to believe that their bets were likely to pay off.

However, as the landscape has become more mature and crowded, there have been fewer true white spaces. Even recent breakthrough technologies such as generative artificial intelligence (GenAI) and mixed reality headsets are sometimes greeted with skepticism. As a result, the stance and potential opportunity for strategic research is now less about building brand new technology to capture open opportunities, and more about defending against competitors.

High Interest Rates and Reduced Cash for Investments

Another major recent shift has been the increased cost of money (Sunderji 2022). Interest rates were at historic lows in the 2010s before creeping up a bit right before 2020, when they then dropped dramatically as COVID threatened to upend the economy. Historically low interest rates were a crucial enabler for major product investments that big tech companies were making. They made it much easier to invest in major new product groups, hiring massive cross-functional teams to support these investments. As interest rates have shot up since 2022, it has become more costly to spin up new offerings, which has in turn impacted the calculus of which product bets are worth making (Walker and Dowd 2024).

Financial Pressures of Shareholders

A separate (but related) shift has been the increased pressure that Wall Street is putting on big tech companies to demonstrate financial returns (Picchi 2022). For

years, investors were willing to invest in the compelling visions of tech companies, betting on profits to be realized down the road. In this context, a key product strategy question was "*What should we build that users really need and will create strong user growth for us, regardless of whether it actually makes us money*?"

The overall tenor of the investors, however, has changed markedly. As one Meta shareholder wrote, "*Such a franchise can print money and shareholders want that money returned to them, not frittered away on head count and blue-sky bets*" (Ip 2023). The new imperative for many technology companies is less about finding new product opportunities to grow user bases, and more about what will make the company money today (Forbes 2022). Businesses have responded, in turn, by looking for safer growth opportunities, for example, eking more out of core products rather than building up untested ones.

The Rise of Data Science

A recent report from Stripe Partners (2023) emphasized the rise in tech companies of high-volume, scalable A/B experiments and the powerful data science analyses that can translate them into directional insights on how to move key user metrics. Running thousands of these experiments at once is an efficient way to derive actionable learnings that can directly impact numbers that the business cares about. This observation is an incredibly important one for making sense of our discipline's current state and potential future; it was also echoed strongly among many of the product managers and data science leaders who we interviewed. Data science answers many of the *what* questions with authority and speed that researchers cannot match. Of course, our interviewees point out that there are limitations to what data science can and cannot answer through even the most advanced experiment designs. For example, only having access to the data of existing users (not potential ones), and not always having the right technical infrastructure to capture and analyze impactful data. But especially for mature products in mature technology companies, which prioritize efficiency and optimization, it is data science—not research—that increasingly powers design changes and corresponding financial returns.

Such macro shifts in the overall business landscape have had substantial impacts on the way big tech companies operate, how they evaluate strategic bets, and the broader cultural ethos that drives day-to-day decision making. They have also changed the role and opportunity of strategic ethnographic research. Back in the "golden days," deep ethnographic research on foundational human experiences (think "belonging" or "wayfinding") was actually a relatively efficient way to identify new opportunities to build innovative and exciting new technologies. However, with less white space in the market overall, increased focus on risk mitigation and

optimization rather than breakthrough innovation, increased financial pressures, and adjacent insights functions that are highly efficient at answering core product questions, strategic ethnographic research may need a new strategy.

Section 2: How Has Our Discipline Responded?

In the current labor market, many research practitioners have been asking hard questions about how we got here and where we might be going. Much of the discourse can be categorized into arguments that call out cross-functional stakeholders or leaders for not recognizing our discipline's inherent value (see, for example, Spool 2023), or arguments that look inward to diagnose what "went wrong" for us researchers (Antin 2023).

Among many of the voices in the latter camp, we see a common call to action: To prove our value to business, we need to focus more on "strategic" work. What strikes us, however, is that many of the arguments pushing for research to be more strategic lack a clear perspective on what we mean when we talk about strategy. This is an issue that also came through strongly in our interviews with researchers and cross-functional stakeholders: We found substantial variation in what people mean when they say strategic, and we believe that this may be making it harder for us to actually drive more strategic impact.

Some of the common slippages we see in the discourse are a conflation with "future-facing," which we believe is problematic because a lot of important strategic decisions are very focused on the here and now. Another is a conflation with "complex" or "broad in scope," which we believe is problematic because there are plenty of research projects that tackle complex issues but are actually quite tactical. Another is a conflation with "impactful," which we believe is problematic because impact often just means something happens because of your work—a new feature launched or sunsetted, or a product change moved key metrics. Another very common one we hear is "actionable," which we believe is problematic because there is plenty of valuable strategic input that is actually not actionable, and a lot of tactical input that is highly actionable. Such other adjectives floating around in the discourse may or may not have anything to do with whether the outputs of said research inform actually strategic decisions for the business.

This semantic confusion is problematic for a variety of reasons. For one, "strategy" shows up all over researcher performance guidelines and evaluations. So not having a shared understanding and vocabulary for what strategy means has real consequences for individual career trajectories. But more important is that without a

common understanding of what strategy means and how we should be more strategic within our organizations, we likely will collectively struggle to drive the kind of impact that we aspire to.

Confusion about what we mean when we talk about strategy is not unique to the technology world. A simple Google search will produce dozens of different definitions and interpretations of strategy. What strikes us, however, is that many in the tech world have relatively little familiarity with the long history of thinking, theorizing, and framework building related to strategy. Before we outline our working definition of strategy and its implications for ethnographic researchers, let us briefly situate it in a bit of historical context.

Section 3: What Is Strategy?

So what is strategy—and what is it not? Any aspirations for the future of strategic research depend on a clear take on what strategy means, how to develop it, and how to *do* it. As a formal discipline, strategy first emerged in the 1970s, when Bruce Henderson (1979) at the Boston Consulting Group championed efficiency as a goal for businesses, enabling them to outcompete rivals through pricing models and other techniques. By winning at efficiency (the thinking went), these businesses would grow in scale, compounding their advantages relative to their competitors, and accelerating a positive feedback loop.

But Michael Porter's (1980) theory of competitive advantage soon expanded the lens of what strategy entails. He began to seriously think about different approaches to obtaining competitive advantage beyond prices and efficiency plays, including differentiating from their competitors through unique value propositions and brand positions. By the 1990s and into the 2000s, much of the conversation about strategy in business circles had substantially expanded to include discussions about unmet customer needs and customer satisfaction. Such topics are now core content in the broader business world, from business school training programs to the Harvard Business Review.

Within the landscape of strategic frameworks, we believe a particularly useful one for applied ethnographers is one popularized by the strategist Roger Martin[2]. In his and collaborator Alan Lafley's book *Playing to Win*, strategy is described as an integrated set of decisions collectively executed to create sustainable competitive advantage over rivals (Lafley and Martin 2013). In human terms, a strategy is a bundle of bets about a business, a customer, and the landscape in which they interact. These are bets at the juncture of capital allocation, product features,

customer segmentation, go-to-market channels, timing, brand positioning, and organizational structure.

Their framework outlines an integrated cascade of five types of decisions that collectively yield a comprehensive strategy:

1. Set a *Winning Aspiration*, a clarifying purpose that all following decisions will be measured against.
2. Make a decision about *Where to Play*, an intentional commitment to your core offerings, in which selected market categories, which geographies, and for which audiences.
3. The third step, *How to Win*, is about how to execute in this space in order to realize your *Winning Aspiration*.
4. Now with a plan to win, the fourth step asks *What Capabilities Must You Have?,* what skills and tools must be in place to realize a winning strategy.
5. The last step asks *What Management Systems Are Required?* In other words, organizational structure is needed in order to manifest your plan?

All five pieces have to harmonize. The strategy is tuned into a comprehensive focus by looking up and back down the cascade again, iterating and fine tuning as decisions evolve. Crucially, Martin distinguishes this way of thinking about strategy from a "plan" or a goal (e.g., moving X metric Y percentage points). From our interviews with tech workers, we heard that a lot of product leaders conflate strategy with plans. According to Martin:

> "Crucially, [a] plan does not specify how all these initiatives will link together to accomplish a given outcome. The essence of strategy is to compel the thing you do not control to do what you need it to do. If you think about a company, what does it control? It controls how many people it hires, how many capital dollars it puts in place, where it sells its products—all of those things are in its control. What is completely outside its control? Customers. What strategy is, is your way of compelling those people you do not control, the customers, to feel that buying your product or service is the best thing for them to do." (Martin, Storgaard, and Lau 2023, p. 3)

Now, the framework proposed in *Playing to Win* has its limitations. For example, we believe there is an opportunity to expand this framework to more explicitly consider how to build strategies that win given the complexities of modern business environments. New factors worth considering include network effects, social acceptability of new technologies, collective and group needs—not just individual

user or customer needs. However, we believe this framework is one of the most helpful for applied ethnographers in technology to consider, for three key reasons:

1. It frames strategy as a practice of *decisions*, rather than a set of plans or goal metrics, which tend to dominate product strategy discussions in many tech contexts.
2. It conceptualizes these decisions as integrated, rather than siloed and without implications for other decisions that the business needs to make.
3. It conceptualizes strategy as inherently contextual. Customers and their needs are considered within the broader landscape of the market, other products, and other companies with other positions.

Importantly, Lafley and Martin emphasize that this framework can be exercised by anyone in any context—from a CEO crafting a generational strategy to a salesperson on the shop floor—who thinks about which customers to target, where to position themselves, and which pitch to lead with. We can use it as well, and will apply this strategy framework to ethnographic researchers in tech in Section 5, treating ourselves as a *"company in need of a strong strategy."*

We believe this classic strategy framework can be particularly helpful for driving impact in tech companies precisely because many product teams within these companies don't yet have the muscle and intuition for it. After decades of strategizing through low-cost bets and scaling up whatever sticks into high-revenue returns, they are pivoting on the spot. It's not just researchers who don't have a strong strategic muscle, it's also plenty of leaders and cross-functional stakeholders. This, we believe, represents an opportunity for applied ethnographers to help tech organizations actually think strategically and do strategy differently.

Section 4: The Challenges to Success

Our interviews with dozens of research and cross-functional tech workers surfaced two key categories of challenges for ethnographic researchers aspiring to influence strategic decisions. The first centers on the organization: How it is structured and how decisions get made. The second is about the researcher's orientation toward users and user data, rather than an orientation toward the business and commercial data. We acknowledge these challenges because they represent real constraints on the ability of researchers to drive strategic impact. But with some awareness, these factors can be worked around.

Organizational Barriers

For many of the people we interviewed, a core challenge to strategic research is cemented in org charts and organizational cultures where research functions operate in service to product design, which in turn, works in service to product management. This type of structure often results in a culture where research is looked upon to validate existing hypotheses that product managers have and have been working on with their design counterparts. A related observation we heard from respondents was that for many engineering-led product cultures (which are common in tech), problem solving is highly hypothesis driven. Getting engineering buy-in on projects can often require having a very clear hypothesis about what the problem even is in the first place, and offering testable solutions. This kind of organizational structure and culture can result in less appetite for research that may reject existing hypotheses, let alone answer more existential questions about what a product's winning aspiration or positioning should be.

Interestingly, we heard in interviews that some of the most "strategic" research is actually being done outside of product teams, in adjacent insights teams like product strategy or product marketing. These teams often have direct lines of communication with senior leadership, rather than having to go first through design and product management. Some tech companies have even built foresight practices into core business orgs like finance (e.g., at Amazon) to ensure that they have direct line of sight to senior business leaders. While many tech companies champion the flatness of their org charts and their bottoms-up approach to decision-making, this does not always reflect reality. Truly strategic decisions tend to be made very high up on the ladder before they are passed down to different teams to execute. In the best cases, it still takes an incredible amount of patience and sometimes a little luck for research insights to work their way into strategic decision-making.

An Orientation toward Humans before Business

A second core challenge we found in our interviews is that some researchers spend the majority of their energy focusing on user data, and much less time on the commercial data that senior leadership evaluates to inform strategic decisions. Such data may describe which products are generating money and which have not yet demonstrated ROI, potential costs and benefits of a restructuring, where competitors are investing and how much, and so on.

In some cases, we heard that researchers are even intentionally shielded from this kind of commercial data. This impedes researchers at all phases of project work. At

the start, it limits researchers from coming to a fresh understanding of the business's needs and opportunities through an ethnographic read of business data. Later on, it prevents researchers from making recommendations intelligible to their colleagues who currently make strategic business decisions.

Lastly, many ethnographic researchers view themselves as advocates for users first, and in service to the business second. This philosophy is noble. However, across our interview participants, we heard a strong recommendation to shift this mindset—particularly among in-house researchers—and view ourselves to be in service to the business as well.

Section 5: A Strategy for Strategic Ethnographic Research

Below we outline a potential strategy for ethnographic researchers who aspire to do great *strategic* research, using the Playing to Win framework to structure our observations. Now, it should be noted that we use the strategic framework loosely, and these observations should be read as "tips and tricks" and considerations rather than strict "dos and don'ts."

One overarching takeaway that helps set the scene is a conceptual one. Strategy can be usefully thought of as a verb, not just an adjective (as in "strategic") or a noun (as in "this product's strategy is XYZ.") We believe this conceptual shift can be helpful because it reframes strategy as a *practice* of decision-making for researchers and the organizations they're a part of, less so about the type of research projects we take on.

1. What Should Our Winning Aspiration Be?

The highest-level decision in the Playing to Win framework is: "*What should we set as our vision for ourselves?"* So for us, the question is: *"What should ethnography's winning aspiration be within big tech?"* Based on our interviews, we have discovered that many tech leaders across the functional spectrum do not have particularly well-established habits when it comes to thinking strategically and practicing strategy. As a result, we suggest that a winning aspiration for ethnographic researchers in tech could be to grow our teams' strategy muscles.

As applied ethnographers, we have been trained to think holistically and in an integrated way, paying attention to broader contexts. This positions us very well to grow the strategy muscle within our product teams. This aspiration also applies to researchers regardless of career level. Early career researchers can direct attention around strategic concerns by raising them to their cross-functional stakeholders (and put them on the agenda), as more senior researchers can facilitate conversations with

leadership from the right angles. People managers, too, have a role; they can lead by example in creating a team culture of immersion in commercial data, and kicking down doors that can inspire researchers, their stakeholders, and the broader organization.

2. Where Should We Play?

The next bucket of decisions in the Playing to Win framework is: *"Where should we play??"* Several compelling answers to this decision surfaced in our interviews with researchers and cross-functional stakeholders that can be divided into two categories depending on what kind of product team you support:

4. **Where to play if you are in a position to choose your team**. For example, if you are on the agency or freelancer side, are in a flexible role in-house where you can work with multiple product teams, or looking for your next position.
5. **Where to play if you sit in a mature product organization.** For example, "cash cows" that may be legacy products and/or highly monetized, so that product tradeoffs have immediate and direct financial implications for the company.

As we outline below, where you sit influences what strategy practically means in your context. Interestingly, we observed that many researchers don't tend to reflect much on what kind of product organization they are supporting, and therefore, what strategy means in their context and the role ethnography can plan in that context. A common thread, however, regardless of where you sit, is to focus on "territories" where making a mistake is costly to the business.

Where to Play if You Are in a Position to Choose Your Team

If you are in a position to choose what kind of team you support, we heard that some of the most impactful "territories" to play in are hardware, breakthrough technologies, and policy.

Focus on hardware. Some respondents emphasized that hardware technologies are generally more contextual than pure software. This means that there are more questions about where to play and how to win that are dependent on the broader physical and social context of use. Similarly, because hardware often has a longer product lifecycle and requires financial expenditures for manufacturing, it is comparatively costly to pivot if the business has made a wrong bet. These factors

make hardware a prime territory to play in for applied ethnographers looking to drive strategic impact.

Focus on breakthrough technologies. Other respondents pointed to the opportunity to play in spaces that are focused on very nascent breakthrough technologies. Roles supporting nascent products, like GenAI, are rich territories for applied ethnographers for many reasons. For example, in these product spaces there tends to be less behavioral data available, which means there may be more need to understand why people are doing what they're doing with the technology. Similarly, our interviews indicated that these product teams tend to have less intuition and pre-baked hypotheses about what will win in the market, increasing the appetite for ethnographic research to explore what users need. Finally, senior leadership often has less of a sense for what these technologies can concretely do for the business, increasing the need for a truly strategic perspective that connects dots among the market landscape, user needs, and the core business.

Go into policy orgs. Other respondents explained that they found more success driving strategic decision making by moving out of product research organizations entirely and into more policy-oriented teams. The reason being that in many big tech companies, strategic questions related to the risks and rewards of corporate positioning, brand perception, as well as unique value propositions tend to be driven more by policy and communications organizations than product groups. This creates more appetite for researchers trained to understand how the broader context impacts people's perceptions and behaviors.

Where To Play If You Are in a Mature Product Organization

If you are in a mature, "cash cow" product organization, we recommend focusing on risks to the business.

Focus core offerings on "risks," not just innovation opportunities. One successful approach to deliver more strategic work has been to reposition the value of research's core offering by highlighting risks to the business, not just opportunities for product growth. This is a space where our insights have unique power. Applied ethnographers have a window into people's lives that is broader than the moment when a user has their thumb on the screen or their nose in their laptop. They can see risks such as broader cultural trends or burgeoning alternatives coming that others may not.

3. How Will We Win?

The next bucket of decisions in the Playing to Win framework is: *"How should we win as we play in those spaces?"* Two core recommendations surfaced through the interviews that can help us better see how to win: Integrating many types of data and applying an ethnographic lens toward our own organization.

Integrate Many Types of Data to Connect the Dots

Two key factors emerged from our interviews that emphasize the importance of triangulating data to influence decision makers:

1. Senior leaders making decisions on strategy are inundated with insights from many different disciplines at the same time, on the same topics. From data science and product management to marketing and research, many different functions are charged with influencing leadership decisions.
2. To have confidence in strategic decisions like where to play and how to win, senior leaders need to sense how large an opportunity is and the potential tradeoffs of investing in one area over another.

Because of these two factors, the most senior ethnographic researchers we interviewed highlighted the importance of synthesizing data and insights across functions. To the extent possible, this analysis should include financial and commercial data to ensure that the message leadership hears about key strategic decisions is unified and clearly commercially relevant.

Turn Your Ethnographic Sensibilities inward, toward the Organization

Many of the most senior ethnographic researchers we interviewed highlighted the value of intensely observing the internal dynamics, power structures, financial performance, and communication channels. Observing organizational rituals like quarterly earnings calls can help researchers position their insights and recommendations in ways that are grounded in the open strategic concerns of the business.

4. What capabilities must be in place?

The next bucket of decisions in the Playing to Win framework is: *"What capabilities must be in place to deliver on our winning aspiration, the decisions on where to play, and the decisions on how to win in those spaces?"* Our interviews indicate that key capabilities

required to drive strategic decision making include commercial acumen, the ability to interpret user engagement metrics, and, of course, core applied ethnographic research skills.

Commercial Acumen

To influence strategic decisions in big tech organizations, many applied ethnographers have an opportunity to get better at understanding and interpreting commercial metrics. This involves understanding financial metrics like annual recurring revenue (and revenue growth) and key costs (e.g., for engineering headcount, content moderators, etc.), as well as market metrics like venture capital investments in a given product space.

Understanding of User Engagement Metrics

For many product organizations, user engagement data is one of the most critical inputs for strategic decision making. Being fluent and up-to-date with basic engagement metrics like monthly active users (MAU), daily active users (DAU), and more product-specific engagement metrics is critical for being seen as credible partners to product leaders making decisions on where to play and how to win. And while some in our community are already expert at interpreting and weaving in log data, as a collective we seem to have an opportunity to get even better at interpreting log data (Anderson et al 2009).

Ethnographic Sensibility

Many researchers, and especially applied ethnographers, have built their insights through excellence at interpreting different types of data, from the provocative space between what we see people do and what we hear people say. Our value has never come from the execution of a chosen research method but rather from the way that we see our data and apply it to problems.

More than any of our cross-functional stakeholders, we tend to bring an ethnographic sensibility that turns ambiguities from one perspective into a question that can be filled in from a different perspective. We need to expand our use of this capability as we fold new types of data and points of view into our synthesis of what matters to people and what matters to businesses. We can use this capability to see things that nobody else could or would.

Section 6: The Promise of Strategic Ethnographic Research

As business strategies evolve and the circumstances of research within companies mutate as a result, ethnographers are uniquely positioned to both ride and propel these waves. Since its emergence as a professional discipline, research has articulated its own specific value proposition to companies through reference to its proximity to users. When companies wanted to strategize more empathetically by thinking like their customers, we were there. We used our research skills to understand people's pains and preferences and then amplified their voices throughout the product development process. We have drawn from ethnography's openness to push a point of view—the user's point of view—that, at best, challenged conventional wisdom in just the right way.

But the business of technology has shifted, data science has progressed, and the meaning of strategy has changed within these companies. It is time for ethnography to expand beyond a value proposition that centers empathetic understanding of users toward one that draws from more of our strengths, one that combines our research skills with problem-solving sensibilities that ethnography encourages. Specifically, we see an opportunity for researchers to expand their focus from understanding how users interact with their products to understanding how businesses interact with users and the broader market.

Our interviews have convinced us that researchers within the EPIC community are actually in an excellent position to drive this change. This is because ethnography combines both a research method with an interpretive sensibility. Through it, we tell stories that are unexpected but true. These stories open up new questions, point to new opportunities, and identify new risks. We think inside and outside of different contexts, from the lives of our interlocutors to the initiatives of our companies, seeing each from a sideways angle. In doing so, we remain open to surprises that break open a single point of view into something that nobody else would have recognized. We also present a perspective that can keep leadership awake even as the inertia of their maturity seems to hum their imagination to sleep. Ethnography's unique technique of making the familiar strange and the strange familiar is the one that meets this new moment.

About the Authors

Morgan Ramsey Elliot is a partner at ReD Associates, and is based out of New York City. He also spent three years working as a senior researcher and research manager at Meta.

Charlie Lotterman is a consultant at ReD Associates.

Cameron Wu is a consultant at ReD Associates.

Notes

We are grateful for the members of our community who shared their experiences with us and helped us sharpen our thinking as we prepared this text. We are also grateful to our EPIC curators for their generous time and thoughtful feedback.

1. Other companies like Apple, Tesla, Nvidia, and other large-cap tech companies are certainly big but have not historically been home to large cohorts of applied ethnographers. At the same time, we believe that the argument in this paper—that researchers need to apply a strategic approach themselves to building the strategy muscle of their companies—is one that has broad relevance to any researcher working in any company, from a big tech firm or a start up to an automobile company or any other type of business where research happens.

2. Disclosure: ReD is collaborating with Roger Martin on building new perspectives on the future of business strategy.

References Cited

Anderson, Ken, Dawn Nafus, Tye Rattenbury, and Ryan Aipperspach. "Numbers have Qualities Too: Experiences with Ethno-mining." *EPIC Proceedings* 2009, 123–40. https://epicpeople.org/numbers-have-qualities-too-experiences-with-ethno-mining/

Antin, Judd. "The UX Research Reckoning Is Here." *Medium*, May 30, 2023. https://medium.com/onebigthought/the-ux-research-reckoning-is-here-c63710ea4084

Forbes. "Why Were Tech Stocks down in 2022-and How Long Will the Slump Last?" *Forbes*, January 20, 2023. https://www.forbes.com/sites/qai/2023/01/19/why-were-tech-stocks-down-in-2022-and-how-long-will-the-slump-last/

Henderson, Bruce D. *Henderson on Corporate Strategy.* Cambridge, Mass: Abt Books, 1979.

Ip, Greg. "As Big Tech's Growth and Innovation Slow, Its Market Dominance Endures." *Wall Street Journal*, February 8, 2023. https://www.wsj.com/articles/as-big-techs-growth-and-innovation-slow-its-market-dominance-endures-11675871487

Kushner, David. "Facebook Philosophy: Move Fast and Break Things." *IEEE Spectrum*, June 1, 2011. https://spectrum.ieee.org/facebook-philosophy-move-fast-and-break-things

Lafley, Alan G., and Roger Martin. *Playing to Win: How Strategy Really Works*. Harvard Business Press, 2013.

Martin, Roger, Iago Storgaard, and Filip Lau. "Strategy in 1000 Words." ReD Associates, June 13, 2023. https://rogerlmartin.com/docs/default-source/default-document-library/6-13-2023-red---roger-martin---strategy-in-1000-words.pdf?sfvrsn=b8ca0382_0

McGee, Matt. "Confirmed: Google to Stop Showing Ads on Right Side of Desktop Search Results Worldwide." *Search Engine Land*, February 19, 2016. https://searchengineland.com/google-no-ads-right-side-of-desktop-search-results-242997

Palandrani, Pedro. "A Decade of Change: How Tech Evolved in the 2010s and What's in Store for the 2020s." *Global X ETFs*, February 10, 2020. https://www.globalxetfs.com/a-decade-of-change-how-tech-evolved-in-the-2010s-and-whats-in-store-for-the-2020s/

Picchi, Aimee. "Meta's Value Has Plunged by $700 Billion. Wall Street Calls It a 'Train Wreck.'" *CBS News*, October 28, 2022. https://www.cbsnews.com/news/meta-stock-down-earnings-700-billion-in-lost-value/

Popken, Ben. "How US Venture Capital Has Grown in the Last 15 Years." *Flippa*, February 7, 2024. https://flippa.com/blog/how-us-venture-capital-has-grown-in-the-last-15-years/#:~:text=From%202008%20to%202022%2C%20the,class%20in%20its%20own%20right

Porter, Michael. *Competitive Strategy: Techniques for Analyzing Industries and Competitors*. New York, NY, 1980.

Richter, Felix. "Charted: There Are More Mobile Phones than People in the World." *World Economic Forum*, April 11, 2023. https://www.weforum.org/agenda/2023/04/charted-there-are-more-phones-than-people-in-the-world/

Spool, Jared M. "Strategic UX Research Is the next Big Thing." *Medium*, July 17, 2023. https://uxdesign.cc/strategic-ux-research-is-the-next-big-thing-915e025c3d8b

Stripe Partners. "The New Research Playbook." *Stripe Partners Viewpoint*, June 1, 2023. https://www.stripepartners.com/viewpoint/the-new-research-playbook-product-design-and-research-leaders-identify-3-shifts-that-will-unlock-the-value-of-research/

Sunderji, Aziz. "Fed's Aggressive Rate Hikes Are a Game Changer." *Wall Street Journal*, November 21, 2022. https://www.wsj.com/articles/feds-aggressive-rate-hikes-are-a-game-changer-11669006579

Teare, Gené. "What We Learned about Venture Funding During the 2008 Financial Crisis and the Pandemic as The Markets Face Fresh Turmoil." *Crunchbase*, March 31, 2022.

https://news.crunchbase.com/business/2022-vc-funding-outlook-compare-2008-financial-crisis-pandemic/

Walker, Peter, and Kevin Dowd. "State of Private Markets: Q4 2023." *Carta*, February 5, 2024. https://carta.com/blog/state-of-private-markets-q4-2023/

Why Don't We Post (Like We Used To)?: Multidimensional Context Collapse and Teenagers' Generative Responses to the Original Paradigm of Social Media

KRISTIN SARMIENTO, *Stripe Partners*
ISABELLE COTTON, *Stripe Partners*
JOSH TERRY, *Stripe Partners*
LEA VENTURA, *Meta*
SHAYLI JIMENEZ, *Meta*

With insights from original, multinational research and analysis of key social science studies, this paper deepens our understanding of how and why social media is changing. We explore the evolution of social media and "public broadcasting" over the past two decades, arguing that teen participants have shifted away from the original social media paradigm and towards more private and ephemeral forms of sharing. Building on existing theories of "context collapse", we examining how traditional boundaries of audience, time, and meaning blur on social platforms. Rich ethnographic research shows how teen participants have adapted to multidimensional context collapse through multilayered, generative strategies that challenge the perceived passivity of their "empty profiles", allowing them to reclaim social media as a place that reflects their generation's values.

Introduction

When EPIC was founded in 2005, the foundations of social media were still in their infancy: blogging platforms like WordPress dominated, and "The Facebook" was only just gaining traction in the dorms of Harvard. In 2004, the authors of this paper were still in their teens, eavesdropping on their older siblings' cordless telephone conversations about the latest school gossip (cf. March and Fleuriot 2005) or starting their own MySpace accounts (cf. Faulkner and Mellican 2007).

When the first users of social networks began to explore platforms like Myspace and Facebook, they encountered platforms that rapidly scaled up the number of people with whom they could connect and share about themselves. One-to-many communication—once the domain of traditional media broadcasting—was now something anyone could do. Anthropologists such as Daniel Miller in his *Why We Post* series studied this "public broadcasting" affordance of social media (2016). The original social media paradigm was a space where people could post about their lives to audiences known and unknown.

Sarmiento, Kristin, Isabelle Cotton, Josh Terry, Lea Ventura, and Shayli Jimenez. "Why Don't We Post (Like We Used To)?: Multidimensional Context Collapse and Teenagers' Generative Responses to the Original Paradigm of Social Media." *EPIC Proceedings* (2024): 335–363. https://www.epicpeople.org/why-dont-we-post-multidimensional-context-collapse.
ISSN 1559-8918.

Based on a multi-country ethnographic study conducted in early 2023 with teens, this paper argues that there has been a significant shift in the way teen participants use social media today in comparison to the original paradigm. In our study, a large proportion of participants no longer posted to social media feeds publicly, or at all—their profiles on various platforms such as TikTok, YouTube, Instagram, and Twitter were empty. If Miller were posing his question of *Why We Post* today, perhaps it would be *Why Don't We?*

To answer this question, we situate today's social media use within a historical context. We examine the evolution of social media over the past two decades, arguing that its technological affordances and mainstream adoption have facilitated ever-increasing forms of "public broadcasting" and subsequent "context collapse" on these platforms.

In its simplest definition, context collapse refers to the flattening of audiences into a single context (Marwick and boyd 2010) in digital spaces. It has the effect of destabilizing an individual's self-presentation, collapsing identities and behaviors that are normally kept separated for distinct audiences. Context collapse is inherent to "mediated technologies" (boyd 2008) and profoundly shapes the way that people interact online (Abidin 2021; boyd 2014; Brandtzaeg and Lüders 2018; Hogan 2010; Wesch 2009), evidenced particularly strongly in our study with teen participants. Although scholars such as boyd identified the phenomenon of context collapse as early as 2008, its effects seem to have become more pronounced with the evolution of social media technology and the affordances we have today.

Based on our original ethnographic research, we propose an expanded framework of multidimensional context collapse encompassing:

1. Flattening of audiences
2. Distortion of time
3. Instability of meaning
4. Conflicting purposes of platforms

Furthermore, our analysis reveals the generative, nuanced strategies that teen participants employ in response to context collapse, allowing them to reclaim social media as a place that reflects their generation's values. Teens in our study were adept at:

1. Managing audiences
2. Hacking time
3. Playing with meaning
4. Engaging intentionally with platforms

Our core insights were developed during an initial study of teen participants' social media use in Spain, the UK, and South Korea across a variety of platforms including YouTube, Twitter, WhatsApp, Snapchat, Pinterest, TikTok, Discord, Facebook, BeReal, Reddit, Instagram, LINE, Kakao Talk, and Cyworld. However, subsequent studies with teen participants in six additional countries (Japan, Brazil, France, Italy, Australia, and the US) demonstrate that these insights hold across a variety of locales. This suggests a broader shift in the way younger generations use social media.

Although numerous scholars over the years have elucidated the ways in which teens have crafted adaptive strategies for engaging with social technology (Abidin 2020; boyd and Marwick 2011; Gopffarth et al. 2022; Ito et al. 2010; Rangaswamy and Yamsani 2011), in this paper, we highlight specific behaviors from teen participants that demonstrate how they operate, by default, with self-awareness of multidimensional context collapse. In particular, we discuss the now commonplace practice of "not posting". When teen participants spoke about "posting", they equated it to the public, permanent sharing of content about themselves to social media feeds—content that remains on their profile. The majority of study participants tended to avoid this form of self-disclosure. Instead, we observed that most sharing behavior is now private and ephemeral by default, and public and permanent only through effort (with deliberate strategies).

We propose that the patterns observed in our research of participants' empty profile grids across a variety of apps do not reflect an abstinence from social media for those teens, but rather an intentional engagement with it that is a generative response to context collapse. Their strategies and alternative forms of "posting" underscore the resilience and intentionality of young people in shaping their digital experiences (Gopffarth et al. 2022), challenging notions of passivity and vulnerability. To answer the question *Why Don't We Post?*, we must look beyond the behavior of public posting to the multifaceted approach that they use to express themselves and connect with others online.

To conclude our paper, we will discuss possible new directions for understanding teen experiences with social technology, which help us to view the relationship between teens and technology not as deterministic, but as negotiated.

The Evolution of Social Media

A review of two decades of social media (with a focus on globally popular platforms) demonstrates how its affordances have given rise to increasing degrees of "public broadcasting", that is, the dissemination of content or other digital "records" to large audiences online.

Social Networking's Early Years: New and Extended Forms of Social Connectivity (2004–2009)

In the years leading up to the launch of "The Facebook" at Harvard University in 2004, teenagers were already making use of preexisting technologies that enabled digital connection with friends and peers. In particular, cordless phones and cellphones introduced more mobile and "on demand" formats of communication and socialization. March and Fleuriot (2005) illuminated how these phones allowed teenagers in the US and UK to carve out spaces to connect with friends out of earshot of family members. Similarly, Asokan (2008) highlighted how teenagers in India developed subtle ways of using technology like cellphones to navigate the blurring of personal and public spaces, creating "their own space in the heart of social activity at home".

Extended forms of online social activity began to take off with platforms like WordPress and YouTube. Blogging and vlogging became increasingly popular ways of "getting noticed, showing-off, [and] being overheard" by global audiences (Faulkner and Melican 2007). These technologies enabled a diverse array of individuals and communities to express themselves—creatively, entrepreneurially, or even in the pursuit of social change.

At the same time, the practice of blogging began to trigger reflection and engagement with questions of online identity management. Now, content creators needed to consider how to present themselves to a variety of audiences, both known and imagined (boyd 2014; Faulkner and Melican 2007; Marwick and boyd 2010), and how to convey a message that would experience narrative changes as it got passed along to more friends.

Amongst the early social networking sites—or "networked publics" as danah boyd referred to them (2008)—MySpace was a site that became central to peer sociality. In the words of one of boyd's participants, "If you're not on MySpace, you don't exist." Alongside MySpace, Facebook was gaining traction: it went from 12 million users in December 2006 to 100 million in less than two years (Brügger 2015).

Social networking was becoming a global phenomenon, evidenced by a number of similar sites launched in other countries, such as Cyworld in South Korea in 2005 (boyd 2008; Horst and Miller 2012). However, none had attained the reach of MySpace and Facebook, where teenagers relished their newfound ability to create elaborate online profiles, glean more information about their friends, and display different dimensions of themselves (boyd 2008; Ito et al. 2010).

While teenagers primarily used these sites to connect with close friends and people they already knew, these technologies allowed "many teens [to] move beyond small-scale intimate friend groups to build 'always-on' networked publics inhabited by their peers" (Ito et al. 2010). Social networking sites introduced public facing aspects to social connection and self-presentation that had not existed previously. Implicit social dynamics were made explicit. Features like publicly visible friend lists and the ability to comment on profiles allowed others to "overhear" a conversation to which they may not have been privy in the physical realm. When Facebook introduced the News Feed feature in September 2006, this process of seeing actions taken by friends on the site became core to the experience (Ito et al. 2010).

From Social Networking to Social Media: Mobile, Ubiquitous Access to Visual Content Creation and Distribution (2010–2014)

By the early 2010s, smartphones and mobile connectivity had become mainstream (boyd 2014; Madden et al. 2013), allowing constant access to the Internet (Campbell 2013; Ling 2012). As Rangaswamy and Yamsani (2011) highlight in their exploration of teen online practices in India, many teens' first experience of the internet was on a mobile phone.

The "mobile internet", which connected users synchronously, asynchronously, and while in transit (Campbell 2013), played a crucial role in facilitating the ubiquitous presence and embedded practices of online social connection in the daily lives of teens (boyd 2014; Horst and Miller 2012). Sites like Facebook and Twitter became central to their coming of age experiences. No longer subcultural, they were now normative in Western contexts (boyd 2014; Horst and Miller 2012).

It was at this point, in the 2010s, that a shift in terminology was needed, to signal a transition from social networking to social *media*. Most social networking sites had evolved to popularize the creation and sharing of content, and the launch of platforms like Instagram, Snapchat, Tumblr, and Vine prompted a move towards visual and image-based content (boyd 2014; Highfield and Leaver 2016; Miller et al. 2016).

The default design of these social media platforms was their support of sharing content with broad audiences, facilitating new methods of distribution "by explicitly

or implicitly encouraging the sharing of links, providing reblogging or favoriting tools that [reposted] images or texts, or by making it easy to copy and paste content from one place to another" (boyd 2014). boyd notes that teens embraced these practices and were enthused about uploading photos, tagging people, and commenting online because it provided a means to extend the enjoyment of shared experiences with friends (boyd 2014). Smartphones facilitated constant content creation, while social media sites facilitated its broadcasting (Highfield and Leaver 2016; Marwick 2015).

In this period of social media growth, visual content introduced a new way to communicate, fashion, and perform identities (Highfield and Leaver 2016). However, questions of online identity expression were about to become even more complicated with social media's next innovation: the algorithmic feed.

Algorithmic Social Media: Personalized Feeds and the Rise of Influencer Culture (2015–2019)

The nature of information sharing, self-presentation, content consumption, and online social connection changed dramatically with the more widespread shift from chronological to algorithmic feeds in the mid 2010s on major platforms such as Twitter and Instagram (Schulz 2023). Now content that appeared to users was personalized by the algorithm, which excluded some pieces of it while highlighting others (Gillepsie 2014). Algorithms upranked content that was aligned with individuals' interests, and posts that had mass appeal across platforms were referred to as "trending" or popular. Upranking created the possibility that any trending piece of content could potentially be launched into others' feeds, beyond the anticipated audiences.

Algorithms "manage our interactions on social networking sites" (Gillepsie 2014) when they tailor what we see and therefore interact with. However, the underlying mechanism is a two-way street—a feedback loop between people's behaviors on the app (clicks or taps, likes, views, comments, etc.) and the algorithms' subsequent recommendations (Gillespie 2014; Klug et al. 2023). In other words, "What we see is no longer what we get. What we get is what we did and that is what we see" (Bucher 2018).

However, algorithms cause a second, more significant effect: not only do they allow users to shape what they themselves see, but also what others see. People can engage in practices that "amplify their efforts" and ensure their content gets picked up by an algorithm and distributed to larger audiences (Abidin 2020; Gillespie 2014),

for example, through the strategic use of hashtags. Gillepsie (2014) highlights how teens on Facebook would "tag their status updates with unrelated brand names, in the hopes that Facebook [would] privilege those updates in their friends' feeds." Thus, people are "not just joining a conversation…[they] are redesigning [their] expression so as to be better recognized and distributed by…search [algorithms]." What we come to know about others online is both socially, and arguably, platform constructed.

The affordances of algorithmic social media set the stage for, and were inextricably linked to, the rise of influencers on social media in the mid 2010s, or what many scholars term "micro-celebrities" (Abidin 2015, 2020; Marwick 2015; Senft 2013). Influencers' practices were aimed at maximum visibility and exposure, in order to attract and compete for mass audiences historically limited to broadcast media (Abidin 2015, 2020). While influencers' direct, commercial strategy included integrating "advertorials" into social media posts to promote sponsored products, the more tacit social strategy underlying their practices was that of crafting an aspirational online persona and identity for their viewers (Abidin 2016).

Although YouTube was an important platform for influencer vlogging content, much of the curated, photo-based persona content was generated on Instagram during this period. Instagram micro-celebrities had amassed vast audiences of followers, and their highly popular content was being pushed to the feeds of hundreds of thousands of people everyday. But popular is not the same thing as viral—and with the rise of TikTok, public broadcasting was about to take on a whole new meaning.

Viral Social Media: Remixing Drives New Forms of Visibility and Influencer Practice (2020–Present)

TikTok was launched in 2016, grew significantly from 2018, and then experienced a massive surge in teen usership in 2020, coinciding with increased internet use during the pandemic (Abidin 2020; Klug et al. 2023; Wei 2021). Its algorithm ushered in what Abidin (2020) referred to as a "new frontier of social media" in which TikTokers were "actively and very quickly adapting [to] the latest trends and viral practices…to aim for the For You Page, or the 'golden ticket' that would allow one to gain an immense number of followers overnight".

On TikTok, the nature of influencer culture shifted. Whereas Instagram was largely predicated on the careful curation of a "singular coherent persona or style" by "'staging'...an 'Instagrammable' lifestyle that was aspirational and pristine", TikTok influencers crafted "relatable" performances that felt entertaining and accessible to wide audiences, through a kind of "calibrated amateurism". Wei (2021) sums this up

by saying "Whereas Instagram is performative, TikTok is performative and self-aware."

In tracing the evolution of social media over the past two decades, it becomes clear that the scale of public broadcasting has increased over time, owing to the progression in the technical capabilities of the mobile internet and these platforms, as well as their mainstream adoption and global success. This, in turn, has created the conditions for increased context collapse. Social media platforms bring together multiple, diverse audiences into one space and make it possible for information to be accessed across space and time at an accelerated rate. When users engage with today's social media platforms, they find themselves needing to manage a greater number of collapsed social contexts online.

We now turn to the present era through the lens of our ethnographic research. After introducing our study's methodology, we propose a framework that reflects teen participants' experience of multidimensional context collapse, and the generative strategies that they employ to navigate it.

An Ethnographic Study of The Emerging Paradigm of Social Media for Today's Generation of Teen Participants

Methodology

In March 2023, we carried out a multi-method ethnographic study of teen participants' social media use across a variety of platforms including YouTube, Twitter, WhatsApp, Snapchat, Pinterest, TikTok, Discord, Facebook, BeReal, Reddit, Instagram, LINE, Kakao Talk, and Cyworld. The study was a collaboration between Stripe Partners and Meta. We conducted research with 127 teen participants (ages 13-17) in South Korea, Spain, and the UK[1], to bring a global perspective to emerging teen social media practices. Our research methodology comprised a diary study, interviews with friendship groups, as well as participant observation in public spaces popular with teens.

The diary study generated a wealth of digital artifacts, painting a picture of the daily online practices of study participants—blank social media profiles, playful avatars, intentionally vague usernames, blurry photos, insider memes and various group chats. This data was analyzed thematically, allowing us to explore teen participants' aesthetics, cultural codes, and the values these represented.

In our interviews with friendship groups, we explored their social media values through several activities such as app ecosystem mapping, explorations of online

norms and etiquette (in the words of a participant: "cringey and "not cringey" behaviors), and ideation exercises that prompted them to design the "best" or "worst" social media app. We analyzed this data using a grounded theory approach, coding themes from both interview responses and the artifacts generated in the activities.

For our immersions into public spaces, young adult participants guided us through locations popular with youth such as arcades, shopping malls, basketball courts, and Korean cafes. Observing study participants' socialization within these spaces and taking part in a number of their favorite activities allowed us to gain a more holistic perspective of the connection between their online and offline practices, as well as embed ourselves more fully into their cultural contexts. The data from these immersions was incorporated into our analysis of the interviews.

A Framework for Multidimensional Context Collapse

Rooted in our ethnographic work with teen participants, we propose a context collapse framework that consists of four key dimensions:

1. Flattening of audiences: in contrast to offline experiences, online audiences are significantly more vast and mixed, and social circles are often bundled together.
2. Distortion of time: time no longer possesses a linear or sequential nature on social media—the past does not stay in the past, the present doesn't stay in the present.
3. Instability of meaning: content can circulate unpredictably, with the poster having varying levels of awareness of how this is taking place. The meaning of the content can be continually remixed and reinterpreted.
4. Conflicting purposes of platforms: social media platforms bring together multiple purposes for their use, with one example cited by teen participants being the personal vs commercial uses of platforms.

All four dimensions play a significant role in study participants' experiences of context collapse on social media today. At the same time, the constituent elements are highly interconnected, creating a complex set of considerations for engaging online.

For example, when teen participants decide to post a photo of themselves to a social media feed, they are aware that multiple audiences may view it, including friends but also family members, teachers, and potential employers. For many participants, this is reason enough to stop posting entirely: Isabella[2] from Spain told

us, “If you publish something, it's totally open, anyone can see it. That’s why I don’t want to publish stuff.”

Added to this, their photo may still be visible to these audiences at a future moment in time, representing them in a way that is no longer consistent with certain characteristics or interests. In South Korea, many teen participants spoke about the desire to obscure their face. In one group interview, participants discussed the lengths they would go to, either covering their face in photos or only posing in group selfies. Ye Joon explained that pictures of their face might circle back in an undesirable way: “People might really look at your face. They might capture it. If it's captured it can be used in the future against you. You don’t want people to think you're a child.”

Furthermore, study participants are wary that their content might be reshared into other online contexts, and even edited by others in such a way that it no longer reflects the original message intended by the photo.

Given this complexity, our four dimensional framework is not an attempt to reduce context collapse to a strict formula or show it as a static phenomenon. Rather, our aim will be to illustrate each dimension’s nuances and the roles they play together in creating teen participants’ experiences of context collapse online.

Our ethnographic study revealed that in response to each dimension of context collapse, participants develop strategies to mitigate, or even take advantage of its effects.

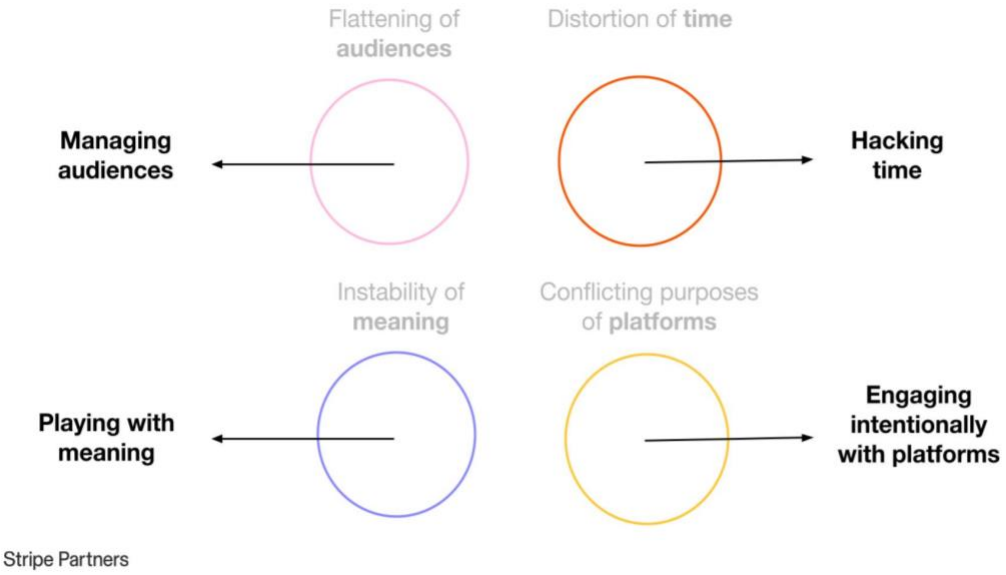

Study participants:

1. Manage audiences by using different apps, creating multiple accounts, and meticulously (re)adjusting the settings of follower lists and individual pieces of content to correspond to nuanced social circles.
2. Hack time by defaulting to ephemeral forms of interaction, curating their published memories, carving out spaces for private memory banks, and developing content in the present with a view towards the future.
3. Play with meaning by decentering themselves in posts, subverting platform norms, communicating indirectly, and encoding messages online.
4. Engage intentionally with platforms by humanizing and personifying app algorithms to help them rationalize decisions and retain a sense of agency.

In all of these strategies, an important feature of many teen participants' behavior is the tendency to avoid public posting to social media feeds, which represents a distancing from the original paradigm of social media. Study participants' diary submissions featured a plethora of screenshots of blank profiles, and during interviews, the phrase "I don't post much" came up constantly. However, this practice cannot be viewed as a wholesale rejection of social media. On the contrary, teen participants find many other ways of engaging meaningfully and intentionally with friends, peers, and their wider communities on these platforms. In other words, rather than passively accepting context collapse, study participants actively adapt their behaviors to reclaim social media as a place of their own (boyd 2014).

Below we explore the dimensions of context collapse in detail. For each dimension, we will highlight previous scholarly work that has contributed to our understanding, the perspectives and experiences of teens in our study, and the generative strategies that they use in response.

Context Collapse Dimension #1: Flattening of Audiences

Context collapse has often been closely associated with the collapse of audiences online, and multiple scholars have examined the subsequent effect on self-presentation and expression in digital spaces (boyd 2008; Hogan 2010; Marwick and boyd 2010; Wesch 2009). boyd (2008) described context collapse as the challenge of delineating potential different audiences online—the boundaries that are normally present in offline experiences do not apply in online spaces, resulting in a bundling of audiences. For example, teen participants described how those in positions of authority, such as parents, teachers, coaches, or even potential employers were all possible witnesses to their digital interactions.

Much of the scholarly understanding around the effects of audience collapse is rooted in Goffman's framework (1959) of social situations and performances. In the absence of clear boundaries online between front and back stages, and between different audiences, social contexts become collapsed, disrupting the process of impression management. The design of many social media features for public broadcasting has a tendency to bring what would otherwise be backstage behaviors (such as more "authentic" or personal expression) onto the front stage, turning them into performances which require users to reckon with how to present to a mix of audiences.

Michael Wesch's work on YouTube vloggers (2009) described the phenomenon of audience collapse in this way: "The problem is not lack of context. It is context collapse: an infinite number of contexts collapsing upon one another into that single moment of recording." Applying Goffman's theory of stages and performances, he described how this phenomenon created a "crisis of self-presentation" for vloggers. They had to imagine an infinite number of potential audiences and "address anybody, everybody, and maybe even nobody all at once."

We found that these observations regarding the flattening of audiences held true for teen participants in all markets, who feel it is difficult to delineate different social circles online. In offline contexts, teen participants have intricate social ecosystems with varying degrees of relational closeness between individuals and groups. However, when it comes to social media, study participants are highly aware that apps do not have structures that reflect these nuanced degrees of social intimacy.

Online, the boundaries that exist between social groups in the physical world have no clear digital equivalent. Any digital interaction, be it sharing content, commenting, or even liking a post can be broadcast to an audience where many forms of social diversity are collapsed into one. Some platforms have settings that allow for sharing with different audiences, such as "Close Friends" on Instagram or "Best Friends" on Snapchat, but these introduce binary categorizations that force, and make public, social distinctions that are not aligned with or expressed explicitly in offline settings (as Ito et al. 2010 have noted). When study participants were asked to design their dream social media app, one group of South Korean participants went well beyond the binary Close Friends feature, designing a social intimacy scale from 1-100 that would intelligently sort what content would be shared with who, creating nuanced layers of access.

Teen participants are conscious of the mix of age groups in online audiences, or those who "aren't the same age as them". This was especially highlighted amongst

South Korean participants, where strong age hierarchies and respect for those who are older than oneself (not just adults, but older teens) are an embedded aspect of their culture. In South Korea, it is most common to socialize within one's year group—in fact, one of our research partners informed us that the Korean word for "friend" (친구 (chingu)) is generally only used for someone of the same age. This attention to age means that teen participants avoid being on apps where audiences are thought to be "too old" or "too young", and consider it appropriate to use certain apps at different stages of adolescence.

Furthermore, study participants are confronted with the fact that their parents are sometimes on the same apps as them, including in the years before they themselves arrive on the app. One participant lamented over what he described as "moms posting pictures of their kids on social media, trying to expose them, like 'First day of school, fam!'".

In previous work on context collapse, boyd (2008) pointed out that in many cases, online audiences are not only bundled together, but also unknown or invisible. Thus, she often referred to how her teen participants would "imagine" their audiences on social media (2014). However, we observed that participants in our study take this one step further—they often *assume*, by default, that their content will be seen by any member of these vast audiences, and this mindset shapes their behavior and forms of engagement on many different apps—in particular, their tendency to avoid posting to social media feeds.

The Flattening of Audiences Online Leads Teen Participants to Manage Audiences

In response to audience collapse, teen participants are constantly managing audiences using a variety of apps, tools, and settings.

Study participants find ways to create multiple layers of social intimacy online, crafting separate spaces for their different social groupings and for their own personal use. For example, study participants choose to use different apps to keep audiences distinct and to tailor their online self-presentations accordingly. For most participants, Snapchat is largely used to interact with friendship circles, while Instagram is used to add broader audiences to their network such as acquaintances or to advertise entrepreneurial pursuits. On the other hand, TikTok and Twitter are often used to "be the audience" rather than "be the performer": these are platforms where they can be entertained or receive in-real-time updates about interests or favorite creators.

Within a single app, one "tried and true" way of carving out audience boundaries is through the creation of multiple accounts (Ito et al. 2010). On apps such as Instagram, having two accounts is commonplace, one for presenting a more "curated" set of photos or videos, usually to somewhat larger audiences, and one that is a more "personal" account, often described by teen participants as their "silly" or "spam" account that they use to send casual updates and engage with closer friends.

However, two accounts is just the starting point for many study participants. One of our participants, Ha-Eun, related that she has "seven different accounts, which include three for different friend groups, one for my interest in drawing and also one just for studying". Her approach is emblematic of how participants across the study use separate accounts for different purposes and audiences. In a similar way, Julia from Spain informed us of the typical process: "You start off with a main account, where you follow loads of people, and then you make a more private one." Her friend Alberto agreed, commenting that "You start caring less and less about the main account. I feel we all just use our various private ones, because that's where the real friends are and where I can show them what I really want them to look at."

In South Korea, many study participants create accounts expressly for following K-pop stars and engaging with fandoms, in order to keep this part of their identity separate from other self-presentations. Study participants told us that some people set up Twitter accounts specifically for this purpose: "People keep their secret hobbies on Twitter." This practice not only helps them to avoid potential embarrassment if it isn't considered "cool" to be a fan of a particular celebrity, but also provides a means to enjoy relationship building with a wider community of people who share their passions or interests.

Finally, some participants also create completely private or "secret" accounts with no followers at all, using it as a "personal diary" to store digital memories or journal their thoughts.

At a more micro level within the apps, teen participants are adept at adjusting the boundaries around audiences through the meticulous use of settings like "Close Friends" on Instagram. This feature's design introduces an audience binary that does not always align with the nuanced social ecosystems teen participants belong to. As such, we observed participants regularly readjusting the Close Friends list for individual posts, adding or removing users to ensure that their content will reach a specific set of people each time. Doing so is an effective, albeit involved way to navigate audience collapse, but it also serves a secondary purpose: it allows them to

play with social inclusivity and exclusivity, whereby "in group" status can cultivate relational bonding with specific individuals or groups.

These behaviors illustrate that despite the challenges of flattened and bundled audiences online, study participants find ways to unbundle them, express themselves, and develop shared experiences with their various social groups. Although study participants' public-facing profiles often contain no posts, a significant amount of posting activity is taking place on private accounts or on those with very select audiences or specific purposes.

Context Collapse Dimension #2: Distortion of Time

Time operates differently online—the boundaries between the past, present and future are blurred. As Brandtzdaeg and Lüders (2018) noted, time no longer possesses a linear or sequential nature on social media: the past does not stay in the past, the present doesn't stay in the present. This characteristic of digital spaces contributes to context collapse, which Michael Wesch highlighted in his 2009 study of the webcam and YouTube videos: "The images, actions, and words captured by the [webcam] lens at any moment can be transported to anywhere on the planet and preserved (the performer must assume) for all time. The little glass lens becomes the gateway to a black hole sucking all of time and space—virtually all possible contexts—in on itself." This temporal blurring affects how users manage their identity and performances, since it forces people to consider their self-presentation over time and not just in the moment (Brandtzdaeg and Lüders 2018; Wesch 2009).

Time distortion online results from a number of properties inherent to social media, such as what boyd (2008) terms "persistence" and "searchability." Persistence refers to how "online expressions are automatically recorded and archived", meaning that the potential to view or interact with what has been digitally inscribed extends in time far beyond the moment of inscription. Searchability refers to how digital content can be retrieved by a user at will. These characteristics of online interactions stand in stark contrast to the synchronicity and ephemerality of offline interactions (boyd 2008, 2010; Hogan 2010). One participant we spoke to in the UK articulated this in her reflections on Snapchat, an app where chats aren't saved: "Snapchat is the most like real life…in real life, if you have a conversation, it disappears afterwards."

Teen participants are highly conscious of the "persistence" of digital recordings, but they characterize this phenomenon even more strongly, consistently referring to interactions, particularly posting to feeds and messaging, as "permanent". The permanence of inscriptions online runs counter to the fact many teens in our study simply desire to connect over trivial, everyday moments with their friends. As Jenny from Spain told us, "I do have a main account where I post things that are actually

nice but most of the time I like sending pictures of something like me brushing my teeth to my friends." Teen participants intend for the sharing of the regular and mundane to be something ephemeral, a passing comment rather than a mark of their identity for years to come.

The permanence of social media content is particularly evident to teen participants given the amount of change that they experience during the journey of adolescence, both personally and socially. As many developmental scholars have noted (McNelles and Connolly 1999; Kilford et al. 2016), the teenage years are characterized by a significant number of transitions, both micro and macro, and teens are continually navigating the shifting landscape of identity formation and social maturation. Against this backdrop, the effects of time distortion on social media are felt very acutely by teen participants, to the point that terminology around these effects have become common parlance. For example, South Korean teens from our study use the term "dark history" to describe the idea that anything from your past can at some point "pop up and embarrass you" in the future. The affordances of algorithms make content not only searchable, but "discoverable" at any time (Abidin 2021).

The Distortion of Time Leads Teen Participants to Hack Time

In the midst of time distortion online, teen participants utilize three principal techniques to counteract or mitigate the effects of online permanence and discoverability: using ephemeral features for daily interactions, carefully curating publicly posted memories, and fashioning out spaces for private nostalgia. In general, the mindset of study participants is to develop content in the present with a view towards the future.

While the social media profiles of teens in our study often appear derelict, containing no signs of their everyday lives, other spaces online are teeming with regular activity. Indeed, for the majority of daily interactions, teen participants gravitate to platforms or features that facilitate ephemerality and imitate the nature of offline conversations. For example, messaging features on Snapchat such as disappearing photos and dynamic chat groups have led to it becoming the primary domain of many UK and Spanish participants' social lives online. The ephemeral nature of the content in these spaces reduces the sense of pressure associated with more permanent posting to social media feeds. Participants across markets referenced ongoing "photo chains" of random pictures of the floor or a portion of

their face that they would send to a friend. As one participant said, "You just can't stop sending them!"

Although teens in our study largely avoid online permanence, there are occasions in which they do choose to share content in more public, persistent spaces. This is typically in relation to memories. Several participants spoke about reserving the act of posting on their profiles for special moments that they felt particularly proud of. However, given that teen participants are experiencing constant changes as part of growing up, many of them routinely engage in the practice of "cleaning" their profiles by deleting published memories that do not match how they want to present themselves in their current life stage: "Once a month I go through my Instagram highlights and clean them up" - Woo-Jin, South Korea. Thus, their empty profiles do not necessarily signify that they have always been empty, but rather that participants are monitoring and editing them regularly to continually adjust their self-presentations.

In addition to sharing their memories with other people, study participants also carve out more private spaces for personal nostalgia that enable them to look back on their past and relive certain moments in time. For example, participants archive posts on Instagram or create and save drafts on TikTok which are visible only to themselves. Participants in our study often have hundreds, even thousands of drafts, demonstrating a highly active engagement with social media, despite any evidence of public broadcasting. As Lucy from the UK put it, "I have hundreds of drafts and just one post. Drafts are an easy way to look back on time with my friends without the pressure of posting."

Study participants use draft spaces as "enhanced camera rolls" that also allow them to experiment with different modes of self-presentation: "I save the ones I like the most to my 'Favorites'—the videos that I want to go back to" - Estella, Spain. Although TikTok is known as a public broadcasting-centric platform, the teen participants we spoke with do not actively, regularly post there. Most often, they enjoy scrolling through their draft videos or photos on their own or with the friends they had recorded them with to reminisce and bond over shared experiences. Stephanie from the UK had made this a highly regular practice: "I love making TikTok videos and just keep them in my drafts. I have over 4000 drafts and I go back and watch them all the time and go 'awww', it's so fun."

Participants' strategies for hacking time allow them to actually use time collapse to their advantage: whereas offline interactions are truly fleeting, generating digital memories allows them to traverse time and space within their own private time capsules. Our research demonstrates how teen participants shift fluidly between varying degrees of temporality online. They hack the architecture of social media

platforms to both remediate the effects of online permanence in public spaces, and also to leverage the temporal affordances of these technologies for both social and reflexive practices.

Context Collapse Dimension #3: Instability of Meaning

Even in everyday communication, we often say that "message received is not always message sent" and that our intended meaning can be "taken out of context". However, digital spaces like social media amplify and accelerate the reinterpretation of meaning in profound ways. Previous scholarly work has not tended to focus on "meaning" as a part of context collapse, but we argue that it should be included as a dimension. Teen participants are highly aware of its instability online—a piece of content can circulate in unpredictable ways and take on a life of its own. Once teen participants put something out into the world of social media, they know that it can be difficult or in some cases even impossible to be fully taken back. Even if they later decide to delete it, it may have already been screenshotted or disseminated. In any digital space, from the most public to the most private, there is always the possibility that one's message can be reshared and reinterpreted by someone else. As Luke, 17, from the UK warned us: "Social media is pressure: everything you say can be controversial and taken out of context."

boyd's property of "replicability" (2008) is an example of the instability of meaning, where the ability to duplicate a digital record allows a user to transfer it from one context to another. Abidin (2021) updated this property to "decodability", referring to how "content can be duplicated but may not be contextually intelligible". However, we suggest that for our study participants, the most salient understanding of this facet of context collapse is "remixability". Platforms such as TikTok have ushered in a new set of tools, behaviors, and norms that have remixing at their core. These foster creativity and playfulness, but also magnify the effects of context collapse because of how easy it is for users to quickly edit, reinterpret, and redistribute other users' content (Abidin 2020; Leaver et al. 2020; Klug et al. 2023; Wei 2021).

Furthermore, with an abundance of social media apps, there is significant circulation and cross-pollination of content between and across them. A piece of content with a particular meaning that is created on one app can end up on a completely different app with its own distinct set of norms and audiences. Through the combination of algorithmic technology and users' ability to remix, reshare or bring attention to content (for example, via comments and likes), a double context

collapse can result: there is the broadcasting of content to unknown audiences, as well as the instability of meaning that can result from these audiences' interpretation or reinterpretation of the content. Ji-Won in South Korea described how this reality made her reluctant to post selfies: "People might screenshot my selfie and use it in the future. They might evaluate or judge my look. In the future the selfie might not look nice."

From the perspective of study participants, we also heard that social media platforms destabilize meaning because of the tensions that exist between being interpreted as "authentic" or "cringey". While the common rhetoric of self-presentation is "just be yourself", they find this maxim constraining and challenging to realize, given that their content can be remixed and presented in different ways by others on the platform, and also because of the flattening of audiences online. They understand that people have various opinions on what constitutes the desired level of authenticity, that is, alignment between one's real life and their online expression. As Marwick and boyd (2010) have noted, "there is no such thing as universal authenticity; rather, the authentic is a localized, temporally situated social construct that varies widely based on community." As such, teen participants feel like they are walking an authenticity interpretation tightrope, often leading some of them to avoid posting on social media feeds altogether. However, this doesn't stop them from finding other ways to express themselves.

The Instability of Meaning Leads Teen Participants to Play with Meaning

It was particularly fascinating to explore how teen participants played with meaning in response to the frequent decontextualization of what they shared online. To preemptively shield from being misinterpreted, we noted how study participants intentionally worked to create distance between themselves and their "messages" in online spaces. They often use a variety of image, sound, editing tools and app settings in service of three important strategies: decentering the self, subversion of platform norms, and indirect communication.

The rise of influencers alongside affordances for public broadcasting has led to a hyper-focus on personal brand, curated aesthetics, and self-promotion on social media, often in the form of "selfies" (Abidin 2016; Marwick 2015). Within this context, we observed how study participants across markets sought ways to deflect focus from the self when sharing personal or biographical content. Although they want to provide meaningful glimpses of their personalities, interests and special moments, they don't want to come across as vain or self-absorbed, or be associated with influencer culture.

As a consequence, we encountered numerous examples of selfies that were blurred or where faces had been intentionally obscured, as well as a preference for photos of scenery that gave indications of participants' environments but didn't feature themselves explicitly. Creating some level of vagueness or distance between the self and what is shared helps protect against reinterpretation of personal content, essentially facilitating scenarios of "plausible deniability". Another practice we observed was using black cover photos for Instagram highlights, labeled with ambiguous titles or no titles at all. By subduing the aesthetics of these photo highlights, teen participants hope they might "fly under the radar" (Abidin 2021), more likely to only be noticed by friends "in the know".

The second strategy for playing with meaning is the deliberate, but often subtle, subversion of aesthetic and behavioral norms on platforms. A number of participants discussed engaging in "anti-curation" practices on secondary profiles, created to intentionally post photos or videos that do not align with the "typical" aesthetics of apps such as Instagram. We noted that participants loved to use tactics like these, as well as irony, satire, and humor to reclaim the context and meaning of shared content. For example, memes often create a set of fast moving "in-jokes" that can be only understood by insiders but not by the broader public.

The third form of playing with meaning is indirect communication. This strategy allows teen participants to create a protective layer of ambiguity while they seek to initiate or strengthen social relationships. For example, study participants use BeReal app's structure for reciprocal photo sharing to discreetly find out what their friends are doing. Instead of posting a photo to share something about themselves, they upload a random one just to trigger the ability to view others' posts. This subtle form of engagement is preferred to asking their friends outright, "What are you up to?"

Participants also go to great lengths on other social media apps to communicate indirectly with or detect interest from romantic crushes. On Snapchat, one participant, Mariela, UK, explained that she used the app's settings to add her crush to a "group" that contained no other users, allowing her to send him a specific photo without him knowing that it wasn't sent to anyone else. Meanwhile, on Instagram, Sofia, Spain, described a well-thought out process of posting a photo, archiving it, and then unarchiving it, at which point the algorithm would no longer push the photo to people's feeds. She would then wait to see if her crush liked the photo, as that signaled he had purposefully visited her profile, rather than seeing it on his feed. In a slightly less involved process, Jared, UK, described turning his "Ghost mode"

location feature on and off on Snapchat, to generate intrigue about his whereabouts and potentially invite inquiries.

Finally, we also observed many forms of encoding messages online, such as social steganography (Abidin 2021; boyd and Marwick 2011), the practice of hiding messages in plain sight. For example, sharing song lyrics or even overlaying a song onto a post enables teen participants to communicate a message that only a select few in their social media audiences will understand, as it is tied to some kind of shared experience or insider knowledge. This strategy allows participants to ensure their message will be interpreted correctly by the intended individuals, and it also simultaneously enhances their social bonds with those individuals. Playing with meaning was truly a playful and meaningful strategy for so many of our teen participants.

Context Collapse Dimension #4: Conflicting Purposes of Platforms

Throughout our research with teen participants, we were struck by not only their levels of self-awareness when it came to self-presentation, but also their awareness of how different platforms influence their interactions.

Social media apps often seem to have conflicting purposes and norms for their use, resulting in tensions around the sorts of behaviors or interactions that users should engage in. Platforms are simultaneously spaces for sharing personal updates, shopping for products, following celebrities, interacting with official school accounts, and advancing social causes, which introduces uncertainty and complexity as to what they are truly for and how they should be used.

While some scholars have previously examined the ways in which platform norms impact people's social media use (Abidin 2020; boyd 2014; Miller et al. 2016; Szabla and Blommaert 2020; Wexler et al. 2019), there has been less scholarly consideration for how the conflicting purposes of platforms also represent a key dimension of context collapse, particularly for teens navigating today's social media landscape.

Daniel Miller et. al (2016) posited that "we should be careful in presuming that there are properties of the platforms that are responsible for, or in some sense cause, the associations that we observe with platforms", but we observed that many of our teen participants were vocal about the active role of the app architectures themselves. Across markets, participants are conscious of "the algorithm" dictating what they watch and how long they spend on the app. They are cognizant of posts being monitored, accounts being "shadowbanned", and certain content or interactions being celebrated more than others. Teens in our study often feel that platforms expect them to fulfill certain "roles" that they aren't always willing to play:

> I used to want to be an influencer, but then I kind of gave up. I made £800 from TikTok from advertising people's products. But then I got shadowbanned. I was tired anyway with how the platform was continually trying to get me to make content.
>
> — 16, F, UK

This intrinsic awareness of the role the app plays in their interactions can also be viewed in relation to the rise of influencer culture and the attention economy that was peaking as this generation of teenagers arrived on social media. Teen participants are aware that social media is as much of a commercial space (where businesses operate and generate profit) as it is a deeply personal and social space (where meaningful connections with friends, peers, and even themselves occur). The conflation of these two seemingly incongruent purposes is a reality that plays out in each day-to-day interaction for teen participants on these apps: at what point does posting photos for my friends to see become a monetizable interaction? At what point am I scrolling for enjoyment or scrolling because of the algorithm? Maria from Spain described this dilemma and her response this way: "I used to spend so much time looking at influencers so I created a new account just for adding my friends, and none of the influencers."

Academics such as Marwick and boyd pointed out as early as 2010 that spaces involving "micro-celebrity practice" such as Twitter combined "public-facing and interpersonal interaction", the result of which were "new tensions and conflicts" for the networked audience within it. They went on to ask "if 'public' space is becoming synonymous with 'commercial'", a question increasingly at the forefront of teen participants' experiences online today. Participants consistently, and consciously, grapple with the conflict between public and personal, as well as "authentic" interactions and commercialized ones. Arriving on and operating within a social media platform means a collision between these binaries, and study participants expressed growing uncertainty about apps whose commercially-driven "inner-workings" felt difficult to understand.

The Conflicting Purposes of Platforms Leads Teen Participants to Engage Intentionally with Platforms

As teen participants encounter conflicts between personal and commercial purposes of social media, as well as the influence of algorithms on their interactions, they think critically about how to engage with platforms from a place of agency and intentionality.

When study participants spoke about the various platforms they used, they often portrayed apps as "actors" in their interactions. In this sense, they humanize and personify apps, which helps them to rationalize decisions and gain a sense of agency through the promise of a two-way relationship. The "algorithm" is an entity that they can "get to know", an entity with which they can negotiate, influence or shape (Gillespie 2014). It monitors and evaluates their content, and teen participants modify their content to avoid the app's "gaze".

This strategy can be understood as a means by which teen participants attempt to contextualize the intentions of apps they are using. They feel they have to fill in the gaps between what they believe platforms can do (shadowbanning, etc.) and what platforms actually do. In the absence of the level of transparency study participants expect from social media platforms, engaging with an "app" as an "actor" helps them to define a role that feels "authentic" to them: one teen from our study purposefully created seven different accounts on Instagram in order to generate "seven different algorithms" that catered to her diverse interests. Other teen participants use platforms like TikTok through their phone's browser rather than through the app. In doing so, they feel they can still enjoy using the app, without it learning too much about them and capturing more of their attention.

For some study participants, the perceived pressure they feel from apps to behave in a certain way leads them to delete the apps entirely. As one teen from the study recounted, "I got rid of TikTok because I felt there was a lot of pressure to make videos and that made it less fun. That kind of all went away when I deleted it—there wasn't as much pressure anymore."

Teen participants' response to the collapse of platform purposes—particularly the conflict between personal use and commercialized influencer culture—is particularly evidenced by their gravitation away from posting-oriented platforms and more towards messaging and gaming platforms:

> "I don't see any need to be constantly sharing [things on social media]. If I have to share something, I'll send it to other people through messages"
> – Pablo, Spain

For many teens in the study, apps like WhatsApp, Snapchat and Discord, which center on messaging with friends and family, are foundational, while apps which focus on content, like TikTok and Instagram, are more optional.

This reorientation towards smaller circle social interaction plays out hand in hand with the phenomenon of empty social media profiles across many different apps. Conscious that their posts to social media feeds are subject to commercial purposes or unclear algorithmic influence, teen participants opt for interactions that feel

immune to these processes. In some cases, participants were vocal about their outright rejection of public broadcasting, while for others, their strategic shifts in behavior were more subtle. Throughout our study, we witnessed a broad inventory of creative strategies that teen participants employ to mitigate multidimensional context collapse, and even in some instances, use context collapse to their advantage.

Conclusion

This paper has provided a foundational understanding of teen social media use through historical, theoretical, and international ethnographic lenses. While a historical view shows that context collapse has always been implicit in online interactions, our ethnographic research reveals that online context collapse is arguably more prevalent and complex than ever and that teen participants operate by default from a place of high awareness and attunement to this phenomenon.

We have argued for an expanded understanding of multidimensional context collapse that encompasses not only the flattening of audiences and the distortion of time, but also the instability of meaning and the conflicting purposes of platforms. Through our context collapse framework, we have provided a lens that can help future ethnographic practitioners understand young people's experiences with social media more deeply.

The teen participants from our study operate with a mindset of private and ephemeral by default, and public and permanent through effort (only with deliberate strategies). This current paradigm often appears to present as an avoidance of "posting", but does not necessarily imply an avoidance of social media. On the contrary, our research reveals a wealth of generative strategies that participants use to meaningfully engage with social media on their own terms. It is our hope that our framework for context collapse can also aid future practitioners in uncovering the adaptive and innovative behaviors of users who navigate context collapse online.

Finally, in revealing the negotiations between users and context collapse, at times tense, at others playful, we wish to challenge the notion of social media as either a utopian or dystopian technology, two polarizing extremes that are often popularized in public discourse. Our research demonstrates that both users and social media technology are active agents in the shaping of online experiences (Abidin 2021; Davis and Jurgenson 2014; Gillespie 2014; Szabla and Blommaert 2020), and as such, this two-way relationship creates a far more complex reality than may be initially assumed. Social media has the capacity to amplify a range of content and experiences, both positive and negative (boyd 2014), but people also have an

incredible capacity to navigate, strategize, and reclaim these online spaces for meaningful purposes.

As we look forward to the next era of social media, we consider both the past and the present, and ponder, will it continue to be called social media by the next generation of users? Where might their strategies lead, and how will these shape the future of these platforms? One thing is for certain—we must continue to generate new approaches for meeting young people where they're at, taking notice of what they want to share, their challenges, and their responses. We should consider how we can more proactively and thoughtfully design platforms that support what is important to them. We conclude not with a prediction but rather an open question for the community: how can we pay attention to both what is said and also unsaid, done and not done? What are people communicating when they post, but also, what are they conveying when they don't?

About the Authors

Kristin Sarmiento is an applied cultural anthropologist consulting for Stripe Partners. She specializes in using ethnographic approaches to understand the values and needs that underlie human behavior, and applying these insights to inform business strategy, decision making, and product development. Notable projects include foundational and UX research with youth for tech innovation companies. She holds a BA in Anthropology (double minor in Linguistics and Spanish) from Mount Royal University, Canada. kristin.sarmiento@stripepartners.com

Isabelle Cotton is a digital anthropologist and social researcher. Her work centers around using anthropology as a transformative tool of empathy in digital spaces. Before leading on research at Stripe Partners, she worked at a Gen-Z insights agency that platformed minority voices. Especially when working with young audiences, she prioritizes creating a research process that is mutually beneficial for both clients and participants. Isabelle holds an MSc in Digital Anthropology from UCL. isabelle.cotton@stripepartners.com

Josh Terry is a consultant with a background in anthropology and history. He is particularly interested in the social and cultural dimensions of emerging digital spaces and much of his current work focuses on younger generations' relationships with digital platforms. Josh holds an MSc in Social and Cultural Anthropology from UCL. His work and studies have been driven by a fascination in the way societies, cultures and the relationships people form both shape and are shaped by emerging technologies. josh.terry@stripepartners.com

Lea Ventura is part of a team at Instagram that supports the responsible growth of teens on the platform. With over 10 years of experience as a licensed child and adolescent clinical psychologist, Lea has a strong background in mental and behavioral health. She holds a PhD in Clinical Psychology and an undergraduate degree in Cognitive Neuroscience. At Instagram, Lea combines clinical expertise with UX research to promote wellbeing in tech and ethical & supportive product development for younger users.

Shayli Jimenez is a Senior UX Researcher at Meta's Reality Labs, specializing in developing research-based strategies for emerging technologies like GenAI and platforms such as Horizon and the Metaverse. With almost 10 years of experience working across Meta, Shayli excels in identifying priorities and delivering top-tier research insights. She holds a Bachelor's degree in Public Communication and Graphic Design from American University.

Notes

We wish to thank Annie Lambla (Lead User Researcher, Coloplast), Cath Richardson (Partner, Stripe Partners), Simon Roberts (Partner, Stripe Partners), and Kristin Hendrix (VP, Head of Research, Instagram) for their time and thoughtful comments in reviewing the drafts of this paper.

This work does not reflect the views, opinions, or official position of Meta. This paper is not based on any internal log data, quantitative studies that size or show social media behaviors at a population level, or studies about adult users.

1. We recruited a representative mix of SECs, ethnicities and genders from urban settings (cities and suburban regions), as well as a mix of social media user types (different apps, different levels of engagement).

2. All participant names have been pseudonymized in this paper.

References Cited

Abidin, C. 2015. "Communicative ❤ Intimacies: Influencers and Perceived Interconnectedness." *International Journal of Communication* 9:21–39.

Abidin, C. 2016. "'Aren't These Just Young, Rich Women Doing Vain Things Online?': Influencer Selfies as Subversive Frivolity." *Social Media + Society* 2(2). https://doi.org/10.1177/20563051166413

Abidin, C. 2020. "Mapping Internet Celebrity on TikTok: Exploring Attention Economies and Visibility Labours." *Media International Australia* 177(1): 30–45.

Abidin, C. 2021. "From 'Networked Publics' to 'Refracted Publics': A Companion Framework for Researching 'Below the Radar' Studies." *Social Media + Society* 7(2). https://doi.org/10.1177/205630512098445

Asokan, A. 2008. "The Space Between Mine and Ours: Exploring the Subtle Spaces Between the Private and the Shared in India." *Ethnographic Praxis in Industry Conference Proceedings* 1:175–191.

boyd, d. 2008. "Taken Out of Context: American Teen Sociality in Networked Publics." PhD dissertation, University of California, Berkeley.

boyd, d. 2014. *It's Complicated: The Social Lives of Networked Teens*. New Haven, CT: Yale University Press.

boyd, d., and Marwick, A. 2011. "Social Privacy in Networked Publics: Teens' Attitudes, Practices, and Strategies." *Journal of Communication* 61(4): 102–121.

Brandtzaeg, P. B., and Lüders, M. 2018. "Time Collapse in Social Media: Extending the Context Collapse." *Social Media + Society* 4(1).

Brügger, N. 2015. "A Brief History of Facebook as a Media Text: The Development of an Empty Structure." *First Monday* 20(5) https://firstmonday.org/ojs/index.php/fm/article/view/5423.

Bucher, T. 2018. "If…Then: Algorithmic Power and Politics." *New Media & Society* 20(11):4197–4213.

Campbell, S.W. 2013. "Mobile Media and Communication: A New Field, or Just a New Journal?" *Mobile Media & Communication* 1(1): 8–13.

Davis, L., and Jurgenson, N. 2014. "Context Collapse: Theorizing Context Collusions and Collisions." *Information, Communication & Society* 17(4):476–485. https://doi.org/10.1080/1369118X.2014.888458.

Faulkner, S., and Melican, J. 2007. "Getting Noticed, Showing-Off, Being Overheard: Amateurs, Authors and Artists Inventing and Reinventing Themselves in Online." *Ethnographic Praxis in Industry Conference Proceedings* 1:51–65.

Gillespie, T. 2014. "The Relevance of Algorithms." In *Media Technologies: Essays on Communication, Materiality, and Society*, edited by Tarleton Gillespie, Pablo J. Boczkowski, and Kirsten A. Foot, 167–194. MIT Press.

Goffman, Erving. 1959. *The Presentation of Self in Everyday Life*. New York: Doubleday.

Gopffarth, J., Jablonsky, R., and Richardson, C. 2022. "Building Resilient Futures in the Virtual Everyday: Virtual Worlds and the Social Resilience of Teens During COVID-19." *Ethnographic Praxis in Industry Conference Proceedings* 1:117–136.

Highfield, T., and Leaver, T. 2016. "Instagrammatics and Digital Methods: Studying Visual Social Media, from Selfies and GIFs to Memes and Emoji." *Communication Research and Practice* 2(1):47–62.

Hogan, B. 2010. "The Presentation of Self in the Age of Social Media: Distinguishing Performances and Exhibitions Online." *Bulletin of Science, Technology & Society* 30(6):377–386.

Horst, H.A., and Miller, D. 2012. *Digital Anthropology*. London: Berg.

Ito, M., Baumer, S. Bittanti, M., boyd, d., Cody, R., Herr-Stephenson, B., Horst, H. A., Lange, P. G., Mahendran, D., Martinez, K. Z., Pascoe, C.J., Perkel, D., Robinson, L., Sims, C., Tripp, L. 2010. *Hanging Out, Messing Around, and Geeking Out: Kids Living and Learning with New Media*. Cambridge, MA: MIT Press.

Kilford, E., Garrett, E., and Blakemore, S-J. 2016. "The Development of Social Cognition in Adolescence: An Integrated Perspective." *Neuroscience and Biobehavioral Reviews* 70:106–120. https://doi.org/10.1016/j.neubiorev.2016.08.016

Klug, D., Evans, M., and Kaufman, G. 2023. "How TikTok Served as a Platform for Young People to Share and Cope with Lived COVID-19 Experiences." *MedieKultur: Journal of Media and Communication Research* 38:152–170. https://doi.org/10.7146/mk.v38i73.128463

Leaver, T., Highfield, T., and Abidin, C. 2020. *Instagram: Visual Social Media Cultures*. Cambridge: Polity.

Ling, R. 2012. *Taken for Grantedness: The Embedding of Mobile Communication into Society*. Cambridge, MA: MIT Press.

Madden, M., Lenhart, A., Duggan, M., and Smith, A. 2013. "Teens, Social Media, and Privacy." Pew Research Center, May 21 https://www.pewresearch.org/internet/2013/05/21/teens-social-media-and-privacy/.

March, W., and Fleuriot, C. 2005. "The Worst Technology for Girls?" *Ethnographic Praxis in Industry Conference Proceedings* 1:165–172.

Marwick, A. 2015. "Instafame: Luxury Selfies in the Attention Economy." In *The Drama of Social Life*, edited by T. Merrin, 107–126. Routledge.

Marwick, A., and boyd, d. 2010. "I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience." *New Media & Society* 20:1–20. https://doi.org/10.1177/1461444810365313

McNelles, L.R., and Connolly, J.A. 1999. "Intimacy between adolescent friends: Age and gender differences in intimate affect and intimate behaviors." *Journal of Research on Adolescence* 9(2):143–159. https://doi.org/10.1207/s15327795jra0902_2

Miller, D., Costa, E., Haynes, N., McDonal, T., Nicolescu, R., Sinanan, J., Spyer, J., Venkatraman, S., Wang, X. 2016. *How the World Changed Social Media*. London: UCL Press.

Rangaswamy, N., and Yamsani, S. 2011. "'Mental Kartha Hai' or 'Its Blowing My Mind': Evolution of the Mobile Internet in an Indian Slum." *Ethnographic Praxis in Industry Conference Proceedings* 2011:285–298.

Schulz, C. 2023. "A New Algorithmic Imaginary." *Media, Culture & Society* 45(3):646–655. https://doi.org/10.1177/01634437221136014

Senft, T. M. 2013. "Microcelebrity and the Branded Self." In *A Companion to New Media Dynamics*, edited by John Hartley, Jean Burgess, and Axel Bruns. Oxford: Blackwell Publishing Ltd. https://doi.org/10.1002/9781118321607.ch22

Szabla, M., and Blommaert, J. 2020. "Does context really collapse in social media interaction?" *Applied Linguistics Review* 11(2):251–279. https://doi.org/10.1515/applirev-2017-0119

Wei, E. 2021. "American Idle." Accessed July 17, 2024. https://www.eugenewei.com/blog/american-idle

Wexler, M., Yunzhijun, Y., and Bridson, S. 2019. "Putting Context Collapse in Context." *Journal of Ideology* 40(1):3.

Wesch, M. 2009. "Youtube and You: Experiences of self-awareness in the context collapse of the recording webcam." *Explorations in Media Ecology*, 8(2):19-34.

Collaborative Generation: Making Sense of Today's Stories to Catalyze a Thriving Tomorrow for All

JENNY FLINDERS, *University of Kansas Center for Public Partnerships and Research*
JACQUELINE COUNTS, *University of Kansas Center for Public Partnerships and Research*

This article shares an innovative, community-centered method for research, sensemaking, and innovation in social services that has been used by communities and agencies in six states. 'Our Tomorrows' is a framework for capturing real-time signals while stimulating community-based solutions to complex problems. The method includes three components that establish an ongoing feedback loop between citizens and decision makers: intensive narrative data collection with web-based tools, community sensemaking workshops, and community action labs. This article covers the principles and methodology of the approach, then walks through an example of its application and outcomes in an early childhood program. Since its inception, Our Tomorrows has leveraged distributed ethnography data tools to listen to over 10,000 voices from across the country and put the everyday experiences of individuals and families at the center of local programs, policies, and systems.

Introduction

The Center for Public Partnerships is housed within the Achievement and Assessment Institute at the University of Kansas. We are a multi-disciplinary staff of over 100 researchers, data scientists, strategic planners, and communications experts serving the broad mission of optimizing the well-being of children, youth, and families. We partner with state agencies, non-profit organizations, and community organizations to provide strategic leadership and support, research and evaluation, and data management services.

As change-makers and systems-builders, we recognize that our work is built on the foundation of daily lived experience. In 2019, we launched the Our Tomorrows project to elevate the nuance and texture of daily interactions with systems and each other and mobilize local voices as a human sensor network. Rooted in complexity theory (Snowden and Boone 2007), Our Tomorrows seeks to capture real-time signals of changing dynamics in the ways families get the social supports they need while stimulating community-based solutions to "wicked problems" (Zivkovic 2021) at all levels of our social systems. Since its inception, the Our Tomorrows project has leveraged distributed ethnography (Joshi et al. 2023) data tools (such as SenseMaker and Sprying.io) to listen to nearly 10,000 voices from across the country, putting the

Flinders, Jenny, and Jacqueline Counts. "Collaborative Generation: Making Sense of Today's Stories to Catalyze a Thriving Tomorrow for All." *EPIC Proceedings* (2024): 364–376. https://www.epicpeople.org/collaborative-generation. ISSN 1559-8918.

everyday, real experiences of individuals and families at the center of local programs, policies, and systems.

Through this important work we have witnessed real ground up change in practice. By leveraging the active sensemaking approach and distributive ethnographic practices, participants in each of the Our Tomorrows frameworks (constructed with rigorous community input) share their stories of thriving or just surviving and then make sense of those experiences in their own words and ways, gathering rich data, both on participants' experiences and, crucially, what those experiences mean to them. Patterns from these responses and their associated stories are then returned to communities through collaborative sensemaking sessions where community participants deepen collective understanding of experiential patterns, helping generate solutions and come together over the ancestral practice of shared storytelling. In the almost ten years that we have been practicing the active sensemaking approach, we have seen individuals, families, community leaders, funders, and state representatives come together at the change making table to uncover the "unknown unknowns" and create real solutions that address the things that most impact the daily lives of citizens. Whether recognizing the cost of living in a pandemic without access to technology or reliable Internet, the barriers of providing licensed in-home childcare as a home renter or identifying the ways that current programmatic requirements exclude historically oppressed and vulnerable families, our community sensemaking approach has agitated complacent systems and uncovered needed changes that can be addressed today.

Context

We are located in Kansas and have strong partnerships and do community-based work across the state. Kansas is the second most rural state in the country. In our rural and frontier areas, we have communities struggling with aging populations, out-migration to urban areas, increasing poverty, and disinvestment. Our urban areas face ongoing challenges with housing costs and food insecurity (All In For Kansas Kids 2024). The political environment is traditionally conservative and skeptical of social services, but there is also a history of leadership on children's issues. Kansas was the first state in the nation to establish a Children's Cabinet in 1980, and one of few states to devote Tobacco Master Settlement dollars to services to children through the establishment of the Children's Initiatives Fund in 1999.

We have been working as a university partner providing evaluation and strategic support to social service organizations and state agencies for 20 years. We believe strongly in the importance of good data and empirical investigation to support

thoughtful decision-making, and in the urgent need to make transformational systems change to support thriving families and communities. We also have a front seat to the dynamics that prevent standard evaluation practices and contexts from supporting the radical change our current circumstances demand. Survey fatigue is real. Vulnerable populations are weary of extractive techniques without feedback loops. Individuals struggle to envision innovations that are vastly different from the status quo. Imagination to change the system is often low. It takes time to develop trust and foster imagination that leads to ideation and action, and most evaluation efforts do not make space for that time or that kind of thinking. Abductive reasoning is not common in applied settings in which conclusive reports are the norm. Provocative questions make partners uncomfortable and are outside the expectations and experience of public servants. Incorporating liberating structures and futures thinking into sensemaking texts is essential to change participant mindset to a more generative state that focuses on a markedly different future.

We are always on the lookout for opportunities to address these tensions and try new strategies to support systems change, frequently drawing on the tools provided by complexity theory and foresight methods. We began experimenting with sensemaking tools with the goal of bridging the gap between abstract understandings of how to effect change and our commitment to the knowledge and wisdom of individuals and communities and locally driven solutions.

The first project in 2015 was Lemonade for Life, a training program for home visitors and other family support professionals to understand what Adverse Childhood Experiences (ACEs) were; to learn about their own ACES and how these might be affecting their work; and to talk about ACEs with families they serve. ACEs are stressful or traumatic events that occur before a person is 18 years of age and include abuse, neglect, household dysfunction., exposure to violence, substance use, disorder and others (Centers for Disease Control 2021). We searched for innovative evaluation techniques other than typically used quantitative surveys or qualitative interviews to better understand how individuals' own ACEs and their understanding of them interacted with how they were building relationships with families. Several web searches and discussions led us to the Cyenfin Framework (Snowden 2007) and SenseMaker understand complex social problems. We piloted the tool with several early cohorts participating in the early phases of Lemonade for Life and found the method to be promising to understand nuances during sensemaking sessions.

In 2019, the State of Kansas received a Preschool Development Grant to conduct a needs assessment and craft a strategic plan for the early childhood system.

The goal of the project was to better align across the system by putting the needs and experiences of young children and their families at the center. This focus demanded a new approach to identifying and understanding the lived experience of families with young children as they navigate their daily lives and the systems intended to serve them. Responding to this need, the grant leadership team of state agencies decided to harness the power of Our Tomorrows' innovative Community Sensemaking Approach to map families' lived experiences and create policies and programming adaptive to families' needs. From a complexity perspective, the overarching goal was to developing a 'human sensor network,' embedding citizen feedback loops and sensemaking processes into governance, and complexity-informed intervention via portfolios of safe-to-fail probes. To do this, the Our Tomorrows project rolled out in three phases:

1. An intensive story collection effort to gather experiences from families across Kansas
2. Community Sensemaking Workshops to make sense of the stories and the patterns that emerged
3. Community Action Labs to test out small-scale innovations designed to address identified needs of families

Methods

The Our Tomorrows Project

Our Tomorrows is a narrative-based data collection project. We primarily utilize data platforms like SenseMaker and Sprying.io which can support collection of multifaceted data in the form of open text boxes, triads, sliders, and traditional multiple-choice formats. We have also developed paper versions of the framework to facilitate story collection during in-person group settings.

Frameworks are built around open-ended prompts inviting respondents to share a short story relevant to a topic or theme. One such prompt that Our Tomorrows has fielded for many years is: "Remember a time when you felt like your family or another family you know was thriving, or just surviving. Share an experience that describes what was happening at that time in the family." Respondents are then given a series of triads which allow them to characterize the story they just shared. Examples:

- "In the experience you shared, views on raising a family were shaped by: a) Personal experience b) Friends and family c) Tradition";

- "In the experience you shared, who should have been responsible for making sure that kids thrived? a) The family b) Friends and community c) government."

Respondents can code their story anywhere within a triangle, allowing them to indicate which of the three phrases apply and to what extent each is a factor, depending on the placement. Frameworks can also include more traditional survey items, providing both qualitative and quantitative data for each observation.

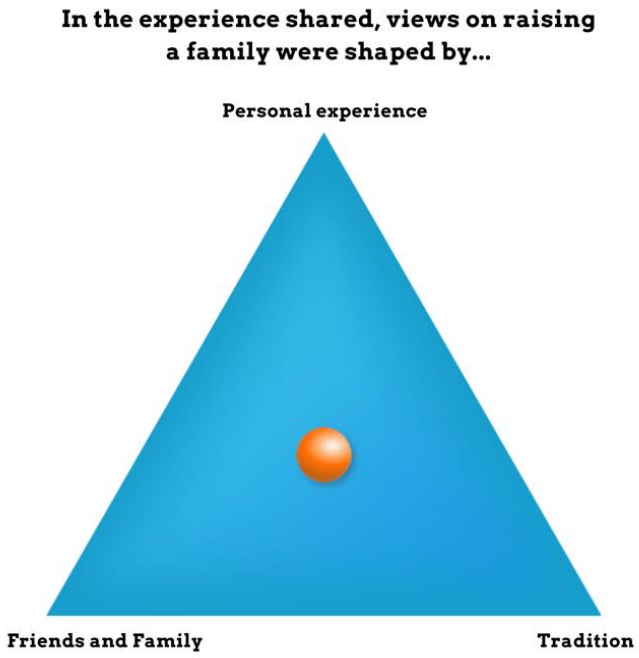

Image 1. Triad question from Our Tomorrows sensemaking survey, 2018. Credit: University of Kansas Center for Public Partnerships and Research.

Story Collection

State leaders, including a new Governor and leadership of the Kansas Children's Cabinet and Trust Fund, Kansas Department of Health and Environment, Kansas Department for Children and Families, and Kansas State Department of Education, aligned to call for statewide participation in the Our Tomorrows project. With this backing, the Our Tomorrows project created a network of story collectors that contributed to a statewide story bank that decision-makers could draw on to better understand the lived experience of Kansans. Citizen Journalist Training provided organizations an understanding of the story collection method, and the Our

Tomorrows team co-developed story collection strategies with each partner. Families submitted their stories online, on paper, and through interviews conducted by Citizen Journalists. This ambitious effort collected 2,279 stories from all 105 Kansas counties, capturing experiences of families in frontier, rural, and urban communities.

Analysis and Community Sensemaking

A core principle of the Our Tomorrows project is the importance of returning data to the people that shared it. Our goal is to establish a feedback loop between story sharing, sensemaking, and resulting actions. After our initial period of intensive story collection, we returned anonymous stories were then returned to communities, families, service providers, and other early childhood stakeholders at fourteen Community Sensemaking Workshops. This phase engages community members and partner organizations to uncover insights with the explicit goal of taking action.

To prepare for community sensemaking, the Our Tomorrows team conducted an initial review and analysis of the dataset. We developed the Sensemaking Analysis and Visualization Dashboard (SAVVY), which helped us identify patterns emerging from families' experiences. Quantitative data gathered through series of triangles, sliders, a resource canvas, and multiple-choice questions provided self-signification to each narrative. When aggregated, this self-signification uncovered patterns around four themes: Bright Spots, Disruptors, Unmet Needs, and Unheard Voices.

The Our Tomorrows team developed Community Sensemaking Guides based on patterns in the data for each of the six regions in Kansas. The Guides were intended to orient participants to the project and their community's data to enable deep conversations and prompt ideas for action. They included a description of Kansas Early Childhood Systems Building, an interpretation guide, summary data describing respondents and narratives, and initial patterns in the four thematic areas. We also compiled Story Packs for each region with groups of stories that share a similar signification in the framework. For example, a Story Pack could be all stories from a particular region where individuals selected the "Thoughtful Planning" section of a triad when characterizing how decisions were made.

The two-hour Community Sensemaking workshops were built around interpreting the Story Packs. Participants reviewed the Story Packs and sorted them into thematic categories rooted in the ways respondents interpreted their own stories, identifying emerging patterns in the conditions under which families thrive. Discussion focused on surprises, expectations, and ways to take action.

Community Action

In the final phase of the project, we launched Community Action Labs to support local initiatives that were quick, local, and inexpensive (up to $2,000). Community Action Labs created a structure for communities to test innovative ideas to address family needs as uncovered by the Our Tomorrows process, encouraging adaptive, localized strategies and a culture of experimentation to make progress on shared goals. In this spirit, we established the following principles to guide Community Action Labs:

- Big ideas start with small ripples.
- Anyone can take action and make a difference.
- Stories and families' experiences fuel action.
- Locals know best.
- There are many paths to our shared destination.

The Our Tomorrows innovation was inspired by a partnership with the Observatory of Public Sector Innovations (OPSI) Anticipatory Innovation Governance Program and the Cynefin Centre for Applied Complexity. Our Tomorrows consulted with OPSI and the Institute for the Future on developing Community Action Labs to incorporate Facets of Innovation and futures and foresight methodologies. As a result, each Community Action Lab Actionable application was categorized along the Facets to provide insight on the disposition to innovation across the state. Our Tomorrows has laid groundwork to introduce anticipatory innovation to state decision-makers while providing avenues at the community level for immediate participation. The Cynefin Centre for Applied Complexity consists of a network of SenseMaker practitioners that have provided valuable guidance on story collection management and sensemaking workshop facilitation.

Impact

The Our Tomorrows Community Sensemaking Approach is an innovative application of complexity-informed methods toward citizen engagement in four ways:

1. It is the first instance of an ongoing sensemaking feedback loop between citizens and decision-makers across an entire early childhood system.
2. Every person is empowered to act according to their skillset and level of authority by asking themselves, "What can I do tomorrow to create more

stories like the ones I want to see and fewer like the ones I don't?" This "fractal engagement" puts problem-solving power in the hands of communities, not just high-level decision-makers.

3. Sensemaking data is returned to communities for analysis and action planning in a comprehensible and accessible way through Community Sensemaking Workshops.
4. Community Action Labs crowdsourced a portfolio of safe-to-fail experiments for complexity-informed intervention strategy through small grants.

Measuring impact from sensemaking work is challenging given the emphasis on local level community-based conversations and solutions. Thanks to the Our Tomorrows development of direct-to-community Action Labs and facilitation of community sensemaking sessions, the Center for Public Partnerships and Research has observed and documented several key areas of impact from the evolution of its active sensemaking work over the past eight years. These critical impacts include the use of sensemaking data for systems alignment and policy reform, investment in community organizations and members for safe-to-fail solutions, an increase in diverse community member participation and youth engagement, and active participation and support from state government agencies and policymakers.

The Our Tomorrows Story Bank provides a de-politicized lens for discussions about core issues that often devolve into partisan debates, like health care. By framing dialogue with stories of thriving or surviving, people across the political spectrum can think about problems from the perspective of families. Then, they can think about what they have the capacity to change. This work has prompted Kansas state agencies and early childhood stakeholders to use SenseMaker data for systems alignment, workforce development, adaptive program management, and building political will for systemic reforms. Specifically, five (5) state agencies with high-level decision-makers that are interested in complexity-informed intervention strategies, innovation, and futures methodologies.

Community members across the state participated in Community Action Labs to test innovative ideas developed through the sensemaking process, with investment in these micro-grants from the Kansas Children's Cabinet and Trust Fund. In total, forty-six (46) individuals or organizations proposed local solutions for Community Action Labs, doubling the expected response. The Labs allowed them to safely take a risk on new ideas without jeopardizing pre-existing funding, relationships, or organizational boundaries. Overall, all stakeholders have learned to apply complexity principles and embed SenseMaker into their day-to-day operations along the way.

Finally, Our Tomorrows resulted in youth engagement and new dialogue on deep cultural issues. One citizen journalist was a 13-year-old who went door-to-door asking people to "make their community a better place" by sharing a story. Upon hearing of this effort, a state legislator unexpectedly and emotionally shared the youth's story at a state meeting. This was a pivotal moment that led to an increased commitment from state leadership to center family experiences to inform decision-making. The youth was then invited to join a panel and share his hopes for his community and has been an inspiration for others across the state. Whether recognizing the cost of living in a pandemic without access to technology or reliable Internet, the barriers of providing licensed in-home childcare as a home renter or identifying the ways that current programmatic requirements exclude historically oppressed and vulnerable families, our community sensemaking approach has agitated complacent systems and uncovered needed changes that can be addressed today.

Lessons Learned

Challenges and Fails

The tension between meeting community partners 'where they were' and adopting new methods for community engagement styles was a constant challenge. Although there was universal interest in trying something new, people were unsure how to begin or were stuck in old ways of working. To address this problem, Our Tomorrows pursued the 'adjacent possible' by breaking down big ideas into manageable steps. Emerging goals of state leadership, feedback from community partners, and technical infrastructure challenges required abrupt pivots and creative solutions at scale without time for testing. Our Tomorrows communicated vision, principles, and introduced new vocabulary to maintain coherence and provide stability amidst this uncertainty.

Conditions for Success

Open-minded leadership and adequate infrastructure for grassroots participation were the most important conditions for success. The support of the Governor's Office and state agency leaders resulted in a statewide commitment to the SenseMaking process that spread to elected officials, state boards, advisory groups, and advocacy organizations. With this support from the top, Our Tomorrows began an intensive partner on-boarding process to build local capacity for story collection,

sensemaking, and Community Action Labs. The strong relationship with local partners created a bottom-up demand for the Community Sensemaking Approach that increased leadership's investment in the innovation. This dialectic introduced the trust and stability to the process needed for sustainable change.

Replication

The Community Sensemaking Approach can be replicated by organizations, agencies, or governments that seek to use citizens' experiences to drive complexity-informed change. With appropriate capacity and onboarding, 'sensemaking' organizations can adopt the SenseMaker tools, data visualization infrastructure, and strategy developed by Our Tomorrows to bolster community listening and social innovation. Our Tomorrows partners are replicating the approach locally by integrating community feedback loops into their day-to-day organizational practices. We have discussed a direct replication of Our Tomorrows in other states that have received federal grants to strengthen their early childhood systems. We are also exploring a social innovation platform collaboration with the Agirre Lehendakaria Center in the Basque Country (Spain) and sharing our approach with the members of the Cynefin Centre for Applied Complexity.

Lessons Learned

Implementing the Community Sensemaking Approach requires that practitioners play a leadership role to get others to join in a shared struggle to solve a complex problem. Lessons learned were:

1) People need to understand the why, how, and what of a process to feel secure enough to take an innovative risk. "Breadcrumbing" is an approach we developed to educate partners about innovative ways of doing things without overwhelming them with jargon and academic language. We introduced the 'Why' of the Community Sensemaking Approach. Then participants experienced the 'How' by completing program activities. Through guided reflection afterwards, we provided the language of the innovation to describe the 'What.' This staged process introduces complexity concepts in a consumable and respectful manner.

2) Survey fatigue is real. Vulnerable populations are weary of extractive techniques without feedback loops. Integrating story collection into routine activities and going to where people are is essential. Additionally, communication strategies must be adjusted based on the audience. The statewide project required that we use top-down and bottom-up approaches to establish feedback loops. In response, we developed "fractal knowledge management" techniques to share the same ideas in a

variety of ways to provide coherence across the system while not overwhelming people who had less shared context.

3) The project team must use complexity techniques to deliver the project and be an exemplar for others. For example, Our Tomorrows utilized the Cynefin framework for situational assessment and as a guide to adjust our practices accordingly. We began the Community Action Lab process with a long application like a request for proposals. After some confusion from our partners, we recognized that we were approaching the application as a 'clear' problem rather than a complex one. We adjusted our approach to reflect the heuristics for action in the complex domain and created a three-question application to probe for unexpected ideas. By loosening constraints, the Labs achieved greater engagement. In the end, this resulted in a locally-driven innovation portfolio that was an iterative process built on trust and supportive coaching. Individuals struggle to envision innovation that are vastly different from the status quo. Imagination and change the system is often low. It takes time to develop trust and foster imagination that leads to ideation and action.

5) Abductive reasoning is not common in applied settings in which conclusive reports are the norm. Provocative questions make partners uncomfortable and are outside the expectations and experience of public servants. Funders are used to compliance models and to receiving reports with a set of (often expected) recommendations. Incorporating liberating structures and futures thinking into sensemaking texts needs is essential to change participant mindset to a more generative state that focuses on a markedly different future.

How the Project Has Evolved

Over the past six years, the Our Tomorrows team at the Center for Public Partnerships and Research at the University of Kansas has co-created and tested an adapted framework focused on child and family wellbeing instead of the more traditional "child welfare" approach to prevention work. This framework, which started with the Our Tomorrows 2.0 framework in Kansas, has been workshopped in varying ways by communities and agency groups in Washington, San Diego, Minnesota, and in currently ongoing ways by additional Kansas child wellbeing programs in Kansas. Every additional community and state partner that utilized this framework had the opportunity to review, test, and adapt it to best meet the unique local needs. This iterative approach to framework design across communities, identities, and levels of government resulted in a socially validated survey framework that is helping leaders at all levels of the child and family wellbeing system better serve all families in communities across the country and is helping to center the

voices of "lived experience experts" in policy, funding, and programmatic decision-making.

Additionally, recent iterations of the active sensemaking model have led the way on the progression of action labs from micro-funded community grants to more policy-based action labs. Facilitated by the state of Minnesota's MN StoryCollective project, policy-centered action labs bring together staff and leaders from multiple state and local agencies to discuss Minnesota's sensemaking data and stories in order to convert their "a-ha" moments into quick action. By ensuring that action labs include the active participation of those who can make quick decisions, access funding, and change policies, Minnesota's action lab model is using the active sensemaking approach to connect community issues to state level problem solvers for efficient, effective change.

Our active sensemaking team has recently engaged in adaptive and emergent design principles to temporally expand its narrative prompting. More specifically, a current project in partnership with another KU research center team is seeking narrative responses about both past youthful experiences and current adult situations, drawing connections between the things we learn as children to the outcomes we may experience in our careers as adults.

The CPPR team is exploring narrative prompting for the creation of future scenarios or stories—asking people to imagine possible futures in the year 2040. This approach is generating wider and more creative possibilities for our future and helping us understand the different ways people perceive of foresight planning, consider future possibilities, and balance the needs of today with the hopes of tomorrow. These innovative approaches to narrative prompting across generations and temporal states are indicative of a sea change in how we communicate with each other, understand our realities, and utilize core ethnographic and complex methods to make sense of what has happened, what is needed, and what is possible.

About the Authors

Jenny Flinders, MSed, leads research projects at the University of Kansas Center for Public Partnerships and Research, that center community voice in the policymaking process by leveraging complex sensemaking processes. She also leads indigenous community partnerships that seek to strengthen intergenerational wellbeing. She is a PhD candidate in Education Policy, currently researching how family choice discourse is leveraged for control across the education system. Contact at jflinders@ku.edu.

Jacqueline Counts, MSW, PhD., is the Director University of Kansas Center for Public Partnerships and Research. She leads a multi-disciplinary team of social workers, social scientists, psychologists,

sociologists, and other creative staff to come together to create a better world. She has been principal investigator on over 50 grants. She uses a complexity lens to understand wicked social problems and applies futures and foresight methods to transform systems. Contact at jcounts@ku.edu.

Note

We are grateful to Dr. Jessica Sprague-Jones for her exceptional writing support, which greatly enhanced the clarity and coherence of our case study. Her expertise and dedication have been instrumental in shaping the dissemination of our work.

References Cited

All In for Kansas Kids. 2024. Kansas early childhood system needs assessment. Prepared by the University of Kansas Center for Public Partnerships and Research on behalf of the Kansas Children's Cabinet and Trust Fund. https://kschildrenscabinet.org/wp-content/uploads/2024/02/Needs-Assessment-2024.pdf

Centers for Disease Control and Prevention. 2021. Adverse Childhood Experiences (ACEs). Retrieved from https://www.cdc.gov/violenceprevention/aces/index.html

Joshi, Deepa, Anna Panagiotou, Meera Bisht, Upandha Udalagama, and Alexandra Schindler. 2023. "Digital Ethnography? Our Experiences in the Use of Sensemaker for Understanding Gendered Climate Vulnerabilities amongst Marginalized Agrarian Communities." *Sustainability* 15, no. 9: 7196. https://doi.org/10.3390/su15097196

Snowden, David J., and M. E. Boone. 2007. "A Leader's Framework for Decision Making." *Harvard Business Review* 85, no. 11: 68. https://doi.org/10.3390/su15097196

Zivkovic, Sharon. 2021. "A Practitioner Tool for Developing and Measuring the Results of Interventions." *Journal of Strategic Innovation and Sustainability* 16, no. 1. https://doi.org/10.33423/jsis.v16i1.4183

Foundations of Practice: An Ethnographic Exploration of Software Documentation and its Cultural Implications for an Agile Startup

NICOLE TIETZ-SOKOLSKAYA, Principal Engineer, Remesh
SUZANNE WALSH, Research Consultant, Remesh

This case study offers an in-depth analysis of the foundational role of documentation for software quality, efficiency, and organizational success. It includes an inside look how a start-up overcame inefficiencies and scaling challenges by implementing a sociotechnical documentation practice, and 10 recommendations for implementation or assessment based the case study data and a wider analysis of research and expertise. Time pressures of agile workflows that prioritize writing code, coupled with rapidly scaling software and startups, create an environment where documentation often takes a back seat. But its value to organizations and engineering teams as a source of knowledge and a practice of knowledge sharing becomes clear when critical inefficiencies and even chaos to prompt an overhaul of norms and behaviors. The experience at Remesh demonstrates the value of an ethnographic approach to gain a holistic understanding of engineering and work practices and to implement positive change.

Transforming Culture: Remesh's Journey to Comprehensive Software Documentation Practices

Remesh, a SaaS research platform, has undergone a significant cultural transformation through the adoption of comprehensive software documentation practices since the pandemic. This shift has not only enhanced operational efficiency but also fostered collaboration, inclusivity, and technical excellence within the organization. This case study delves into the personal and organizational values that have emerged from interviews with Remesh engineers, underscoring the profound impact that thorough documentation has had on the company's culture. By emphasizing formality, collaboration, effective onboarding, tailored documentation, automation, and a cultural recognition of its importance, Remesh's journey towards better documentation practices highlights a strategic approach to fostering a supportive and efficient work environment.

Remesh was founded in 2014 by Andrew Konya, current Chief Science Officer. Konya realized during his PhD program the significant gap in effectively

Tietz-Sokolskaya, Nicole, and Suzanne Walsh. "Foundations of Practice: An Ethnographic Exploration of Software Documentation and its Cultural Implications for an Agile Startup." *EPIC Proceedings* (2024): 377–395. https://www.epicpeople.org/foundations-of-practice-ethnographic-software-documentation. ISSN 1559-8918.

representing the collective voice of people. The initial concept was to develop a platform that could facilitate many-to-one communication, enabling large groups to engage in meaningful dialogue and represent collective responses. The company began in earnest when it joined the 2014 FlashStarts accelerator program in Cleveland, Ohio, which provided the necessary resources to build a prototype. The company secured further growth and funding through its participation in the 2015 Techstars Barclays accelerator program, where Gary Ellis, current CEO, joined the organization as a co-founder.

The platform is designed to facilitate real-time qualitative research and insights gathering, leveraging artificial intelligence and machine learning to engage with large groups of participants simultaneously. This enables organizations to gain deeper and more actionable insights from their target audience. Key features include real-time and asynchronous engagement, AI-driven analysis, scalability, versatility and a user-friendly interface. Overall, Remesh offers a powerful solution for market researchers and organizations looking to conduct in-depth qualitative research at scale, offering the ability to gather rich, real-time insights from large groups of participants with the speed of AI.

Building Structured Documentation Practices

Prior to 2019, Remesh's engineering team consisted of three members, Engineers A, B, and C, each operating within their specialized areas. The lack of formal documentation was a significant challenge, particularly after Engineer B rewrote the entire backend without any accompanying documentation. This led to a situation where Engineers A and B were indispensable for any changes or understanding of the system, as knowledge transfer relied solely on direct interactions. The absence of documented current and future system workings left Engineer C and others dependent on informal, inefficient knowledge-sharing methods, creating a bottleneck and hindering productivity.

Recognizing the need for inclusivity and organizational efficiency, Engineer A championed a guiding ethos: all team members should have access to system knowledge, both current and future. This initiative was driven by the realization that undocumented knowledge was unsustainable, especially as the team grew and individuals took time off or left the company.

As Remesh's team expanded and transitioned to remote work during the pandemic, the necessity for formal documentation became apparent. Engineer A

spearheaded efforts to enhance documentation practices, emphasizing their importance due to the challenges posed by remote work. By early 2021, after securing operational support and experiencing the drawbacks of inadequate documentation, the team adopted more structured processes, including design reviews and the use of collaborative tools like Google Drive and Notion. This shift ensured that knowledge was not only documented but also easily accessible and shareable among all team members.

In the early stages, Remesh's documentation was sparse and informal, leading to inefficiencies and repeated efforts. Essential information was often kept in tools like Slack, complicating the process of referencing or explaining code behavior.

> I do maintain a lot of documentation in Slack, but it isn't in the official place where it should be. I remember when we first started, we were told to add anything we found wrong or unclear, but as a newcomer, that was quite challenging.
>
> —Engineer E1

Outdated commands and incomplete records further hindered newcomers' ability to navigate the system and understand past decisions.

The implementation of structured documentation processes marked a turning point. Engineers began systematically documenting design reviews and testing protocols using tools like Notion and TestRail, making documentation accessible to all team members and providing clear and comprehensive guidelines. The value of comprehensive "how-to" guides and detailed pull request descriptions became more accepted, which facilitated knowledge transfer and aided in code reviews. By focusing on both the "why" and the "how" behind decisions, documentation ensures that even non-technical team members can understand and contribute to basic operations, maintaining clarity and supporting effective collaboration across departments.

> I chose to use Notion for documentation, considering that the audience was not engineers. My approach included two components – I focused first on explaining why I made certain decisions and how things worked, providing links for future operations.
>
> —Engineer E2

Remesh's adoption of formal design review processes further solidified their documentation practices. Engineers systematically documented high-level design

decisions, architectural choices, and technical approaches, fostering consistency and clarity in problem-solving. The transition from co-located to geographically distributed teams during the global pandemic necessitated new documentation processes to maintain efficiency. Collaborative tools like Google Drive and Notion became integral, allowing for real-time feedback and iteration. Additionally, GitHub is used for documentation to align with coding processes, demonstrating Remesh's commitment to robust documentation practices that support growth and operational efficiency.

The team's journey involved experimenting with various documentation processes and iterating based on feedback and observed pain points. The current documentation practices are well-regarded within the team, with most members satisfied with the outcomes.

> At our company, documentation is highly valued. We strive to make our code readable and our documentation clear, reflecting our care for each other's work. This practice facilitates easier replication and modification of work, fostering a supportive and efficient work environment.
>
> —Engineer E3

The engineering team and leadership recognize the importance of a robust documentation process to support the company's growth and operational efficiency, reflecting a well-evolved practice that meets the organization's needs.

Software Documentation: A Scan of the Literature

The goal of software engineering is to produce working software that solves real-world problems. Many activities are required to achieve this goal, with writing code being the most obvious. However, supporting activities like creating automated tests and writing documentation are crucial for project success by decreasing the rate of bugs, increasing confidence in planning, and lowering the probability of schedule delays. Software documentation plays a critical role in the software development lifecycle, serving as a key medium for communication, maintenance, and knowledge sharing.

Documentation can be broadly classified into technical and non-technical categories. Technical documentation includes requirement specifications, architectural design documents, source code comments, test documents, and process descriptions, all of which aim to help software engineers comprehend and effectively

perform their tasks (Garousi et al., 2015). However, determining the appropriate amount and depth of documentation remains a significant challenge. Many development teams either produce too much or too little documentation, leading to incomplete, inconsistent, or outdated documents (Garousi et al., 2015). This issue is compounded by the perception that documentation is an expensive and often undervalued activity (Garousi et al., 2015). Aghajani et al. (2019) highlights the significant gap between the intended and actual use of documentation, suggesting the necessity for more integrated tools that reduce the overhead associated with documentation tasks. They emphasize the importance of embedding these tools within the development environment to streamline the process and ensure that documentation stays up-to-date and useful. Building on this, Aghajani et al. (2020) reveals that developers generally undervalue documentation, viewing it as a secondary task due to the effort required for its maintenance.

Chomal (2015) argues that clear and comprehensive documentation can mitigate risks and enhance project outcomes by providing clear guidelines and reference points. However, Chomal also notes the common tendency for documentation to become outdated quickly, especially in agile environments where iterative development cycles are prevalent. This presents a significant challenge for maintaining the relevance and accuracy of documentation over time. Das et al. (2007) emphasizes the importance of documentation in the maintenance phase of software development. They argue that thorough documentation can significantly ease maintenance by providing detailed insights into the system's design and implementation. Regular updates to documentation are crucial to ensure it remains useful and relevant, highlighting the ongoing effort required to keep documentation aligned with the current state of the software.

De Souza and colleagues (2005) explore the types of documentation essential for effective software maintenance. They argue that concise and accurate documentation is crucial for facilitating maintenance activities and point out that involving all stakeholders in the documentation process can enhance its quality and relevance. This inclusive approach ensures that documentation addresses the needs of various users and maintains its utility throughout the software's lifecycle. Importantly, agile workflows, created in response to the shortcomings of documentation-driven workflows (Tordup Heeagar, 2012), do not easily incorporate documentation practices as the goal of agile methods is to increase the quality of software through flexibility – and the 'proof' of quality is software that works as intended. Document-

driven methods, such as Waterfall, have a similar aim but increase quality through the use of stringent documentation and a lengthy design process. Standards, such as ISO, and maturity models, such as CMMI, are both document-driven and leverage defined documentation as the 'proof' of quality. Neither method is perfect, and both come with strengths and weaknesses. Agile is noted for its innovation and change-embracing flow but can come off the wheels in highly complex environments. Documentation-driven methods are excellent for highly stable projects but can stifle change or innovation (Tordup Heeagar, 2012).

Rai et al. (2022) emphasize the critical role of different types of technical documentation, including system documentation, design documents, source code documentation, and test reports. These documents are crucial for software maintenance and reuse, yet they are often neglected due to time and monetary pressures (Rai et al., 2022; McBurney et al., 2018). High-quality documentation provides numerous benefits, such as better comprehension, faster bug repair, and quicker onboarding of new employees (McBurney et al., 2018). Historically, software testing was done manually, which was expensive, slow, and error-prone (DeSouza et al., 2005). Over time, the industry has shifted towards automated testing, making it ubiquitous and essential for code delivery (DeSouza et al., 2005). In contrast, creating and using documentation is fundamental yet remains in a more nascent state. Software engineers and stakeholders struggle with creating, maintaining, and using documentation, which plays multiple roles in the process of creating and delivering software. Programmers must prioritize their documentation efforts, focusing on the most critical areas of code (McBurney et al., 2018).

Documentation plays a crucial role in knowledge sharing and organizational learning, as highlighted by several studies. Cyr and Choo (2010) emphasize the importance of documentation in facilitating knowledge transfer within organizations, particularly in distributed teams. Their research finds that effective documentation practices can bridge the gap between tacit and explicit knowledge, making information more accessible and actionable for team members. Building on this, Fannoun and Kerins (2019) assess the impact of documentation on organizational learning in software engineering. They argue that comprehensive documentation captures the collective knowledge of the development team, thereby enhancing organizational learning.

Heger et al. (2022) focus on the perceptions and needs of machine learning practitioners regarding data documentation. Their study highlights the necessity of

clear and detailed documentation to support the reproducibility and transparency of machine learning models. They identify several challenges, including the complexity of documenting dynamic datasets and the need for standardized documentation practices. Shilton et al. (2013) emphasize the importance of understanding the sociotechnical dimensions of values for design research, particularly in the context of documentation practices. They argue that documentation is influenced by both technical and social factors, and recognizing these dimensions is crucial for creating documentation that aligns with the values and needs of the organization. This perspective underscores the need for documentation practices that are not only technically sound but also socially informed, ensuring they meet the diverse requirements of all stakeholders involved.

The literature on software documentation underscores its pivotal role in enhancing communication, maintenance, and knowledge sharing within the software development lifecycle. Despite its critical importance, documentation practices face numerous challenges, including keeping information current, determining the appropriate level of detail, and integrating seamlessly with developers' workflows. Effective documentation can mitigate risks, improve project outcomes, and facilitate maintenance by providing clear guidelines and reference points. Furthermore, documentation's role in organizational learning and knowledge sharing is emphasized, with comprehensive documentation capturing the collective knowledge of development teams and enhancing learning processes.

Exploring Documentation Practices at Remesh through Ethnographic Methods: Insights from Participant-Observation and Interviews

Ethnographic methods, rooted in anthropology, sociology, and social psychology, offer valuable insights into understanding documentation practices within organizations. These methods involve immersive, qualitative research techniques that capture the complex interactions between people and systems. Shilton et al. (2013) argue that values in design research are not merely attributes of people or systems but emerge from the interplay between them. This perspective is particularly relevant for studying software documentation, where the values and practices of developers, users, and other stakeholders influence the creation and maintenance of documentation. Ethnographic methods can uncover the "accidental"

nature of values embedded in documentation practices, as observed in the case study of documentation development at Remesh. These methods can also reveal the sociotechnical dimensions of documentation, providing a holistic understanding of how documentation practices evolve and impact software development processes.

In order to understand the impact of bringing a documentation strategy to the Remesh engineering team, the authors employed the foundational ethnographic methods of participant-observation, interviewing, and archival review. As the principal engineer for the organization, Nicole Tietz-Sokolskaya developed the documentation protocol initially adopted by the engineering division in 2020 and has greatly influenced the ways in which the protocol has morphed and expanded over the last four years. This deep knowledge has afforded a close look into the workings of documentation, including the challenges engineers have faced. Along with this expert participant-observation knowledge, the authors interviewed eight engineers currently with the organization, who also experienced the transition to a fully remote environment. These engineers were involved in implementing the protocol at varying levels, some with extensive oversight in creating how and when documentation would be used by individual teams, others with more a 'reviewer' role.

The authors developed an interview question set and followed semi-structured interviewing techniques, allowing for the freedom to pull on unexpected conversational threads that emerged. The interviews were recorded using Zoom and Google Meet software and transcribed using Otter.ai or available embedded transcription. The transcripts were checked for accuracy and uploaded into Taguette and Atlas.ti for coding. The authors were interested in experimenting with open-source software as well as traditionally available coding software in the analysis. The authors coded specifically for the themes of *collaboration, learning,* and *knowledge sharing*, but otherwise identified emergent codes upon multiple readings of each transcript. While the Remesh platform was considered as a data collection and analysis option, the small dataset (n=9) did not lend itself to optimizing the platform, which performs best with at least 15–20 participants.

Findings

Documentation at Remesh is a socio-technical practice that embodies the engineer's values. This relationship is evident in how documentation is used to facilitate communication and collaboration among team members. For example, efforts to ensure that documentation speaks the language of both market researchers

and engineers highlight the importance placed on inclusivity and clarity. The iterative process of creating, reviewing, and refining documentation further underscores a value system that prioritizes thoroughness and continuous improvement.

Value: Collaboration and Inclusivity

A notable cultural shift has been in the collaborative nature of documentation efforts, particularly through the design review process, which involves stakeholders and team members in documenting detailed plans and understanding the rationale behind decisions.

> During a time when our team lacked a clear approach to tackling certain tasks, we decided to implement a design review process that had been successful at another company. This involved talking to stakeholders, like our manager, to understand the details, weighing the pros and cons, and exploring alternative approaches. Once we agreed on a solid plan, we documented the design review.
>
> —Engineer E3

This approach ensures thorough vetting and consensus before development, fostering a supportive and efficient work environment. Distributing responsibilities, such as translation tasks, among team members avoids a single point of failure and promotes effective knowledge sharing.

> To avoid a single point of failure, we distributed the responsibility among all teams, creating translation liaisons who were trained and gradually took over the process. This collaboration ensured the knowledge was shared and documented effectively.
>
> —Engineer E4

This inclusivity aids new hires in integrating quickly, reducing dependency on individual knowledge, and promoting a more collaborative culture. Engineers draft documents and seek feedback from peers and managers, ensuring multiple perspectives are considered, thereby improving the quality and accuracy of documentation. Design reviews provide structured feedback, allowing engineers to refine ideas and solutions based on collective input, enhancing the technical soundness of solutions and fostering a culture of collaboration and knowledge sharing.

Value: Learning and Development

Onboarding at Remesh is facilitated through common training and documentation practices that convey the behaviors and values of the engineering team. These practices are crucial in ensuring new team members grasp the significance of documentation and how it embodies the organization's commitment to transparency, thoroughness, and collaboration.

> When I first joined, onboarding was facilitated by small bug tickets assigned to new engineers. This helped us get familiar with the codebase through practical tasks. [The new] system architecture documentation was particularly helpful in understanding how different services relate to each other, which is essential for new hires.
>
> —Engineer E6

Engineers have developed comprehensive guides to help new hires understand the codebase, tools, and processes, enabling them to quickly become productive and minimizing the reliance on informal knowledge transfer. This effort resulted in the evolution of documentation into living documents, regularly updated and maintained by the team to ensure they stayed relevant and reflected the current state of the codebase and processes. But not everything is documented…

> In engineering, we consider technical debt when deciding what to document. If something is straightforward and easily understood, it might not need extensive documentation. However, complex processes or code that requires significant effort to understand are thoroughly documented to save time and effort in the future.
>
> —Engineer E9

The importance of documentation was further recognized through past experiences where the absence of thorough documentation led to significant issues. Consequently, the company culture evolved to regard documentation as a vital component of the engineering process. This cultural shift ensured that documentation was no longer viewed as an afterthought but as an integral part of the development lifecycle.

Value: Continuous Improvement and Adaptability

The interviews revealed a commitment to continuous improvement and adaptability in documentation practices at Remesh, evolving from manual efforts to fully automated processes.

> We improved our processes, creating centralized, comprehensive documentation to aid new hires and ensure continuity despite team changes. Proper documentation prevents knowledge loss when key team members leave and facilitates smoother onboarding for new members.
>
> —Engineer E5

Initially, documentation was manual and time-consuming, with engineers meticulously documenting their test cases and processes to ensure thorough testing and quality assurance. Over time, tools like JIRA and TestRail became central to managing tasks and test cases, ensuring that documentation remained current and relevant. Further, detailed design reviews provided high-level insights into technical decisions and helped mitigate issues during development.

> I believe our current balance of documentation at Remesh is effective. We have detailed design reviews for new features, high-level system architecture documentation, and practical guides for common issues. This approach ensures that documentation remains useful and relevant without becoming outdated or overwhelming.
>
> —Engineer E6

This iterative process of documentation and feedback ensured the team could adapt to changes efficiently, streamlining efforts and ensuring consistency and accuracy.

Value: Responsibility and Accountability

At Remesh, responsibility and accountability have become core values, with engineers taking personal responsibility for documenting their work. Creating comprehensive design documents for complex projects has ensured that new team members can easily understand the system. This practice not only validates design decisions but also provides a reference for future work, ensuring accountability and reducing the risk of knowledge loss.

> When I joined, there was little to no documentation, making it difficult to understand the codebase. Over time, I advocated for better documentation practices to ensure knowledge was shared and accessible, particularly as the team grew and went remote. This shift culminated in the adoption of a robust design review process that significantly improved our documentation quality and inclusivity.
>
> —Engineer E7

The design review process plays a crucial role in ensuring high-quality documentation and effective communication within the engineering team. It involves multiple iterations and feedback loops, helping to identify potential issues early and incorporate diverse perspectives. Design review documents serve as a key reference for engineers, outlining important design decisions and the rationale behind them. The audience for documentation is carefully considered to ensure the content is relevant and useful. For internal engineering teams, documentation is more technical and detailed, whereas for external stakeholders, it is more high-level and explanatory.

> My approach to writing documentation is guided by understanding the audience and their needs. For internal engineering documents, I might cut certain corners, while external documents for clients require more careful crafting. I start by brain-dumping my thoughts and then organizing them into a coherent structure, often using outlines and diagrams to clarify complex concepts.
>
> —Engineer E7

This audience-centric approach helps ensure that documentation serves its intended purpose effectively, whether for internal alignment or external communication.

Value: Technical Excellence and Precision

Technical excellence and precision in documentation are highly emphasized. Engineers use visual Markdown editors and detailed design reviews to ensure that documentation accurately represents the system's design and functionality.

> A recent design review I conducted focused on importing non-Remesh data into our platform for the first time. This involved translating messages into Remesh message types and figuring out how to store both the data and metadata about the upload process. The design review aimed to organize

> these details and communicate the approach to the team. Additionally, I wrote a research-oriented review to parse another software's' undocumented import format, detailing how it translates to our needs and the properties we would use.
>
> —Engineer E8

This meticulous approach fosters a culture of technical excellence, making it easier to replicate and modify work. Documentation is tailored based on purpose, with high-level overviews for general understanding and detailed technical documents for developers. For complex features or architectural changes, comprehensive design documents capture the rationale, decisions, and implementation details, preserving knowledge and providing a reference for future work.

Remesh balances the need for comprehensive documentation with the agile methodology's emphasis on working software, particularly in projects where technical precision and usability are crucial. Documentation requirements vary along a spectrum influenced by industry standards and regulatory demands, and Remesh finds itself in the middle of this spectrum. Writing documentation is recognized as a distinct skill set requiring both technical and communication abilities.

> Early in my career, I learned the importance of documentation the hard way, through experiences where the lack of documentation led to significant issues. For example, at a former employer, a critical processing script was accidentally deleted, and without documentation, it took weeks to restore functionality.
>
> —Engineer E9

The engineering team prioritizes technical excellence and precision by stipulating documentation as a core competency in performance reviews. It is not possible to be promoted without developing strong documentation skills, highlighting the organization's commitment to high-quality documentation as an integral part of engineering excellence.

Recommendations

Based on the insights from data analysis, the following recommendations are suggested for any start up or organization wrestling with software documentation:

Start with Clear, High-Level Documentation

Begin with high-level documentation that provides an overview of systems and processes. This helps new team members get up to speed quickly and understand the context before diving into specifics. Use collaborative documentation tools like Notion or Google Drive for easy iteration and feedback.

Use Design Review as a Standard Practice

Implement a design review process to document the technical approach for new features, including design decisions, alternatives considered, and potential edge cases. Ensure that design reviews are collaborative, involving feedback from team members and relevant stakeholders.

Tailor Documentation to the Audience

Customize the level of detail based on the audience, with high-level overviews for non-technical stakeholders and detailed technical documentation for engineers. Include both the "why" and "how" in technical documentation to provide context and ensure comprehensibility for readers with different expertise levels.

Maintain and Update Documentation Regularly

Establish a process for regularly reviewing and updating documentation to keep it current, especially for onboarding and design reviews. Encourage team members to update documentation as part of their workflow, ensuring it evolves with the product and processes.

Document Test Plans and Processes

Create detailed test plans and cases to guide QA processes and document feature testing. Tools like TestRail can help organize and maintain these test cases. Ensure test documentation is accessible to both developers and QA personnel, facilitating collaboration and shared understanding.

Formalize Onboarding and Knowledge Sharing

Develop comprehensive onboarding documentation covering environment setup, key systems, and processes to help new team members integrate quickly.

Create shared documentation spaces (e.g., Notion, Confluence) for team members to contribute to and access knowledge easily.

Leverage Post-Mortems and Retrospectives

Conduct post-mortems and retrospectives to document lessons learned from projects and feature implementations, informing future work. Include addendums in design reviews to document significant changes made during implementation.

Encourage Collaborative Documentation Practices

Foster a culture of collaboration in documentation by involving multiple team members in the creation and review process. Use documentation as a tool for mentoring and professional development, especially for junior engineers. Additionally, involve team members at all levels in designing documentation processes to ensure they meet actual needs and are effective. Use feedback mechanisms like 1:1 conversations, small group meetings, and collaborative platforms to gather input from all levels.

Prioritize Critical and Complex Areas for Documentation

Focus detailed documentation efforts on critical and complex codebase areas likely to cause issues or require significant effort to understand. Use documentation to mitigate technical debt, ensuring complex processes and systems are well-documented.

Eliminate Ineffective Documentation and Address Friction Points

Trust your team and eliminate documentation that isn't being done if it doesn't serve its intended purpose, suggesting a need for different solutions. Focus documentation efforts on resolving friction points between teams or in the development process to ensure it serves a clear purpose. Regularly assess and refine documentation processes based on team feedback and observed effectiveness.

Implementing one, some, or all of these recommendations will help startups and other organizations establish robust and effective documentation practices.

About the Authors

Nicole Tietz-Sokolskaya is a software engineer with expertise in performance and software architecture. She has over 10 years of experience at software start-ups and consulting. Currently, she is Principal Software Engineer at Remesh, a SaaS AI-enabled research platform, where she works on platform stability and performance, security, processes, and provides top-level leadership.

Suzanne Walsh is an anthropologist with expertise in business and health research. With over 25 years of experience in academia and consulting, she currently works at Remesh, a SaaS AI-enabled research platform, as a Research Consultant. Suzanne helps clients puzzle through sticky organizational research problems, and leverages machine learning, NLP, and generative AI to understand qualitative data at scale. She has published in high-impact journals and authors thought leadership for Remesh.

References Cited

Aghajani, Emad, Csaba Nagy, Olga Lucero Vega-Márquez, et al. "Software Documentation Issues Unveiled." *Proceedings of the 41st International Conference on Software Engineering*, IEEE Press, 2019, pp. 1199–210. *ACM Digital Library*, https://doi.org/10.1109/ICSE.2019.00122

Aghajani, Emad, Csaba Nagy, Mario Linares-Vásquez, et al. "Software Documentation: The Practitioners' Perspective." *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, Association for Computing Machinery, 2020, pp. 590–601. *ACM Digital Library*, https://doi.org/10.1145/3377811.3380405

Andreeva, Tatiana, and Aino Kianto. "Knowledge Processes, Knowledge-intensity and Innovation: A Moderated Mediation Analysis." *Journal of Knowledge Management*, edited by Gregorio Martín-de Castro, vol. 15, no. 6, Oct. 2011, pp. 1016–34. https://doi.org/10.1108/13673271111179343

Asatiani, Aleksandre, et al. "Sociotechnical Envelopment of Artificial Intelligence: An Approach to Organizational Deployment of Inscrutable Artificial Intelligence Systems." *Journal of the Association for Information Systems*, vol. 22, no. 2, 2021, pp. 325–52. https://doi.org/10.17705/1jais.00664

Chen, Ricky, et al. "Probing Documentation Practices: Reflecting on Students' Conceptions, Values, and Experiences with Documentation in Creative Inquiry." *Proceedings of the 13th Conference on Creativity and Cognition*, Association for Computing Machinery, 2021, pp. 1–14. *ACM Digital Library*, https://doi.org/10.1145/3450741.3465391

Chomal, Vikas S., and Jatinderkumar R. Saini. *Software Project Documentation – An Essence of Software Development. International Journal of Advanced Networking and Applications*, vol. 6, no. 06, May/June 2015, pp. 2563–2572.

Clarke, Steven. "What Is an End User Software Engineer?" In *End-User Software Engineering. Dagstuhl Seminar Proceedings*, vol. 7081, p. 1, Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2007) https://doi.org/10.4230/DagSemProc.07081.26

Cyr, Sylvio, and Chun Wei Choo. "The Individual and Social Dynamics of Knowledge Sharing: An Exploratory Study." *Journal of Documentation*, vol. 66, no. 6, Oct. 2010, pp. 824–46. *DOI.org (Crossref)*, https://doi.org/10.1108/00220411011087832

Das, Sumita, et al. "Understanding Documentation Value in Software Maintenance." *Proceedings of the 2007 Symposium on Computer Human Interaction for the Management of Information Technology*, Association for Computing Machinery, 2007, pp. 2-es. *ACM Digital Library*, https://doi.org/10.1145/1234772.1234790

De Souza, Sergio Cozzetti B., et al. "A Study of the Documentation Essential to Software Maintenance." *Proceedings of the 23rd Annual International Conference on Design of Communication: Documenting & Designing for Pervasive Information*, ACM, 2005, pp. 68–75. *DOI.org (Crossref)*, https://doi.org/10.1145/1085313.1085331

Ding, Wei, et al. "Knowledge-Based Approaches in Software Documentation: A Systematic Literature Review." *Information and Software Technology*, vol. 56, no. 6, June 2014, pp. 545–67. *DOI.org (Crossref)*, https://doi.org/10.1016/j.infsof.2014.01.008

Fannoun, Sufian, and John Kerins. "Towards Organisational Learning Enhancement: Assessing Software Engineering Practice." *The Learning Organization*, vol. 26, no. 1, Jan. 2019, pp. 44–59. *DOI.org (Crossref)*, https://doi.org/10.1108/TLO-09-2018-0149

Garousi, Golara, et al. "Usage and Usefulness of Technical Software Documentation: An Industrial Case Study." *Information and Software Technology*, vol. 57, Jan. 2015, pp. 664–82. *DOI.org (Crossref)*, https://doi.org/10.1016/j.infsof.2014.08.003

Heger, Amy K., et al. "Understanding Machine Learning Practitioners' Data Documentation Perceptions, Needs, Challenges, and Desiderata." *Proceedings of the ACM on Human-Computer Interaction*, vol. 6, no. CSCW2, Nov. 2022, p. 340:1-340:29. *ACM Digital Library*, https://doi.org/10.1145/3555760

Herrmann, Thomas, et al. "Sociotechnical Walkthrough: A Means for Knowledge Integration." *The Learning Organization*, edited by Alex Ramirez, vol. 14, no. 5, July 2007, pp. 450–64. *DOI.org (Crossref)*, https://doi.org/10.1108/09696470710762664

Hicks, Catherine. *It's Like Coding in the Dark: The Need for Learning Cultures within Coding Teams*. Catharsis Consulting, white paper, 2022. https://www.catharsisinsight.com/_files/ugd/fce7f8_2a41aa82670f4f08a3e403d196bcc341.pdf

Huang, Shihong, and Scott Tilley. "Towards a Documentation Maturity Model." *Proceedings of the 21st Annual International Conference on Documentation*, Association for Computing Machinery, 2003, pp. 93–99. *ACM Digital Library*, https://doi.org/10.1145/944868.944888

Ifenthaler, Dirk, et al., editors. *Digital Transformation of Learning Organizations*. Springer International Publishing, 2021. https://doi.org/10.1007/978-3-030-55878-9

Kajko-Mattsson, Mira. "A Survey of Documentation Practice within Corrective Maintenance." *Empirical Software Engineering*, vol. 10, no. 1, Jan. 2005, pp. 31–55. *ACM Digital Library*, https://doi.org/10.1023/B:LIDA.0000048322.42751.ca

Königstorfer, Florian, and Stefan Thalmann. "Software Documentation Is Not Enough! Requirements for the Documentation of AI." *Digital Policy, Regulation and Governance*, vol. 23, no. 5, Nov. 2021, pp. 475–88. https://doi.org/10.1108/DPRG-03-2021-0047

Lethbridge, T. C., et al. "How Software Engineers Use Documentation: The State of the Practice." *IEEE Software*, vol. 20, no. 6, Nov. 2003, pp. 35–39. https://doi.org/10.1109/MS.2003.1241364

McBurney, Paul W., et al. "Towards Prioritizing Documentation Effort." *IEEE Transactions on Software Engineering*, vol. 44, no. 9, Sept. 2018, pp. 897–913. https://doi.org/10.1109/TSE.2017.2716950

Meng, Michael, Stephanie Steinhardt, et al. "How Developers Use API Documentation: An Observation Study." *Communication Design Quarterly*, vol. 7, no. 2, Aug. 2019, pp. 40–49. *ACM Digital Library*, https://doi.org/10.1145/3358931.3358937

Meng, Michael, Stephanie M. Steinhardt, et al. "Optimizing API Documentation: Some Guidelines and Effects." *Proceedings of the 38th ACM International Conference on Design of Communication*, Association for Computing Machinery, 2020, pp. 1–11. *ACM Digital Library*, https://doi.org/10.1145/3380851.3416759

Mohamed, Shakir, et al. "Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence." *Philosophy & Technology*, vol. 33, no. 4, Dec. 2020, pp. 659–84. https://doi.org/10.1007/s13347-020-00405-8

Nakayama, Makoto, et al. "Agility and System Documentation in Large-Scale Enterprise System Projects: A Knowledge Management Perspective." *Procedia Computer Science*, vol. 181, 2021, pp. 386–93. https://doi.org/10.1016/j.procs.2021.01.181

Nassif, Mathieu, and Martin P. Robillard. "A Field Study of Developer Documentation Format." *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, Association for Computing Machinery, 2023, pp. 1–7. *ACM Digital Library*, https://doi.org/10.1145/3544549.3585767

Newman, Christian D., et al. "Automatically Generating Natural Language Documentation for Methods." *2018 IEEE Third International Workshop on Dynamic Software Documentation (DySDoc3)*, 2018, pp. 1–2. *IEEE Xplore*, https://doi.org/10.1109/DySDoc3.2018.00007

Pasuksmit, Jirat, et al. "Towards Just-Enough Documentation for Agile Effort Estimation: What Information Should Be Documented?" *International Conference on Software Maintenance and Evolution* (ICSME), 2021, arXiv:2107.02420

Raglianti, Marco. "Topology of the Documentation Landscape." *Proceedings of the ACM/IEEE 44th International Conference on Software Engineering: Companion Proceedings*, Association for Computing Machinery, 2022, pp. 297–99. *ACM Digital Library*, https://doi.org/10.1145/3510454.3517068

Rai, Sawan, et al. "A Review on Source Code Documentation." *ACM Transactions on Intelligent Systems and Technology*, vol. 13, no. 5, Oct. 2022, pp. 1–44. https://doi.org/10.1145/3519312

Renzl, B. "Trust in Management and Knowledge Sharing: The Mediating Effects of Fear and Knowledge Documentation." *Omega*, vol. 36, no. 2, Apr. 2008, pp. 206–20. https://doi.org/10.1016/j.omega.2006.06.005

Robillard, Martin, et al. "On-Demand Developer Documentation." *2017 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, Shanghai, China, 2017, pp. 479–483. https://doi.org/10.1109/ICSME.2017.17

Shilton, Katie, et al. "Charting Sociotechnical Dimensions of Values for Design Research." *The Information Society*, vol. 29, no. 5, Oct. 2013, pp. 259–71. https://doi.org/10.1080/01972243.2013.825357

Tamburri, Damian A. "Software Architecture Social Debt: Managing the Incommunicability Factor." *IEEE Transactions on Computational Social Systems*, vol. 6, no. 1, Feb. 2019, pp. 20–37. https://doi.org/10.1109/TCSS.2018.2886433

Wagenaar, Gerard, et al. "Working Software over Comprehensive Documentation – Rationales of Agile Teams for Artefacts Usage." *Journal of Software Engineering Research and Development*, vol. 6, Dec. 2018. https://doi.org/10.1186/s40411-018-0051-7

Waheed, Sara, et al. "Improving Knowledge Sharing in Distributed Software Development." *International Journal of Advanced Computer Science and Applications*, vol. 10, no. 6, 2019. https://doi.org/10.14569/IJACSA.2019.0100656

Zhi, Junji, et al. "Cost, Benefits and Quality of Software Development Documentation: A Systematic Mapping." *Journal of Systems and Software*, vol. 99, Jan. 2015, pp. 175–98. https://doi.org/10.1016/j.jss.2014.09.042

Mother Tongue: The Role of Vernacular Knowledge in Chronic Illness and Healthcare

JULIANA SALDARRIAGA, *A Piece of Pie*

This work demonstrates how and why the expertise of non-medical actors, such as patients, caregivers, and patient advocacy groups, is essential for improving patient outcomes and quality of life. 'Vernacular' knowledge, including community beliefs, values, networks, and practices, both challenge and collaborate with institutional scientific knowledge. The value of this interaction of multiple kinds of knowledge is illuminated through a study of spinal muscular atrophy in Colombia, which explores how mother caregivers develop unique expertise, devices, and care practices, and use online social networks to share and learn both clinical and vernacular information. Supported by a broad body of evidence and theory, this paper makes a powerful case for elevating vernacular knowledge and the people and communities who practice it. It also showcases methods and channels that researchers and medical professionals can use to collaborate with caregivers in these non-institutional, but medically essential, spaces.

Introduction

There is an ongoing debate among scholars regarding the concept of the "vernacular". In the field of architecture, anthropologists have discouraged its use, suggesting it has become just another concept to essentialize the Other, much like "primitive" or "folk" (Vellinga 2011). However, environmental studies have argued that expert science is "not enough" to make competent decisions (Turner 2004), and emphasize the importance of local knowledge to address ecological crises. In these studies, vernacular knowledge is defined as an integration of expert science with community beliefs and values, or a negotiation or collaboration between different forms of knowledge, rather than a dichotomy (Simpson et al. 2015).

In this case study, we apply the concept of the vernacular as *negotiation* to healthcare contexts, specifically to the knowledge developed and shared by chronic patients, caregivers, and patient advocacy groups. This concept is appropriate because this knowledge, rather than being in direct *opposition* to institutional medical knowledge, *negotiates* with it by embedding it in specific sociocultural contexts. Vernacular knowledge in healthcare, despite being non-institutional and sometimes non-scientific, is legitimate because it has the potential to achieve what is arguably the primary goal of medical knowledge: improving health outcomes and patients' quality of life.

Saldarriaga, Juliana. "Mother Tongue: The Role of Vernacular Knowledge in Chronic Illness and Healthcare." *EPIC Proceedings* (2024): 396–410. https://www.epicpeople.org/vernacular-knowledge-in-chronic-illness-and-healthcare. ISSN 1559-8918.

At A Piece of Pie, a global consulting firm with experience in medical anthropology and patient-centricity, we applied the concept of vernacular knowledge to a study on spinal muscular atrophy (SMA). SMA is a disorder that affects the cells controlling voluntary muscle movement, leading to muscle weakening and atrophy. Its most common type is typically evident before six months of age, meaning these children require permanent care, usually provided by their mothers. We saw an opportunity in SMA to understand the role of the vernacular knowledge displayed by these mother caregivers. After conducting an ethnographic study with five of them, we learned that their vernacular knowledge is negotiating three seemingly non-negotiable foundations in Western medicine. The first is the "know-it-all" doctor trope, which refers to society's inability to acknowledge other kinds of expertise in medicine. The second is the rigid "hierarchy of evidence" within the current medical paradigm, where statistical evidence is considered more truthful and is preferred over ethnographic evidence. Finally, the third is the exclusive channels through which knowledge is exchanged, namely academic journals and medical conferences.

Negotiating these foundations entails having multiple ways to answer questions such as *who* talks in medicine (experts), *what* they talk about (evidence), and *how* they talk about it (channels). In the following sections, we will provide a general context on SMA, as well as the study's objective and methodology. Afterwards, we delve into each of the negotiated foundations, demonstrating how ethnography (with its mindset and methods) enabled us to detect negotiation between bodies of knowledge rather than opposition. We close our case study arguing these negotiated foundations are not specific to SMA; they apply to chronic illnesses and conditions in general, in the sense that these patients and their caregivers are all developing their own bodies of knowledge.

Study Overview

Disease Context

Spinal muscular atrophy (SMA) is a genetic neuromuscular disorder that affects motor neurons, the cells that control all voluntary muscle movement, leading to muscle weakening and atrophy. There are several types of SMA, classified based on the age of onset, with type 1 being the most common and typically evident before six months of age. The incidence of SMA is calculated to be between 1 in 6,000 to 1 in 11,000 births (Cleveland Clinic, n.d.). Because it is such a rare disorder, diagnostic delay remains a significant challenge, especially when primary care physicians and pediatricians are not familiar with it (Cao 2021; Pera 2020).

There is no cure for SMA, but recent treatment developments aim to stimulate the production of a specific protein that supports motor neurons. These therapies are extremely expensive and, in healthcare systems with limited resources, are typically approved only for patients with a timely diagnosis, whose motor neurons have not yet deteriorated. For those without access to these disease-modifying therapies, treatment focuses on managing symptoms and preventing complications. This usually includes breathing support (sometimes through tracheostomy), feeding tubes or gastrostomy, physical therapy, occupational therapy, and speech-language therapy. Children with type 1 SMA therefore require permanent care, which is most often provided by their mothers.

In contexts with limited access to disease-modifying therapies, as well as support devices, procedures, and therapies, mother caregivers and patient advocacy groups (PAGs) have become a strong, activist, and demanding force. In Latin America, they are widely using social media platforms, such as Instagram, to raise awareness about SMA, fundraise, sign petitions, pressure decision-makers, and illustrate the daily challenges faced by these patients. It is due to this dynamic context that we became motivated to conduct ethnographic research on SMA.

Study Objective and Methodology

The present study was conducted by A Piece of Pie (APoP), a global consulting firm that applies medical anthropology, design services, and innovation to help stakeholders address healthcare challenges. APoP became familiar with SMA due to a project for a pharma company that was interested in learning about the disorder's current situation in Colombia. For this previous project, we engaged with neurologists and other healthcare professionals (HCPs) that diagnose and treat SMA. Through these interactions and complementary desk research, we learned that mother caregivers and PAGs are empowered, self-organized, and well-informed, and that they have successfully used social media to increase their influence and reach. We therefore decided to conduct a study autonomously, this time focusing on mother caregivers and with the hopes of understanding how they came to assume this leadership role via social media. The study was focused on Latin American caregivers, because this was the region we were already familiar with in terms of SMA.

The research methods used were social listening and in-depth interviews. Social listening was conducted to learn what was being said online about SMA in Latin America. Although research teams can use tools such as Hootsuite Insights, on this occasion we conducted a more organic exercise by identifying opinion leaders (both medical and non-medical), searching news articles, blogs, and webinars, and

following specific hashtags on LinkedIn, Instagram, and X. An organic, as opposed to a tool-based, approach is adequate when the topic of interest is highly specific, such as SMA in Latin America. With this approach, we were able to find five robust Instagram accounts of Latin American mothers with children diagnosed with type 1 SMA: three were Colombian (but one of them lived in Florida), one was Brazilian, and one was Venezuelan.

With the social listening exercise, we gathered a broad understanding of common experiences and concerns among the five caregivers. We put this baseline knowledge to the test conducting in-depth interviews (IDIs) with three of them. The three IDIs were conducted remotely and included virtual home tours. We inquired about routines, early signs and diagnosis, access to treatment (including therapies, procedures, and support devices), and their goals and expectations using social media. The results obtained via social listening and IDIs are described in the following section.

Study Results: Negotiated Foundations

In this section, we present the study's results as negotiated foundations, or foundations of Western institutional medicine that, at first glance, appear non-negotiable. We focus on three: the first consists in the socially accepted idea or trope of the know-it-all doctor. The second refers to the rigid way in which different types of evidence are categorized in the current medical paradigm, which tags ethnographic evidence as the least desirable, and the third is the exclusivity of the institutional channels through which medical knowledge is shared (academic journals, medical congresses).

Foundation 1: The Know-It-All Doctor Trope

The speed at which medical knowledge is currently being generated is unprecedented, with estimates suggesting it doubles every 70 days (Densen 2011). In the field, we've seen that this has led HCPs to embrace what Microsoft CEO Satya Nadella calls a learn-it-all mindset, vs. a more traditional know-it-all one (Sweales 2019). And yet HCPs are still perceived as all-knowing, by most patients and by society as a whole. It is difficult to accept that HCPs also have cognitive biases that affect their decision-making process (Marino 2020), or that during patient consultations primary care physicians are asked questions they need to look up answers to (Kahane 2011). This know-it-all doctor trope is historical; it can be traced back to the nineteenth century, when "expert knowledge" embodied in the (male)

HCP disrupted previously "taken-for-granted practices" in Western modern societies (Escobar 2018). The knowledge possessed by these experts established itself as more objective and universal and thus more legitimate (Nieto 2019) and it still lingers in Western healthcare systems.

Our SMA study suggests there is a different kind of expert when it comes to chronic illness: that of the expert patient or caregiver, who has no choice but to become an expert to face daily challenges. One of the ways in which these experts differ from the know-it-all doctor trope is their fluid identity. They switch between knowing and learning, as seen in the case of mothers of patients with SMA. Just as they're constantly learning from HCPs, they too have developed their own solutions out of necessity, which are validated and praised by HCPs. Mother1 described her own technique to remove mucus from her son's airway walls, because her healthcare insurance wouldn't cover vest therapy for patients with SMA. Mother2, while being grateful for the HCPs who advise her and care for her daughter, also shared do-it-yourself (DIY) solutions with us that demonstrate a deep understanding and analysis of the daily needs that patients with SMA have. Take, for example, an adapter that she and her husband attached to their daughter's electric wheelchair, to secure the ventilator required for ventilatory support. Or how, instead of having special shoes made for her daughter's insteps, she simply added platform shoe soles to her orthopedic boots, which has been more practical and cost-effective. Or, finally, a PVC pipe structure built over her daughter's bed, to organize the "mess" caused by the cables of her life-support devices. HCPs might advise on all these devices, but ensuring they're manageable and work together is certainly up to the caregiver expert.

Figure 1. PVC pipe structure built over the bed over a patient with SMA. Credits: @princess_julieta Instagram account.

The ability to switch between knowing and learning makes these mothers well-suited for teamwork with HCPs. Much is being said in healthcare about the need for integrated care and for higher collaboration between stakeholders, an idea that clashes with the know-it-all doctor trope. In our study, we were pleased to witness mothers and nurses working together to define the best therapy routine for the patient (speech-language therapy, occupational therapy, respiratory therapy, etc.). As Mother1 put it, "What scares me the most is going somewhere when I'm not involved in the decision-making process."

Besides this fluid identity, the patient/caregiver expert has other traits worth discussing. While HCPs are socially expected to super-specialize, the knowledge this kind of expert possesses grows in many directions. It is an organic, fragmented growth that comes from managing the multifaceted consequences of a chronic condition, which extend beyond just the clinical aspects. In our SMA study, we witnessed how these mothers inevitably delve into pedagogical debates, when considering the kind of education their children should receive. "I think, until what point is it relevant for my kid to know who conquered who", Mother2 wonders, a sentiment also seen in Mother1 and Mother3, both interested in nurturing the skills their children will require to use future assistive communication technologies. Even

in SMA cases without cognitive impairment, standard education can be unsuitable for these patients; it may interfere with the patient's therapies, and in some contexts, transportation and school facilities are not disabled-friendly.

There are two final traits of the patient/caregiver expert we can appreciate in our study. The first is that they are much more relatable experts than HCPs; just as they use Instagram to share medical information on SMA, they discuss the emotional toll of the condition, and delve into mental health issues in a non-aspirational, non-influencer kind of way. Perhaps Instagram caregivers are where to look for alternatives to what scholars have called hedonic mental health, or mental health that is reduced to aesthetical healthy habits (Fischer 2009). The fact that they also share life "beyond SMA", such as birthdays, holiday celebrations, and daily life, contributes to this relatability. As mentioned in an article about "heartland influencers" from The New Indian Express, "Influencers from small towns, with their modest backgrounds, are resonating with the average user in a way that is genuine and familiar." (New Indian Express 2024). Second, these mothers are demonstrating how becoming an expert in a specific condition also includes normalizing it—not to make it less severe, but to let go of social stigma associated with it. As Mother2 puts it, "For me, my kid is a seven-year-old boy. I do not see him as sick, and I do not attend to his needs with this perception."

With this description of the patient/caregiver expert (in SMA, but in general too), our intention is not to discredit HCPs or institutional medicine. We simply wish to demonstrate it's possible to negotiate seemingly unalterable foundations in Western medicine, by recognizing there can be multiple ways to accumulate and exhibit expertise in this field. This is particularly relevant when we're focusing on knowledge that is exhibited not for personal recognition or prestige, but to improve patients' quality of life. Also worth mentioning is that it is precisely the ethnography's open-mindedness what enables us to see expertise where (according to institutional standards) there shouldn't be.

Foundation 2: Rigid Hierarchy of Evidence

As ethnographers working in healthcare contexts, we have witnessed the rise of evidence-based medicine or EBM. EBM can be described as the current paradigm in medicine, one which prioritizes evidence from biostatistics, engineering, and epidemiology, and questions the reliability and truthfulness of testimonies and anecdotes (Adams 2013). "…For evidence to say anything valid about "how to prevent or treat a known health problem" it must speak the language of statistics and epidemiology […] In this ordering of priorities, one finds a simultaneous discrediting

of other forms of knowledge and evidence and other ways of conveying truth…" (Adams 2013, 57).

It is provocative to conduct ethnographic research under this paradigm. Rather than discrediting statistical evidence, our interest lies in negotiating this rigid hierarchy of knowledges and highlighting the value of ethnographic evidence as powerful in its own right, not as something complementary or secondary. This value is easy to see in our SMA study: Mother3 described how measuring her daughter's heart rate, in additional its clinical benefits, has also become a way of communicating with her daughter. "Her heart rate increases when she doesn't like something. If her heart rate is elevated, we see her face and we can understand something is upsetting her, like she wants a diaper change. It has enabled us to understand how she is filling." To understand the communication properties of heart rate measurements, something unconventional and yet real, is beyond the scope of clinical trials, the evidence type that scores highest in EBM.

Mother3's extensive work with kinesiology tape also highlights the value of ethnographic evidence. She describes how she has used this tape for muscle support in her daughter's hands and feet, and adds she is "not going according to any literature". And yet, with her daughter's visible improvements, she has "proof that it works and could work in more patients as well."

Figure 2. Kinesiology tape use in patient with SMA. Credits: @unidosporluciana Instagram account.

In our SMA study we have two more examples of evidence that, while generated unconventionally according to EBM standards, are still powerful. First, Mother1 told us she used videos provided by Mother2 to "fight" for nebulization. Her health insurance didn't cover nebulization therapy for patients with SMA, but with this video of Mother2's daughter and how she had improved, Mother1 obtained the therapy for her son. We should therefore acknowledge that videos casually posted on social media can impact decision-making in healthcare systems. As Mother1 expressed, "With the video, we are imparting knowledge to others. If they ever tell me there's nothing to do, I now know what I could show them." Second, Mother1 described how a group of pedagogy students reached out to conduct a project with her son and his educational therapies. This way, evidence was generated in the form of a student project, a learning-by-doing space was organically created, and non-medical students were able to work closely with a rare disease that most HCPs such as primary care physicians never see.

As ethnographers working in healthcare contexts, supported by medical anthropology theory, our job is to visualize illness as a phenomenon embedded in specific sociocultural contexts. This requires an exploration of illness not only in its clinical dimension, but in others as well. And because non-clinical dimensions are equally important, so are the types of evidence that inform us about these.

Foundation 3: Exclusive Information Channels

Years of global fieldwork in healthcare contexts have taught us that academic journals and medical congresses are always among HCPs' preferred information sources. Because it's where to find the most up-to-date information, as well as the key opinion leaders of every medical field, it is understandable that these two sources have become the institutional (foundational) way in which medical knowledge is shared. However, approaching HCPs ethnographically has also enabled us to see these sources can be exclusive in many ways. This is best seen in the following barriers:

- Not all HCPs are able to attend medical congresses due to budget or time constraints.
- There can be language barriers for HCPs not completely fluent in English (this applies to academic journals and congresses).
- The most up-to-date information shared in journals and at congresses is often derived from studies and trials conducted on specific regions (Western Europe, North America) and with specific patient profiles.

- HCPs struggle to find time to read all the academic papers available, and to filter or curate those which could be most relevant to them.

By acknowledging these barriers, one becomes motivated to explore alternative channels in which knowledge could be exchanged in healthcare contexts. This was possible in our SMA study: in our social media exercise, we observed how these mothers are successfully using social media—and specifically Instagram—to share both clinical and vernacular information. Clinical, because they are sharing general information about SMA in a more relatable (but not less scientific!) way. Many posts are focused on early symptoms and diagnosis, to inform followers about babies' developmental milestones and warning signs. Vernacular, because they use Instagram to exhibit their DIY solutions previously mentioned (a result of their caregiver expertise), and are constantly recommending how to manage SMA beyond its strictly clinical dimension. This includes detailed descriptions of their children's morning and night routines, therapies that can be effectively done without using specialized equipment, room layout, tips to leave the house and attend public places, etc.

Figure 3. Occupational therapy for SMA patients using household items. Credits: @unidosporluciana Instagram account.

These examples suggest social media used by patient/caregiver experts should be recognized as a valid channel in which knowledge is exchanged in healthcare. In the table below, we show how the Instagram profiles of our study's mothers, by being less exclusive, are not susceptible to the barriers just mentioned, that we've identified in the field.

Table 1. Susceptibility of Social Media to Exclusivity Barriers

Barrier	Social media (Instagram)
Inability to attend medical congresses due to budget or time constraints	Social media posts are completely available to anyone with a smartphone.
Language barriers for HCPs not completely fluent in English	Social media posts are made in the local language and even include colloquial and slang terms. Mother1, for example, posts in Spanglish, appealing to her followers in Miami and Orlando, without disconnecting from Colombian or American populations. In other words, these channels can take advantage of the less formal and rigid relation they have with language.
Lack of patient representation/diversity in most up-to-date and relevant studies and trials	Social media posts, as shown in our SMA study, have a local focus. They talk about the challenges of the local healthcare system and, when there is a new development regarding a disease and its treatments, they discuss the implications of this for their local contexts. Additionally, patient/caregiver experts are providing the exact information their local followers need in their Instagram profiles: they do this by answering questions via stories or during lives.
Lack of time to filter/curate and read the right papers	Social media content can be consumed and can be informative in just a few seconds. On Instagram, this is possible because content in diverse formats beyond text: carrousels, reels, stories, infographics, etc.

With this comparison, we are certainly not discrediting academic journals and medical congresses. After all, the peer-review process required for journal publication is what ensures the quality, veracity, and progress of medical research. Our intention is for HCPs and other healthcare stakeholders to consider the possibility of knowledge channels beyond the foundational/institutional. One could argue that these alternative channels like social media are not targeting HCPs at all, but we can think of at least three reasons why this isn't true: first, they allow HCPs to learn what "real" patients, caregivers, and their families are asking and commenting on regarding a disease or condition. This can inform them on what topics they

should prioritize during consultations. Second, our fieldwork shows that referring patients to "official" online sources (patient associations' social media accounts) is more useful and significant than handing them printed brochures during consultations—especially when a patient has just been informed about their diagnosis.

Third, HCPs have to accept that people are using many sources to inform themselves, beyond the medical consultation. We still encounter HCPs that look down upon patients doing research on "Dr. Google"—a derogatory way to refer to online search— when the attitude they should adopt is instructing patients how to do this responsibly. A study conducted in the UK about vernacular knowledge in contraception concluded that patients "…make sense of a lot of information from different sources, and that different sources of information served different functions. Medical information was gained from the National Health Service (NHS) website, or conversations with clinicians, whereas the more embodied types of knowledge—what a contraceptive might actually be like—was provided by friends and family." (Newton 2024, 1).

A final comment on social media and alternative knowledge channels: they enable patients and caregivers to connect with each other, which is especially relevant in rare diseases such as SMA. These organic communities provide clinical information, but also emotional support. In our study, Mother1 met Mother2 on Instagram and now calls her "my doctor": "If I have a question, I don't call the doctor, I call her."

Discussion

The three negotiated foundations presented here are not specific to our SMA study. We have observed similar dynamics while conducting ethnographic research in diverse regions and with various chronic illnesses and conditions. For instance, in a recent study on hereditary angioedema (HAE) with US patients, we found that patients are not given explicit instructions regarding "trigger" foods they should avoid to prevent the unexplained swelling caused by HAE. Consequently, they become experts through an empirical, trial-and-error process, and are keen to share this feedback with their HCPs. This mirrors the agency demonstrated by SMA caregiver experts and exemplifies vernacular knowledge that leads to significant improvements in quality of life. As Newton (2024) states, "Studying such knowledge can reveal how people conventionally understand the ways in which medical interventions act on the body." (Newton 2024, 1).

Likewise, in a study on atopic dermatitis (AD) in Europe and Asia, our research team encountered a mother caregiver who, after struggling for months to get her toddler to sit still while she applied an ointment, created a nursery rhyme to actively involve him in the treatment routine. As seen on SMA, while the HCP advises on treatment, the caregiver expert finds a way to do it, even if it's an unconventional way from a clinical point of view. In this AD study, we also witnessed how mothers and their teenage daughters with AD are exchanging makeup tips on social media. Makeup isn't saving anyone's lives, but it is an example of vernacular knowledge that improves the quality of life of a teenage girl, and grants her a sense of "normalcy".

These two brief examples, plus decades of fieldwork with chronic patients and caregivers conducted by APoP, demonstrate the value of our SMA study: shining light on vernacular knowledge in healthcare contexts, as alternative knowledge that is improving quality of life, yielding better health outcomes, and empowering patient communities. In this sense, is it knowledge that shouldn't be overlooked by HCPs, considering it can improve the patient-HCP relationship and optimize medical consultations. As expressed by Simpson et al. (2015), building partnerships with vernacular actors is essential to tackle complex problems.

Two final reflections related to vernacular knowledge, which we do not explore in depth in our SMA study but find provocative for further research in alternative knowledges: the first is Chilean sociologist Martín Tironi's argument that the most radical innovation lies within local and ancestral wisdom and practices (Tironi 2023). This is relevant to ethnographers and designers working in industry (in healthcare and beyond), considering business organizations are thirsty for innovation and often expect us to fulfill this. The second is Ehrenreich and English's observation on the historic role women have played in medicine as autonomous, non-institutional healers, or healers of the poor and vulnerable (Ehrenreich and English 2010). This can definitely be used as a starting point to conduct further research on women as carriers of vernacular knowledge.

About the Author

Juliana Saldarriaga is a Colombian anthropologist with experience in business consulting, particularly in healthcare consulting in LATAM and the US. She has a growing interest in medical anthropology and is conducting research on health inequalities and the digital transformation and AI in healthcare. As a project manager at A Piece of Pie, she applies ethnography, systems thinking, and feminist frameworks to work with health insurance companies, hospitals, and the pharmaceutical industry.

References Cited

Adams, Vincanne. 2013. "Evidence-Based Global Public Health." In *When People Come First: Critical Studies in Global Health*, edited by J. Biehl and A. Petryna, 54-90. Princeton: Princeton University Press.

Cao, Yanyan, Miaomiao Cheng, Yujin Qu, Jinli Bai, Xiaoyin Peng, Xiushan Ge, Yuwei Jin, Hong Wang, and Fang Song. 2021. "Factors Associated with Delayed Diagnosis of Spinal Muscular Atrophy in China and Changes in Diagnostic Delay." *Neuromuscular Disorders 31* (6): 519–527. https://doi.org/10.1016/j.nmd.2021.03.002

Cleveland Clinic. n. d. "Spinal Muscular Atrophy (SMA)." Accessed July 7, 2024. https://my.clevelandclinic.org/health/diseases/14505-spinal-muscular-atrophy-sma

Densen, Peter. 2011. "Challenges and Opportunities Facing Medical Education." *Trans Am Clin Climatol Assoc* 122: 48–58.

Ehrenreich, Barbara, and Deirdre English. 2010. *Witches, Midwives and Nurses: A History of Human Healers.* The Feminist Press.

Escobar, Arturo. 2018. *Designs for the Pluriverse: Radical Interdependence, Autonomy, and the Making of Worlds.* Durham: Duke University Press.

Fischer, Mark. 2009. *Capitalist Realism: Is There No Alternative?* Winchester: Zero Books.

Kahane, Steven. 2011. "Must We Appear to Be All-Knowing?" *Can Fam Physician 57* (6): 228–236.

Marino, Brielle A. 2020. "Doctors Don't Know It All. The dangers of cognitive biases in medical practice." *Psychology Today*, June 10, 2020. https://www.psychologytoday.com/us/blog/adventures-in-cognition/202006/doctors-dont-know-it-all

New Indian Express. 2024. "From Village to Viral: India's Heartland Influencers Reshape Social Media Landscape." *New Indian Express*, April 20, 2024. https://www.newindianexpress.com/magazine/2024/Apr/20/from-village-to-viral-indias-heartland-influencers-reshape-social-media-landscape-2

Newton, Victoria L. 2024. "Vernacular Knowledge about Contraception: An Interdisciplinary Perspective on Myths, Misperceptions and Lived Experience." *BMJ Sex Reprod Health 2024* (0): 1–ß2. https://doi.org/10.1136/bmjsrh-2024-202355

Nieto, Mauricio. 2019. *Una Historia de la Verdad en Occidente: Ciencia, Arte, Religión y Política en la Conformación de la Cosmología Moderna.* Fondo de Cultura Económica.

Pera, Maria C., Giorgia Coratti, Beatrice Berti, Adele D'Amico, Maria Sframeli, Emilio Albamonte, Roberto de Sanctis, Sonia Messina, Michela Catteruccia, and Giorgia Brigati. 2020. "Diagnostic

Journey in Spinal Muscular Atrophy: Is It Still an Odyssey?" *PloS One 15* (3). https://doi.org/10.1371/journal.pone.0230677

Sweales, Jo. 2019. How to Introduce a Learn-It-All Culture in Your Business: 3 Steps to Success. *Microsoft Blogger Series*, October 1, 2019. https://www.microsoft.com/en-gb/industry/blog/cross-industry/2019/10/01/introduce-learn-it-all-culture/

Vellinga, Marcel. 2011. "The End of the Vernacular: Anthropology and the Architecture of the Other." *Etnofoor 23* (1): 171–92.

Simpson, Hugh, Rob de Loë, and Jean Andrey. 2015. "Vernacular Knowledge and Water Management—Towards the Integration of Expert Science and Local Knowledge in Ontario, Canada." *Water Alternatives 8* (3): 352–372.

Tironi, Martín. 2023. *Design for More-Than-Human Futures*. London: Routledge.

Turner, Stephen. 2004. "Quasi-science and the State: "Governing Science" in Comparative Perspective." In *The governance of knowledge*, edited by N. Stehr, 241–268. New Brunswick: Transaction Publishers.

The People Say: How We Built a Groundbreaking Public Research Platform to Catalyze Human-Centered Policymaking

CHELSEA MAULDIN, Public Policy Lab
RAJ KOTTAMASU, Public Policy Lab
PETEY ROUTZAHN, Public Policy Lab

The People Say is a pioneering research repository and a model for making rich qualitative data available for policy innovation. This case study describes the rigorous research and data processing techniques used to create the platform, as well as the immediate impact it had in informing policy. It addresses a key challenge for government agencies and NGOs, which increasingly involve the public in formative research, yet face limitations around recruitment, lack of longitudinal studies, and poor engagement policies. Created by Public Policy Lab and SCAN Foundation, the platform is designed to improve policymaking for older adults, with 108 topic and subtopic tags that users can use to filter and associate data with existing quantitative studies. The project also underscores the value of qualitative research in capturing lived experiences and shifting systems to meet the needs of marginalized populations. Future goals include longitudinal research, expanding to other demographics and locations, and supporting civic researchers in creating more public research.

Introduction

Today, government agencies and NGOs are increasingly engaging in formative research with the public. At minimum, they seek to validate digital products before launch. More ambitiously, some attempt to engage the public in the co-design of policy and service delivery.

However, current practice has multiple major limitations:

- recruiting diverse humans is time-consuming and often fails to meet inclusion goals;
- research is typically one-off, rather than meaningfully longitudinal and contextual, producing shallow findings;
- public-sector entities have weak (or actively unethical) policies and practices for engagement, particularly relating to compensation, consent, and participant confidentiality;

Mauldin, Chelsea, Raj Kottamasu, and Petey Routzahn."The People Say: How We Built a Groundbreaking Public Research Platform to Catalyze Human-Centered Policymaking." *EPIC Proceedings* (2024): 411–434. https://www.epicpeople.org/the-people-say-public-research-platform. ISSN 1559-8918.

- research findings are treated as proprietary, causing different agencies and organizations to re-study the same issues and communities, wasting resources and undermining public trust;
- members of the public have limited or no visibility into how their personal data is used, creating cynicism about engagement efforts; and,
- ultimately, the research process extracts value from (often marginalized) people while reserving all real interpretive and decision-making authority to high-power actors, undermining democratic ideals.

In an effort to respond to these issues, the authors' organization, the Public Policy Lab, formed a partnership in 2023 with The SCAN Foundation. The Public Policy Lab (PPL) is a nonprofit organization with a mission to design policy and services that help Americans build better lives. The SCAN Foundation (TSF) is an independent public charity devoted to transforming care for older adults to ensure all of us can age well with purpose.

Together, we committed to conducting a research project and subsequent creation of a public data platform to catalyze human-centered policymaking around the health and wellbeing of older adults. That platform, The People Say (https://thepeoplesay.org/), launched in July 2024. This case study describes the goals of the project, the activities we conducted to develop the platform, the lessons learned conducting the work, and the next steps for the development of the project – and for the field of civic design research.

Project Context and Goals

Current Challenges

Health and aging policies and systems are designed to respond to the needs of system power-holders, rather than that of citizens and consumers. There is a need to elevate what members of the public actually want, if government policies and healthcare-delivery systems are to reflect their needs – and generate awareness and competition among regulators, providers, and payers to respond to those needs.

When advocates, policymakers, and providers do seek to learn from older adults, those engagements often have significant limitations:

Existing Datasets and Surveys Do Not Adequately Focus on Priority Populations

When current research refers to older adults, particularly people dually eligible for Medicare and Medicaid, it often describes this population being heterogeneous – but it doesn't break down the real differences among populations. Research often fails to prioritize marginalized or vulnerable people and tailor research questions and activities to those populations – even though marginalized populations are most at risk of harm if policy fails to account for their needs.

While there's a growing interest in developing 'people-centered' systems and policies, government is ill-equipped to conduct this research, leaving a knowledge gap that needs to be filled. Further, government systems are often set up only to deliver either universal services or ones that have been customized to individual needs, often after burdensome appeals or litigation processes. If major themes related to the needs of different populations can be identified, policy stakeholders can develop 'flavors' of policy that speak to different community or life situations, rather than grappling with individualized services (which can be perceived as unachievable and expensive).

Existing Datasets and Surveys Do Not Adequately Capture the Lived Experiences of Older Adults

In addition, policy recommendations often rely solely on survey and claims data, and stakeholders can generally find refuting data points. A more effective approach is to couple survey and claims data with human-centered design, to not only capture *what* people experience, but also *why* and *how they feel* about their experiences.

However, currently, human-centered research efforts are often structured as one-time activities, even though health access and delivery are experienced as a series of touchpoints over time – leading to limited understanding of how people's requirements evolve as their health, financial, housing, and other personal circumstances change over the decades after they turn 65. If stakeholders can be presented with compelling narratives that illuminate the real needs of older adults-and importantly – how those needs evolve over time – they may be more likely to take action.

Existing Datasets and Surveys Do Not Adequately Address the Urgent Need to Engage in System and Policy Changes Now

America's population is aging rapidly, and neither our healthcare systems nor our social policy is responding with appropriate urgency. Millions of baby boomers are already turning 65, and their children and grandchildren – or ill-prepared safety-net services – are shouldering their care, with poor outcomes for both older adults and society at large. One reason for the lack of urgent action (among many) is that the insights gathered from existing research with older adults often remain siloed within commissioning organizations and agencies, rather than being disseminated to a coalition of engaged stakeholders as compelling data to inform policymaking and healthcare innovation nationwide.

By developing the right kind of public data, from the highest priority populations, while building a coalition of committed partners to use that data, The People Say hopes to shift the conversation on what needs to be done to build a society that works for older adults, their families, and all of us.

Project Components

To address the problems above, we formed a pool of older adults engaged with The People Say for ongoing human-centered research. We developed qualitative data on their preferences and experiences, then launched the platform – designed for use by advocates, policymakers and other stakeholders – to highlight actionable findings and insights about older adults' health and wellbeing needs.

Specifically, we've created:

1. **A research pool of older adults** – representative of all older Americans, but over-sampled on marginalized populations – with the intention and related infrastructure to be able to return to that pool regularly over time, both to conduct follow-on research related to this project and to address other specific research questions;

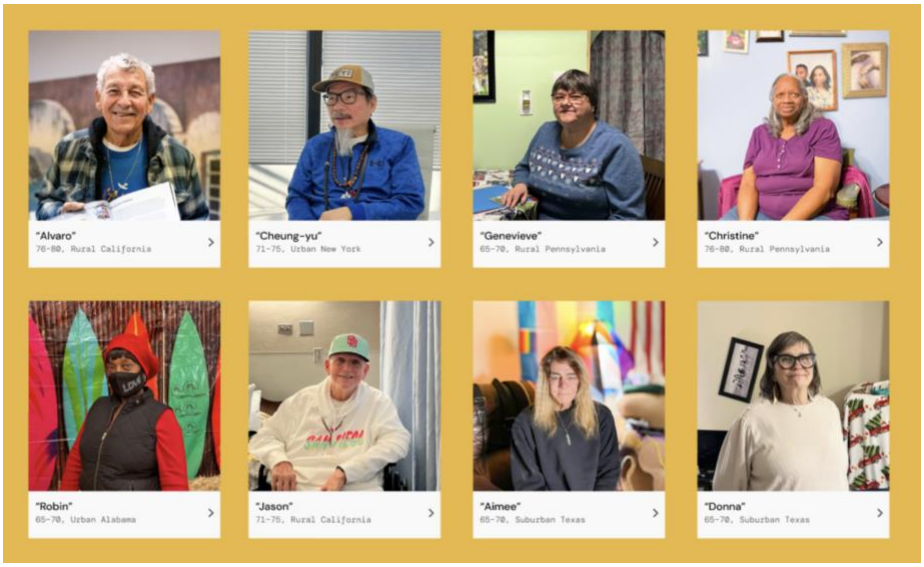

Fig. 1. Profile photos of eight members of the research pool. Image courtesy of the Public Policy Lab.

2. **A repository of tagged research data,** including both synthesized insights from research and direct quotes, transcripts, photographs, audio and video recordings, and other artifacts from research created by members of the pool of older adults, all categorized per a taxonomy developed by the research team; and

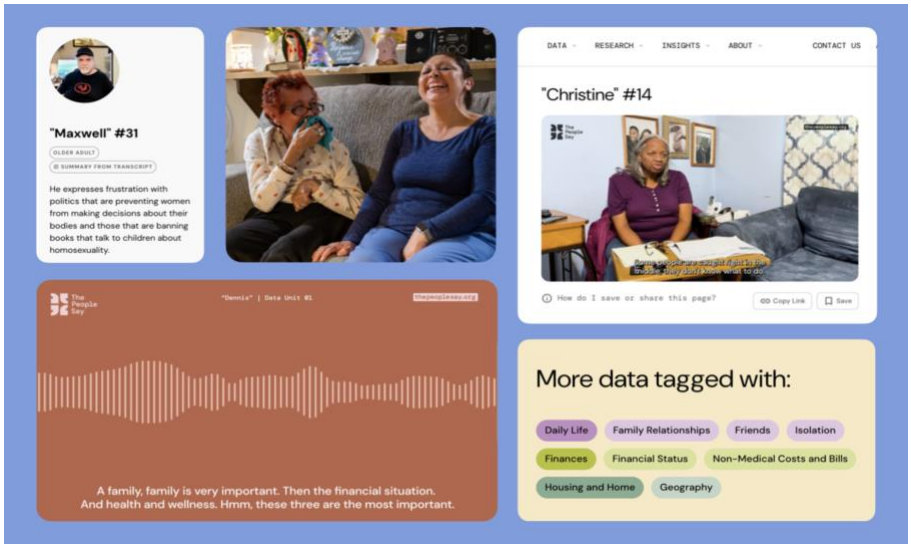

Fig. 2. An array of different types of qualitative data hosted on The People Say. Image courtesy of the Public Policy Lab.

3. The People Say, an online platform that publicly shares out findings and insights generated with those older adults regarding their healthcare access and delivery experiences, designed to highlight opportunities for national policy change and nationwide healthcare-systems improvement.

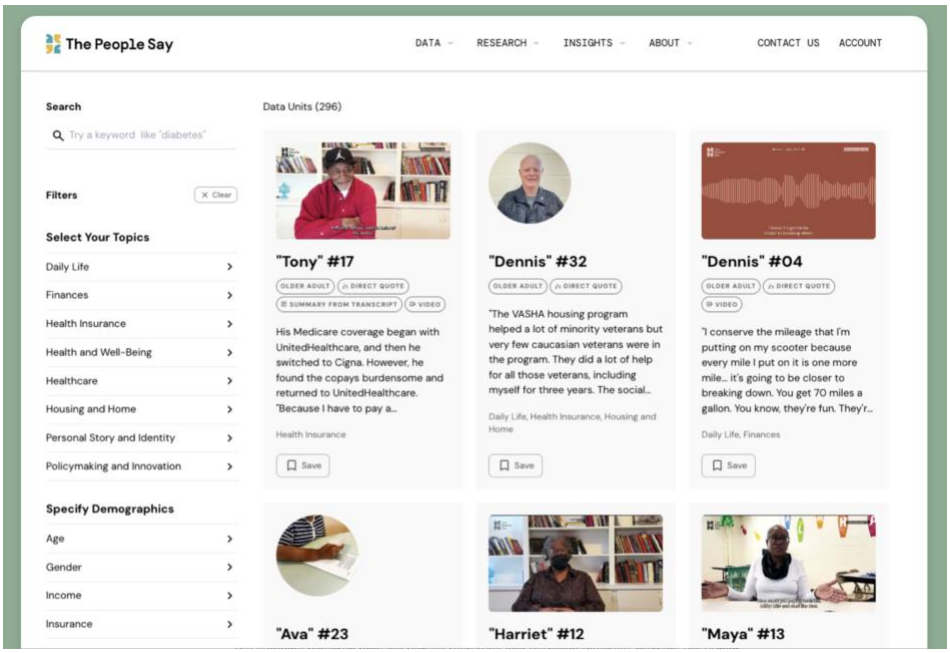

Fig. 3. A screenshot of the search page for The People Say. Image courtesy of the Public Policy Lab.

Project Goals

The overarching goal of The People Say is to support the development of policies, systems, and programs that are more reflective of the needs and preferences of older adults. In support of that goal, in the first year post-launch, we expect the project will have three short-term outcomes for project stakeholders:

More Meaningful Input from Older Adults

Creation of the standing research pool will reduce recruiting and outreach difficulty, allowing PPL, The SCAN Foundation, and partner organizations to more frequently and rapidly seek older adults' input and/or participation in goal-setting and decision-making. (Even before formal platform launch in July 2024, this goal began to come to fruition: PPL and TSF are using the research infrastructure to support the US Department of Health and Human Services' Administration for

Community Living in engaging with older adults to inform the creation of the decennial national Framework on Aging.)

Thicker Data Analysis and Storytelling

Development of the research repository will allow users of The People Say to better identify older adults' needs and preferences over time (as opposed to single-point analysis) and to convey that analysis in more human and compelling ways, given the easy availability of rich visual and audio assets associated with data points.

Broader Reach and Influence

Launch of The People Say will create new opportunities to advocate urgently for improvement in policies and services for older adults and to push for policy change that responds to older adults' lived experiences/needs as well as broader social-change goals. Project outputs may also be adopted by service providers, payers, and other stakeholders to contribute to momentum and evidence for larger policy change.

Project Activities

We conducted the project through six phases of work, from initial scoping and preparing, to research and synthesis, and then to co-design and launch. The project took just short of a year, from kick-off in September 2023 to public launch in July 2024. Activities undertaken over the course of the project are described below.

Preparing for Research

Sample Size

We planned a sample of 65 older adults, in a nod to the estimated population of 65 million Americans over the age of 65 in 2025. Ultimately, we engaged 66 older-adult participants, along with seven caregivers and 13 subject-matter experts, for a total of 86 participants.

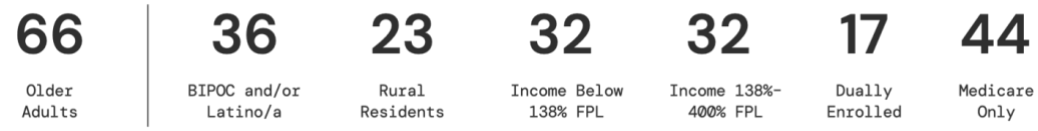

Fig. 4. Our 66 older-adult participants represent multiple racial/ethnic backgrounds, income groups, and geographies. Image courtesy of the Public Policy Lab.

Sample Selection

When initially setting recruiting targets for our sample, we started by identifying a baseline distribution of gender, racial/ethnic, geographic (urban/suburban vs. rural), and income demographics that mirrors the current 65+ population. Beyond seeking to highlight lives and aspirations of all older Americans, however, the Public Policy Lab and The SCAN Foundation share a commitment to focusing on the life experiences, needs, and preferences of populations that experience health disparities – particularly people of color, low-income people, and people who live in rural, medically underserved areas.

To deepen our research on these populations, we decided to oversample for participants who had those backgrounds or characteristics, with the aim of hitting levels equal or greater to the projections of the demographics of America in 2040. That forward-looking focus allows our participant pool to speak to the more diverse future that America is aging into, serving project interests in developing policy now that addresses the needs of the next 15 to 20 years.

See below for current (2020) demographics versus projected 2040 demographics, as well as the target and actual number of participants we recruited, reflective of those different distributions. Note that the racial/ethnic categories in the table below (and the terms used) are drawn from U.S. Census materials.

While a number of our respondents indicated that they were of Indigenous descent, none are enrolled members of a sovereign Tribal nation. In all other instances, we exceeded our targets for engaging lower-income, BIPOC, and rural participants.

Population ↓	2020/2022 Percent of Population	2040 Projected Percent of Population	n= per 2040 Figures	Final Participant Sample
All People 65+	15%	22%	65	66
Hispanic or Latino 65+ (all races)	8.79%	15.00%	≥10	12
Black or African American 65+	9.64%	11.72%	≥8	16
Asian 65+	4.61%	6.38%	≥4	6
Two or More Races 65+	0.94%	1.44%	≥1	2
American Indian & Alaska Native 65+	0.75%	1.03%	≥1	0
Native Hawaiian & Other Pacific Islander 65+	0.13%	0.21%	≥0	0
Rural 65+	20.00%	25.00%	≥16	23
65+ with Income <138% of FPL	17.2%	n/a	≥12	21
All People with Income <400% of FPL	58%	n/a	≥38	64

Fig. 5. Starting with projections of the 2040 65+ population, we then oversampled for priority populations. Image courtesy of the Public Policy Lab.

Research Locations

According to the US Dept. of Health and Human Services' Administration for Community Living, in 2020, 51% of Americans aged 65 and older lived in nine states: California, Florida, Texas, New York, Pennsylvania, Ohio, Illinois, North Carolina, and Michigan. We conducted in-depth human-centered research activities in five of those states – California, Texas, New York, Pennsylvania, and Ohio – as well as in Iowa and Alabama.

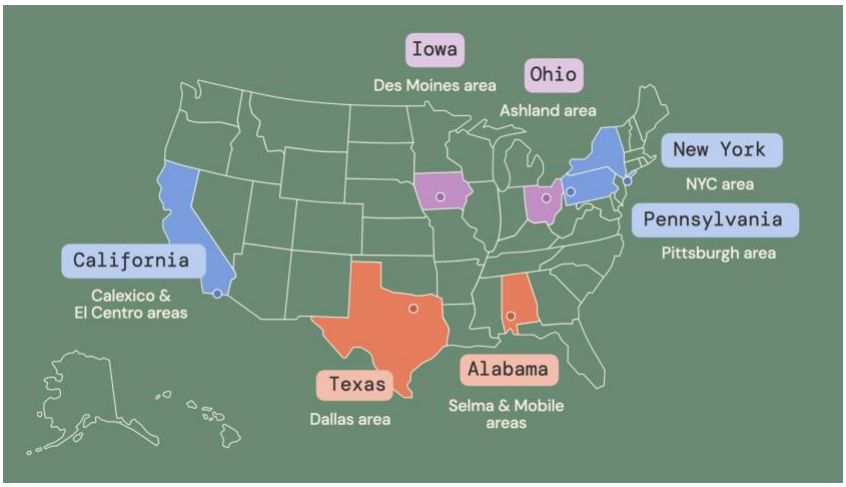

Fig. 6. Research took place in seven diverse communities across the country. Image courtesy of the Public Policy Lab.

The locations represent diverse geographic areas, population densities, political leanings, and U.S. Department of Health and Human Services (HHS) regions. We additionally selected four states in which more than 10.7% of the population is dually eligible for Medicare and Medicaid.

By concentrating research in a set of locations, rather than finding 66 participants in 66 different places, we sought to gather community-level trends, in addition to individual-level experiences. Ultimately, we were able to observe that some topics were more prevalent in each of our locations than in others.

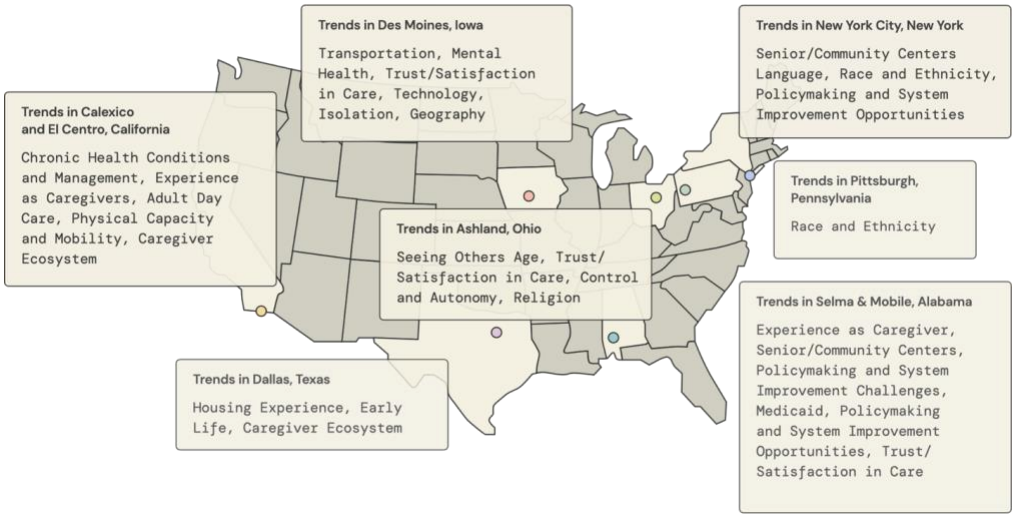

Fig. 7. For each location, some topic areas were more prevalent in that location than across the whole research pool. Image courtesy of the Public Policy Lab.

Inquiry Areas

Inquiry areas are broad topics or categories of knowledge that we hope to better understand through research. Inquiry areas apply to all participant or stakeholder types, and they are used to inform research methodologies, discussion guides, and desk research.

To develop our inquiry areas for the *Public Prospective Longitudinal Understanding Study of 65+ Adults*, PPL collaborated with The SCAN Foundation staff and also examined precedent surveys and studies of older adults. We identified four areas of interest that we believed could be further illuminated through qualitative human-centered design research (see illustration below). These inquiry areas guided our research engagements with older adults, their caregivers, and subject-matter experts:

Table 1. Inquiry Areas for the Public Prospective Longitudinal Understanding Study of 65+ Adults

Older Adults' Lived Experience	Policy and System Interactions
What are older adults' experiences, needs, and preferences around aging?	How do policies and healthcare systems currently respond to older adults' needs and preferences?
How do identity and culture relate to older adults' needs and preferences and the success of existing services and policies in serving them?	How could policies and healthcare systems change to support older adults better?

Research Recruiting

We recruited our cohort of 66 older adults through Area Agencies on Aging, local senior centers, and institutional and academic connections of the project's 25-member advisory committee. Recruiting materials were also distributed on social media and printed and posted at high-traffic areas at community-based organizations' physical locations and other community hubs.

As noted above, our aim was to assemble a research pool that over-represented people of color, lower-income people, and people who live in rural areas. This is a form of intentional sampling bias intended to elevate voices that are often underrepresented in policy decision making. As we had commitments to public sharing of data and to longitudinal research, we also made it explicit during recruiting that participants should be willing to allow us to share their data publicly and to participate in future research.

We were mindful of avoiding other unintended biases – for example, recruiting only individuals who are well connected to community networks or have better-than-typical literacy and technology skills. PPL's practice when working with community recruiting partners is to ask them to connect us to potential participants, but PPL then independently screens individual participants, ensuring that partners don't only provide us only with 'model' participants.

All participant-facing materials were written in plain language to ensure they were accessible to a wide range of participants. Key materials were also translated into Spanish and Chinese (both traditional and simplified) in order to reach participants whose primary language is not English.

Fig. 8. Flyers were produced in English, Spanish, and Chinese, and interviews took place in all three languages. Image courtesy of the Public Policy Lab.

Conducting Research

Semi-Structured Interviews

Pairs of researchers conducted semi-structured interviews in participants' homes or in agreed-upon community spaces, such as senior centers, churches, and libraries.

Before beginning, we walked participants through a comprehensive consent process that explained the project, described how the information collected would be used, and offered contact information for questions or concerns. While all PPL projects include a detailed consent process, this project required special care, given that we were asking participants to consent to public sharing of their data. While we had specifically recruited participants who were open to public sharing, a few participants did decide, during the consent process, that they were not willing to have their face and/or voice be recognizable in shared outputs; we either filmed them without capturing their face, or we subsequently blurred or disguised their visage and/or voice in material posted to the public platform.

Our in-depth interviews focused on older adults' experiences of aging and accessing healthcare and other services. Our goal in semi-structured interviews was to use our inquiry areas as a guide, but not to follow a predetermined script or matrix

of questions. We aim to create an environment where participants feel comfortable, allowing our professional qualitative researchers to explore interesting stories as they emerge during the engagement. We carried lightweight equipment – just a tripod, light, and small microphones – to minimize the intrusiveness of recording.

Teams traveled directly to participants for these in-depth conversations, leaving behind a packet of independent activities to be completed over the following weeks and returned by mail. Our researchers conducted interviews in English, Spanish, or Cantonese, and activity materials were provided in the participants' preferred language.

Fig. 9. A PPL researcher conducts an interview with an older adult. Image courtesy of the Public Policy Lab.

Cultural Probes

The first activity we provided for participants to complete on their own was called *Who I See*. It asked them to map their social network by filling out worksheets about the people they see on a daily, weekly, or monthly basis, including friends, family, and care staff, and also to describe what those interactions are like and how they feel.

The second activity, *My Life Timeline*, prompted older adults to chart their lifetime milestones of aging – along with their hopes for the future – and to provide more detailed explanations about a few of those moments.

Almost all of our participants returned these materials to the research team, and they are fascinating windows into people's past experiences and current social networks. We have included scans of the timelines on the public platform, scrubbed of personally identifying information. The social-network probe artifacts proved too challenging to make publicly available, given references to names and details of non-consented members of our participants' community; these materials were used to inform our insight briefs, however, and could be further analyzed in future rounds of research.

Project Materials Reviews

All older-adult participants were provided with copies of all their data units, in their original language, in advance of site launch. We asked them to indicate if there were any items that they did not want to have published. Based on this review, we removed three units of data, from one participant. All of the nearly 2,400 data units remaining on the site have been approved for public sharing by the participants.

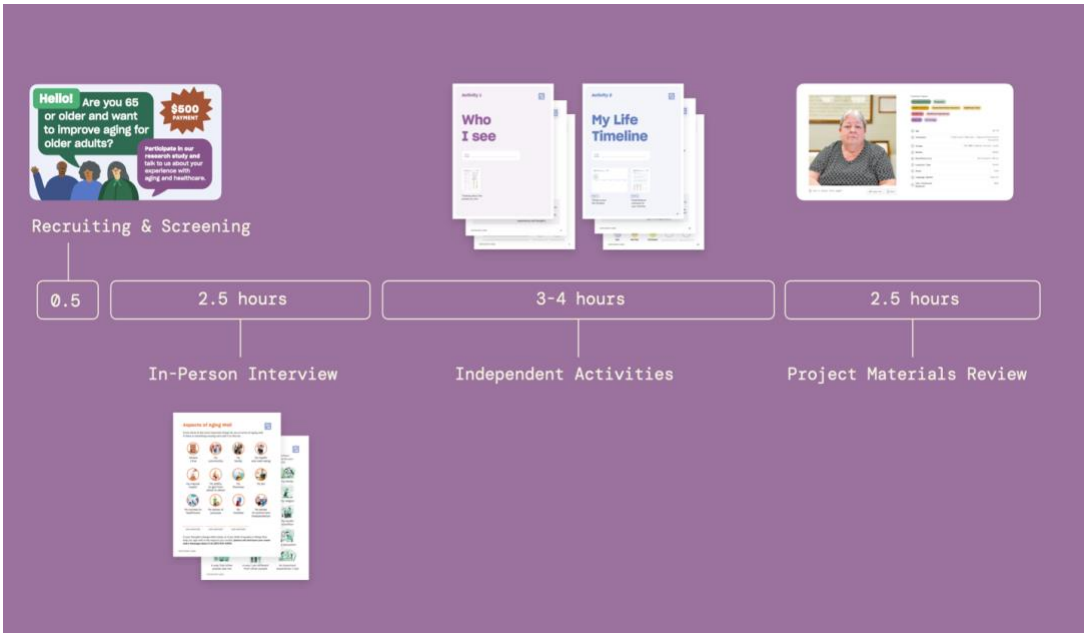

Fig. 10. Participants sat for interviews, completed independent worksheets, and reviewed site data. Image courtesy of the Public Policy Lab.

Research Compensation

Older-adult and caregiver participants who took part in all activities provided about ten hours of their time. Participants received $200 for the initial interview and received up to a total of $500 if they participated in follow-on activities. Expert members of our advisory committee were also asked for about 10 hours of time and were offered an equivalent honorarium (which many declined). Our $50-per-hour compensation rate is intended to match or exceed the rate of pay of an entry-level PPL staff researcher.

Building the Platform

Taxonomy and Tagging

A primary goal of The People Say is to provide access to a qualitative dataset that enriches and expands on the quantitative data about older adults that's typically referenced by policymakers. To make our data more directly associable with existing quantitative data, our 108 topic and subtopic tags were developed with reference to the primary keywords or taxonomies used by five major longitudinal studies of American older adults, specifically:

- The Health and Retirement Survey (HRS) – Aging in the 21st Century: Challenges and Opportunities for Americans
- Medicare Current Beneficiary Survey (MCBS)
- National Health and Aging Trends Study (NHATS)
- Older Americans 2020: Key Indicators of Well-Being
- 2023 Profile of Older Americans

With those studies as a starting point, we developed a taxonomy to categorize the content included on The People Say. Every data unit included in our database is tagged per this taxonomy. You can filter our database by these tags when you search the database; additionally, on any given data-unit page, clicking on any topic or subtopic tag will generate a result of all data on this site with that tag.

To make it easy to find quantitative data that augment our qualitative data, on each one of our data units, we've also provided links to related data in summary reports from the sources above (except the MCBS, which doesn't have an easy-access summary report).

Approximately 40% of our subtopic tags do not have related data in the quantitative surveys. We think this speaks to the meaningful difference in the kind of data that in-person, human-centered research can generate. Our participants shared aspects of their lived experience that survey instruments are unlikely to collect.

Development and Testing

The People Say has been designed primarily for three major user groups: 1) legislators, policymakers, and their staff; 2) federal, state, and local program leaders; and 3) researchers and advocates. Our 25-member advisory committee includes representatives of all of these user types, and they helped us iterate on the tagging taxonomy, even before we began site development. As an example, federal policymakers and program leadership felt that it was important to be quite specific in tagging data related to different various federally funded benefits programs, so we have specific tags for Medicare, Medicaid, veterans' programs, etc. Once in development, we also tested with users recruited from the workplaces of our advisory committee members at two different stages of development.

While focused on those three primary user groups, we've also designed the site to be usable by older adults and anyone seeking to learn how to better support their aging loved ones and community members – and what to expect when growing older themselves. While we did less user testing for these user groups, we did conduct sessions with several of our older-adult participants who use websites (only about a quarter of our participants are comfortable using online tools). We tested their ability to navigate the site, review their data, and also review that of others; this user testing, as expected, provided some useful additional information around accessibility and ease of use.

Launch and Immediate Impact

Launch Event

We launched the site on July 11, 2024, at the U.S. Capitol – an opportunity provided to us based on the relationship between a partner at The SCAN Foundation and the Senate Committee on the Aging. In attendance were staff from two House and Senate committees, member staff from the offices of three senators and four representatives (five Democrats and two Republicans), representatives from

six federal agencies, and staff from 22 research, advocacy, and stakeholder organizations.

We were also joined by two of our research participants and their family members – from Iowa and Pennsylvania – and we conducted a Q&A with them about their stories, their participation in the project, and what they hoped policymakers and program leaders would take away from hearing about their experiences of aging where they live.

Fig. 11. The launch event for The People Say was held on July 11, 2024, at the US Capitol Building. Image courtesy of McCabe Message Partners.

Immediate Impact

Reaction to the launch has been very positive – and it's already having an impact on efforts by the federal government to better incorporate older adults' points of view in policy decision making. The Interagency Coordinating Committee on Healthy Aging and Age-Friendly Communities (ICC), established by the authority of the Older Americans Act, is developing a national framework on aging that can support older Americans in aging in place, while accessing preventive healthcare and long-term services and supports. A key goal of the ICC is to hear directly from older adults – particularly those with the greatest economic and social need – to ensure policies are reflective of needs and preferences.

Working with partners at The SCAN Foundation, a National Plan on Aging Community Engagement Collaborative partner with the ICC, PPL is currently gathering this input during the summer of 2024 by facilitating listening sessions that amplify the voices and lived experiences of older Americans. We're organizing listening sessions in three of the communities where we established relationships during the creation of The People Say, and existing members of our participant pool have been invited to join those sessions. During the listening sessions, older adults will meet with federal officials and discuss aging and the four domains in the national framework: Age-Friendly Communities, Coordinated Housing and Supportive Services, Increased Access to Long-Term Services and Supports, and Aligned Health Care and Supportive Services.

We think that this example – of policymakers' direct engagement with older adults' experiences – points to the immediate value of constituting ongoing research pools for civic research.

Lessons Learned

Development of The People Say has been both immensely challenging – and also terrifically gratifying. We believe that lessons learned during this project will be very helpful in developing similar subsequent work.

Unanticipated Challenges

The project was scoped based on our past experience conducting national-scale qualitative civic research. But even that past experience did not prevent us from experiencing a whole range of challenges, especially relating to the multimedia and public aspects of this project. With hindsight, many of these difficulties seem obvious! However, we apparently had to learn the following lessons the hard way:

Recruitment Partners Could Not Substitute for Team Recruiting in-Person

Recruiting is often time-consuming, but we'd anticipated that warm connections from our advisory committee to on-the-ground organizations that serve older adults would be valuable in connecting us to participants. However, those introductions did not always yield capacity or willingness to help us recruit. Cold outreach to community organizations was less than reliable because most senior centers had no

reason to believe in our legitimacy over the phone, and they typically did not respond to email. Ultimately, we found that in-person intercept recruiting at sites was more valuable than attempting to leverage local staff to recruit.

Data Collection Overwhelmed Data Processing

We have a standard notetaking protocol. Nonetheless, researchers who did not rigorously enter data and quotes into our database as soon as possible post-interview eventually fell far behind in preparing this data for public presentation. If we had been able to stick to our original plan to alternate field-research weeks with stay-at-home weeks, we might have fared better – but that intention collapsed under staff and site scheduling challenges.

Translation Was an Enormous Effort – and Robots Created Problems

Our team members fluent in Spanish and Cantonese translated discussion guides, interview materials, and subsequent outreach materials. These were more time-intensive efforts than anticipated, but still feasible. However, the generation of transcripts, video subtitles, and textual data units for interviews conducted in Spanish and Cantonese was a much larger task than anticipated. We remained committed to presenting all these materials on the site in the original language as well as in English, to preserve their legibility to the participant and other users. We hoped to rely on AI-enabled transcription to do the translation heavy lifting, but AI tools introduced transcription errors that had to be corrected manually – in both languages – for dozens of hours of interview material.

Habits From Non-Filmed Research Created Problems in Filmed Data

We're used to building rapport during interviews by speaking conversationally, addressing people by name, etc. These habits have minimal negative effects when using and scrubbing written transcripts, but proved more problematic in compiling our filmed data. We talked over our participants at times, and both we and participants used personally identifiable information – all of which had to be cleaned up in our final video artifacts.

Film Quality Undershot Our Aspirations

As this project was both an experiment in collecting significant amounts of video data, and also a fast-paced effort with multiple simultaneous teams in the field, we erred on the side of lightweight, low-cost filming using mobile phones. That decision, while expedient, led to predictable variations in video quality. Overall, while we're happy with the research value generated by our video artifacts, we've been disabused of any notion that we might create satisfying artistic filmmaking with amateur tools and technique.

The Ethics of Public Data Creates New Requirements

Preparing multimedia data for a public platform is very different than collecting it for internal use: Each transcript (and its translation, when the original engagement was conducted in a language other than English), as well as all audio/video clips and the timeline cultural probes, had to be thoroughly scrubbed of PII before they could be added to the public platform. Additionally, we decided to print out all data units and ship them to participants for their review and approval in advance of launch – this was not a commitment we made when consenting our participants, but the research team decided it was further valuable assurance that we were not exposing data that any participant would rather not have made public.

As Always, Human Alignment Is the Most Challenging Work Product

We *know* that the hardest work of any project is aligning the mindsets, expectations, and work processes of the partners and the teams involved. And yet: we underestimated the necessary rounds of multi-partner approvals, failed to get our own internal team in advance agreement over the final production and soft-launch timeline, and complicated our software development process with fuzzy feature scoping. Every future project is a chance to not do that again!

Unexpected Successes

Despite the challenges, the work also yielded some unexpected satisfactions.

Outputs Blossomed Past Original Intentions

We'd assured our partners that we'd be able to generate some insight artifacts as outputs of research synthesis. We imagined a set of static materials capturing the

collective experiences of our participants. As we began to develop a set of insight papers, it became easy to iterate on how they might live on the website: they could preface a collection of video clips sorted into subthemes; they could be complemented by a section on the relevance of policy to affecting that insight area. These insight pages have become a primary way to invite both users and casual browsers to draw meaning from the data and to become invested in what it could mean and motivate.

Rich Media Is Super Compelling

We'd hoped that the investment in creating video research would yield greater engagement in participants' needs and aspirations – and so the powerful response generated by the video data has been heartening. We have had many requests to play and use the clips even since we shared a first sampling of them in February 2024.

Listening Is Transformative

Many of our older-adult participants expressed how validating it was to participate in this work, to be asked to tell their stories and to have their stories heard. And the work was also deeply gratifying and even life-changing for our researchers. We got Christmas cards and Lunar New Year's greetings from our participants. We got hugs. And we got the benefits of their decades of life experience, with profound learnings both on what makes the last third of life painful – and what creates joy and meaning.

What's Next

This project is intended to serve as a test case for public multimedia civic research – it's our intention to build on the lessons above by expanding the platform in a variety of ways.

Longitudinal and Expanded Research with Older Adults

An important next step for this work is to conduct subsequent rounds of research with our participants. We believe that tracking the evolving needs of older adults over time will generate important insights into persistent and/or recurrent challenges that can be addressed by policymakers.

We're also interested in expanding our research pool, both to add additional geographic locations and to include additional populations of interest, such as older adults living in residential care settings, older adults in need of at-home care, older adults who are/have been homeless, Indigenous older adults, and/or formerly incarcerated older adults.

Public Civic Research with Other Populations

We developed The People Say with a focus on older adults – but also to demonstrate the potential for public, prospective, longitudinal civic research.

We're interested in applying the same approach to developing other longitudinal research studies with populations of interest, such as moms (re: prenatal, birth, and post-partum experiences and services) and working families (re: access to healthcare/benefits, job quality, education/training opportunities, social connection, and other aspects of economic mobility), or in other specific geographies, such as New York City or Michigan.

Extending the Platform to Other Civic Researchers

In addition to using the platform to house research conducted by the Public Policy Lab or for The SCAN Foundation, we also want to support other civic researchers with a commitment to high-quality, ethical design research to make use of the platform, creating a kind of 'GitHub for civic research' – a repository for the work of multiple civic researchers. Options we're interested in exploring include both adding other research data to the existing The People Say data set and creating new public data sets that use the same platform infrastructure.

Conclusion

While we have further aspirations, we're pleased that The People Say demonstrates a meaningful effort to address the major limitations in civic research practice that we outlined in our introduction above:

- Creation of our standing respondent pool, comprised of a deliberately inclusive group of participants, is already allowing us to more easily and quickly recruit diverse members of the public to participate in civic research, as in the ICC listening session example above.

- Partner organizations are expressing interest in supporting ongoing, longitudinal research with the participants, rather than the one-and-done model that's typically pursued.
- We've demonstrated that it's viable to conduct civic research that respects participant informed consent and agency while generating public outputs.
- We've established a functional model for that participation to be paid – and for compensation of all participation to be equitable.
- We've shown that civic research doesn't have to be proprietary – it's not just possible, but valuable to put civic research data into the public domain. We hope this is a first step to reducing the need to re-study these exact same issues with the same populations.
- Members of the public who participated in the project have visibility into what we collected from them and how their data has been tagged and used in aggregated outputs.

Our ultimate goal, of course, is that our civic research serves to highlight the insight of (often marginalized) people and to shift systems to meet their needs and aspirations. We look forward to building on this precedent to meet that goal.

About the Authors

Chelsea Mauldin (she/her) is a social scientist and designer with a focus on government innovation. She directs the Public Policy Lab, a nonprofit organization that designs better public policy with low-income and marginalized Americans.

Raj Kottamasu (he/him) is a designer, researcher, and strategist balancing priorities for equity, service, and relationships with systematic approaches to collaboration. At the Public Policy Lab, he leads projects and helps to guide design and strategy across the organization. He served as the research lead for The People Say.

Petey Routzahn (he/him) is a hands-on, brains-in designer who creates physical and digital artifacts that help people better understand the world around them and navigate the built environment. He served as the project lead for The People Say.

Acknowledgments

The authors would like to acknowledge the thoughtful feedback of Victoria Lowerson Bredow, our EPIC curator, as well as the many people and organizations that made this project possible, including: our partners at The SCAN Foundation, Narda Ipakchi and Kali Peterson; the rest of the PPL project

team and our tireless corps of video translators and editors; our expert advisory committee; staff at community-based organizations across the United States who provided support during our research fieldwork; Greater Good Studio; McCabe Message Partners; and our web-development partners at Make Repeat. And, especially, our older-adult and caregiver participants – it was an honor to share their insights.

References Cited

Administration for Community Living, U.S. Department of Health and Human Services. 2024. *2023 Profile of Older Americans*. https://acl.gov/sites/default/files/Profile of OA/ACL_ProfileOlderAmericans2023_508.pdf

Centers for Medicare & Medicaid Services Office of Enterprise Data and Analytics. 2023. *Medicare Current Beneficiary Survey*. https://www.cms.gov/data-research/research/medicare-current-beneficiary-survey

Federal Interagency Forum on Aging-Related Statistics. 2020. *Older Americans 2020: Key indicators of well-being*. U.S. Government Printing Office. https://agingstats.gov/docs/LatestReport/OA20_508_10142020.pdf

Freedman VA, and JC Cornman. 2023. *National Health and Aging Trends Study Trends Chart Book: Key Trends, Measures and Detailed Tables*. 2011–2021. https://micda.isr.umich.edu/wp-content/uploads/2023/07/NHATS-Companion-Chartbook-to-Trends-Dashboards-2021.pdf

University of Michigan Health and Retirement Study. 2017. Aging in the 21st Century: Challenges and Opportunities for Americans. https://hrsonline.isr.umich.edu/sitedocs/databook/inc/pdf/HRS-Aging-in-the-21St-Century.pdf

The Transformative Potential of Participatory Research in International Development Programming

OLUCHI AUDU, *YUX*

In 2022, UNICEF launched the Adolescent Girls Programme Strategy to enhance programming for adolescent girls, focusing on their well-being and empowerment. In Northern Nigeria, the Reaching and Empowering Adolescent Girls in Northwest Nigeria program was introduced in Sokoto and Katsina states. YUX conducted formative research to develop Social and Behavior Change strategies addressing harmful social and gender norms impacting girls' education. The study found that entrenched patriarchal norms and practices such as child marriage and purdah significantly hinder girls' educational access. Socioeconomic factors, including poverty and inadequate school facilities, further exacerbate these challenges. Girls expressed a strong desire for education to benefit both themselves and their communities. Through co-creation workshops, the girls proposed strategies like engaging fathers and community leaders to advocate for girls' education. The resulting strategies emphasize community involvement and targeted messaging to challenge harmful norms. UNICEF Nigeria is now implementing these strategies, ensuring they reflect the girls' needs and aspirations while promoting sustainable change. This case study details the significant value that a participatory research approach can bring to international development programming.

Context

Conducting Formative Research on Social and Gender Norms with and for Adolescent Girls

In 2022, the United Nations International Children's Emergency Fund (UNICEF) developed a new vision for programming with and for adolescent girls: the Adolescent Girls Programme Strategy (AGPS) to accelerate action so girls are not left behind. Together with girls and other partners and building on its existing work, UNICEF is promoting evidence-informed, holistic action to deliver more deliberate, girl-centered programming—placing adolescent girls' rights, well-being, voice, agency, and leadership at the very core of what it does.

UNICEF Nigeria was one of 11 initial Country Offices selected to roll out the AGPS and is implementing the Reaching and Empowering Adolescent Girls in Northwest Nigeria (REACH) programme in Sokoto and Katsina states. To support the REACH outcomes, YUX – a pan-African research and design agency – was contracted by the UNICEF Nigeria country office to conduct a formative study on the gender and social norms that affect adolescent girls in Sokoto and Katsina States.

Audu, Oluchi. "The Transformative Potential of Participatory Research in International Development Programming." *EPIC Proceedings* (2024): 435–464. https://www.epicpeople.org/transformative-potential-of-participatory-research-in-international development. ISSN 1559-8918.

The goal of this research was to develop evidence-based Social and Behavior Change (SBC) strategies to address harmful norms impacting various aspects of adolescent girls' lives in Northern Nigeria and to promote their empowerment. This initiative took a holistic approach, encompassing education, digital and financial inclusion, child protection and well-being, and health-seeking behavior. We aimed to understand how entrenched social and gender norms affect these areas and to create intervention plans to address them.

This case study centers on the educational experiences of adolescent girls in Nigeria, emphasizing the significant role of participatory methodologies in uncovering the norms that shape their lives. Through a participatory approach, we not only gained invaluable insights into the social and gender norms impacting these girls but also empowered them to actively contribute to the research. The case study begins with a review of existing data on the influence of these norms on education, followed by findings from our formative research. It then highlights the intervention strategies developed by the girls themselves during a co-creation workshop, showcasing the unique value of their perspectives in addressing the norms affecting their education. Finally, it presents the draft SBC strategy, curated to counter these harmful norms and promote the girls' educational advancement.

Background: The Influence of Social and Gender Norms on the Education of Adolescent Girls in Northern Nigeria

Socio-cultural and gender norms have a significant impact on adolescent girls in the Northern region of Nigeria, shaping their behaviors, opportunities, and overall well-being. The region has a large Muslim population, and both Islamic teachings and its interpretations as well as local cultural practices play a significant role in shaping the lives of girls. One key aspect of the socio-cultural context in Northern Nigeria is the prevalence of patriarchal norms and gender inequality. Traditional gender roles often assign women and girls subordinate positions e.g., caregivers, home makers etc. within the family and society, limiting their access to education, decision-making power, and opportunities for personal and professional growth (Bolarinwa et al., 2022).

Nigeria faces a significant challenge regarding low school attendance, with approximately 10.5 million school-age children not receiving an education. This issue is particularly pronounced in the northern regions of the country, where more than half of female children lack access to schooling (UNICEF Nigeria, 2019). In

Northern Nigeria, 29% of females between the ages of 15-49 are unable to read and write (National Bureau of Statistics and United Nations Children's Fund, 2021). Northwestern Nigeria also holds the unfortunate distinction of having the highest rate of female illiteracy among those aged 15 to 21 in the entire country, with a staggering figure of 70.9% (Nielsen, 2021). The prevailing gender norms in Northern Nigeria contribute to this educational disparity, as communities in Northern Nigeria generally prioritize boys' education over girls'. Girls are typically expected to prioritize domestic duties, including household chores and caregiving, while boys are encouraged to pursue education and careers (Odimegwu et al., 2017).

The prevalence of child marriage among girls in Northern Nigeria also plays a role in lower levels of education access amongst girls in Northern Nigeria. This region has one of the highest rates of child marriage in Nigeria and this has remained a significant issue despite efforts to address it. According to the Nigeria Demographic and Health Survey (NDHS) in 2019, 49% of women aged 20-49 in the region were married before the age of 18 (National Population Commission, 2019). Child marriage emerges as one of the primary factors leading to these high dropout rates, as girls often find themselves forced to leave school prematurely due to this practice (Save the children, 2021).

Furthermore, Purdah, a practice observed in some Muslim and Hindu communities, entails the seclusion of women from public observation (Yusuf, 2014). This cultural and religious tradition requires women to stay out of the public eye, often by remaining within their homes or wearing clothing that covers their bodies and faces when outside. Purdah significantly limits girls' access to education in several ways. Firstly, it often confines girls to their homes, preventing them from physically attending school and participating in a standard educational environment. Additionally, cultural norms discourage or prohibit girls from traveling alone, limiting their transportation options and making consistent school attendance difficult.

In addition, Purdah restricts access to educational resources outside the classroom. Girls miss out on opportunities to visit libraries, community centers, and other institutions that could supplement their learning. They are also less likely to engage in extracurricular activities, or social gatherings that contribute to a well-rounded education. Lastly, Purdah perpetuates a mindset that undervalues girls' education, discouraging families from investing in their daughters' schooling. This leads to lower enrollment rates, higher dropout rates, and limited academic achievement for girls (Yewande & Olawunmi, 2023).

Research Motivation

Given the challenging circumstances and the detrimental impact of social and gender norms on the education of adolescent girls in Northern Nigeria, there was a pressing need to further understand these norms and their effects from the girls' own perspectives. Identifying effective Social and Behavior Change (SBC) strategies to address these harmful norms and empower adolescent girls also became crucial. In response, UNICEF Nigeria country office partnered with the YUX team to conduct formative research in Sokoto and Katsina states. The primary objective was to offer insights that would guide UNICEF Nigeria's programming, particularly in developing strategies to support initiatives for adolescent girls. Our approach was designed to deeply understand the needs and aspirations of the adolescent girls in Northern Nigeria and to gather their perspectives on how they believe they can achieve their goals.

Research Design and Methodology

Capturing the Narratives of Adolescent Girls from Their Perspectives Using a Participatory Research Approach

The initial step in developing evidence-based Social and Behavior Change (SBC) intervention strategies to address harmful norms impacting the education of adolescent girls in Northern Nigeria was to first collect and analyze relevant evidence. Most existing literature on gender and social norms in Northern Nigeria often fails to directly engage with the girls themselves or capture their personal narratives. Therefore, it was crucial for us to ensure that we amplified the voices of these adolescent girls, providing them with the opportunity to share their stories in their own words.

In tackling the complex issues related to gender norms in Northern Nigeria, we acknowledged the importance of not only capturing the perspectives of the girls themselves but also involving them directly in the research process. Hence, we adopted a participatory research methodology, grounded in the belief that involving young women from the communities where the research was conducted would generate more insightful and meaningful results. Rather than conducting research solely for these girls, we aimed to conduct it with them, recognizing their perspectives as invaluable in shaping effective interventions. We decided to engage a cohort of 20 young women (aged 18-27) in each state as integral members of our research team. We collaborated with UNICEF Nigeria to enlist these young women.

Prior to them being fully integrated into our research team, we conducted interactive training sessions in both English and Hausa (the native language of the regions). These sessions included skill-building in Human-Centered Design (HCD), interviewing techniques, and data analysis. After the training sessions concluded, the girls were given official certificates from the YUX Academy to acknowledge their dedication and also potentially enhance their professional standing.

This approach not only equipped the 20 young women with practical skills but also positioned them as active contributors to the research process. They were not just subjects of study but rather partners in designing solutions that could impact their own communities. A participatory approach also meant a more comfortable environment for the adolescent girls being interviewed and surveyed. The peer-to-peer dynamic would foster trust and openness, mitigating traditional power imbalances often present in interviewer-interviewee relationships. This environment was crucial for eliciting candid responses and deep insights into how gender norms affect their lives daily. Furthermore, by involving these young women in the research process, we ensured that the interventions and strategies developed were more relevant and reflective of their realities and aspirations. Their firsthand experiences and perspectives provided nuanced understanding, guiding the development of targeted interventions that could effectively challenge and reshape harmful social and gender norms.

Fig 1: The young women during the training sessions

The training sessions were conducted over a single day, bringing together young women to learn and prepare for the field research they were about to undertake. For most of these young women, this training marked their first introduction to conducting research and mastering interviewing techniques. To ensure the material was accessible and relatable, it was designed in a straightforward and literal manner, drawing parallels to various aspects of their daily lives to explain the research process clearly. The sessions were conducted in their local language, Hausa, to facilitate better understanding and engagement. To create a comfortable learning environment and encourage active participation, the girls were divided into smaller groups. This grouping not only made the sessions more intimate and less intimidating but also allowed the participants to ask more questions and engage in discussions freely. Additionally, pilot studies were conducted to provide practical experience and reinforce the training. These pilots enabled the young women to apply what they had learned in a real-world context, further building their confidence and skills in conducting research and interviews.

Before going into the field, we started off the initial phase of our formative research by conducting a literature review to assess existing data and understand the current landscape regarding harmful social and gender norms affecting the education of adolescent girls in Northern Nigeria. This review helped us identify gaps in the literature, which informed the development of our questions for both quantitative and qualitative research. Additionally, we interviewed six key stakeholders from various teams within UNICEF to gain a clearer understanding of the organization's goals and their implications.

With this preliminary data in hand, we prepared for fieldwork by drafting discussion guides and survey questions, planning logistics, identifying locations for speaking with adolescent girls, and training the 20 young women who would be part of our research team. We presented the initial drafts of the discussion guides and survey questions to our young women researchers, collaborating with them to refine the questions to ensure they were more appropriately tailored for their peers.

After completing our preparation, we conducted field research, beginning in Katsina state before moving on to our second planned location, Sokoto. The research commenced with a total of 20 in-depth interviews with adolescent girls (10 interviews in each state), aged 10 to 19, categorized into three groups: young adolescents (10-13), mid-adolescents (14-16), and older adolescents (17-19). To ensure a diverse range of perspectives, we interviewed girls and young women from various Local Government Areas (LGAs) in both states: Rimi and Mani in Katsina,

and Dange-Shuni and Sokoto North in Sokoto. Our sample included girls from different social contexts, such as pregnant girls, young mothers, married girls, in-school girls, and out-of-school girls.

Each qualitative interview lasted approximately 60 minutes and was primarily conducted by young women researchers, with support from YUX local researchers. The interviews focused on various aspects of the girls' lives, examining how harmful social and gender norms impact them. Additionally, the questions explored the strategies the girls believed would be effective in addressing these norms and empowering themselves.

In the qualitative interviews, we aimed to understand gender and social norms by engaging with not only the adolescent girls but also other key figures in their lives. This included parents or guardians, male siblings, traditional and community leaders, and religious figures. We conducted 8 interviews with caregivers, 4 with male siblings, and 8 with community stakeholders, including religious leaders, government officials, and community leaders. Each interview lasted approximately 60 minutes and was conducted by YUX's local researchers.

Once the qualitative interviews were completed, we proceeded to gather quantitative data using a survey developed in collaboration with the UNICEF team and our young women researchers. Equipped with tablets provided by UNICEF, the researchers collected responses under the guidance of YUX local researchers. The YUX team translated the survey questions into Hausa. Our young women researchers conducted the surveys in person and entered the data using SurveyCTO. We had a total of 548 responses from the adolescent girls in both states.

Through this comprehensive ethnographic approach, we integrated the voices of adolescent girls and their communities, ensuring our Social and Behavior Change (SBC) strategies were evidence-based, culturally relevant, and grounded in the lived experiences of those most affected by harmful social and gender norms. By combining qualitative and quantitative methods, we gained a deeper understanding of the issues. We began with qualitative research to capture the nuanced and contextual insights from the girls and their communities. This allowed us to design more effective quantitative surveys that were statistically relevant across the region. The qualitative phase informed our understanding, helping to tailor our questions and ensuring our quantitative data accurately reflected the complex realities on the ground. This method ensured our strategies were both informed by rich, detailed narratives and supported by robust statistical evidence, creating a comprehensive foundation for impactful SBC initiatives.

By actively involving adolescent girls and their communities in the research process, we aimed to foster a sense of ownership and empowerment among the

participants. This approach was to ensure that the insights we gathered were not only accurate but also deeply reflective of the participants' true experiences and perspectives. Engaging the community in this way would also build trust and rapport, which is crucial for gathering honest and open feedback. Furthermore, the participatory approach would help to identify and address potential biases and blind spots in the research process. The girls and their communities could highlight issues and cultural nuances that external researchers might overlook. This collaborative effort enriched our understanding and ensured that the SBC strategies we developed were not only theoretically sound but also practically feasible and culturally sensitive.

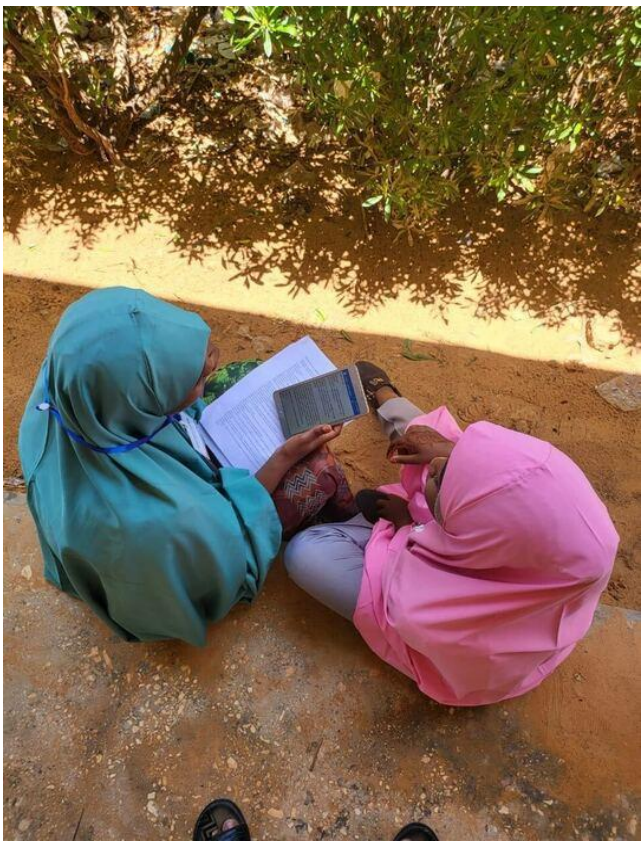

Fig 2: A Young Woman Researcher Administering a Quantitative Survey to an Adolescent Girl

Facilitating Co-creation Workshops to Build SBC Strategy Drafts

Having gained a deeper understanding of the detrimental impacts of these social and gender norms through our comprehensive data analysis, we facilitated co-design workshops. These workshops provided a platform to share our findings with the girls and other stakeholders, fostering collaborative brainstorming sessions to

generate ideas for interventions. We had a total of 4 workshops, 2 in Katsina state and 2 in Sokoto state. Each workshop day was broken into two parts, the main workshops which included the adolescent girls and the second workshop which included the other stakeholders.

Each workshop grouped the adolescent girls into four age-based groups, with facilitation provided by our team of young women researchers, who had received additional training in workshop facilitation prior to the sessions. While these researchers led the workshops with the girls, YUX local researchers facilitated sessions with other stakeholders, including 2 male siblings, 4 caregivers (2 male and 2 female), and 4 community stakeholders. The girls' workshops lasted approximately 7 hours with various breaks, while the sessions with other stakeholders were 2 hours each.

The workshop processes incorporated several ethnographic tools and methods. We utilized "How Might We" (HMW) questions for brainstorming, prioritizing these questions to focus on key issues, and organized participants into groups to facilitate collaborative brainstorming. Voting on prioritized ideas empowered participants in decision-making, and concept prototyping and storyboarding were employed to visualize and refine ideas. This iterative and participatory approach allowed for rich qualitative data collection and ensured that solutions were deeply rooted in the participants' lived experiences and cultural contexts.

Drawing insights from both the field research on norms and impacts and the creative input from the girls during the co-design workshops, we formulated Social and Behavioral Change strategies. These strategies are grounded in understanding various norms and their underlying drivers, aiming to empower adolescent girls within the community by addressing their primary concerns and aspirations. The workshops' methods ensured that the solutions were culturally relevant, inclusive, and adaptable, aligning with ethnographic principles of participant-centered design and ongoing adaptation.

Findings

Perspectives on Education from Adolescent Girls in Northern Nigeria

The girls who were interviewed shared profound insights into the significant values they associate with attending school. They emphasized their desire to contribute positively to society, gain knowledge to support and educate their siblings and future children, and enhance their English language proficiency. Recognizing that Hausa is the widely spoken language in their community, the girls understand the value of learning and speaking English as a means to access broader educational

and professional opportunities, connect with diverse cultures, and participate more fully in the global community. Additionally, they expressed the importance of understanding their rights and learning how to defend themselves effectively in society.

> "When you educate a girl, it's like you educate a whole nation"
> —Out of school girl, 14 years old, Mani, Katsina
> "The only way I can become a female doctor is if I go to school, I can't learn that by staying at home"
> —Out of school girl, 14 years old, Mani, Katsina
> "Education allows girls to know their rights, how to defend themselves and differentiate between right and wrong"
> —University student, 19 years old, Sokoto-North, Sokoto

However, the influence of societal gender norms was evident in how the girls perceived the importance of education. The societal emphasis on domestic roles for women shaped their views, leading them to see schooling as essential for becoming better sisters, wives, and mothers. Some of the girls viewed education as a means to learn how to manage household responsibilities more effectively and gain respect within their familial structures. As one in-school girl from Shuni, Sokoto, expressed, "When you're educated, your husband and co-wives would respect you." This sentiment reflects the belief that education enhances their status and authority within their homes. Similarly, an out-of-school girl from Rimi, Katsina, highlighted the practical benefits of education, stating, "I would like to go to school so I can have answers for my child when he grows up." This perspective reflects the girls' aspirations to be knowledgeable and capable mothers who can support their children's development.

> "When you're educated, your husband and Co-wives would respect you"
> —In-school girl, 18 years, Shuni, Sokoto
> "I would like to go to school so I can have answers for my child when he grows up"
> —Out of school girl, 15 years, Rimi, Katsina

When asked about the importance of education, most of the girls stated the value of education and mentioned that they would willingly attend school if given the opportunity. However, there were still a few girls who had expressed the belief that girls' education is unimportant. Although this viewpoint was held by a small minority, it is crucial to delve into their underlying reasons. Most of these girls viewed school attendance as an obstacle to fulfilling their domestic responsibilities

and childcare duties. As a result, they showed little enthusiasm for education, prioritizing their traditional roles at home over academic pursuits. This viewpoint highlights several deep-rooted cultural and societal issues that need to be addressed to foster gender equality and promote the importance of education for all children, regardless of gender. The perception that education is secondary to domestic responsibilities and childcare duties illustrates the pervasive traditional gender roles within these communities. Girls are often expected to prioritize household chores and caregiving over their personal development and academic achievements. This expectation not only limits their immediate opportunities but also reinforces the notion that a girl's primary value lies in her ability to manage a household and care for family members. Additionally, some girls expressed a strong preference for attending Islamiyya, an Islamic education program focused solely on teachings from the Quran. These girls often viewed formal Western education as a distraction from their religious studies. They believed that the time and effort required for Western education could detract from their commitment to Islamiyya, where they felt they gained more meaningful and culturally relevant knowledge. This perspective highlights a significant cultural divide and the value placed on religious education over Western educational models in their communities.

> "I prefer to go to Islamiyya because there I will actually learn important things"
> —Out of school girl, 18 years, Shuni, Sokoto

Do You Think It's Important for Girls to Go to School?

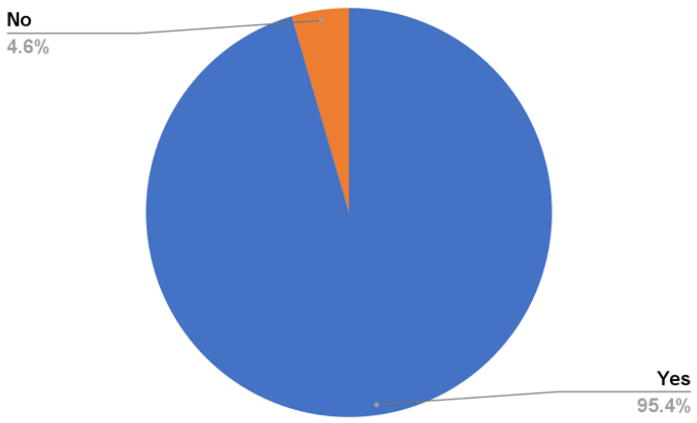

Fig 3: Pie chart visualizing the response to: Do you think it's important for girls to go to school? From 548 responses.

Why Do You Think School Is Not Important for Girls?

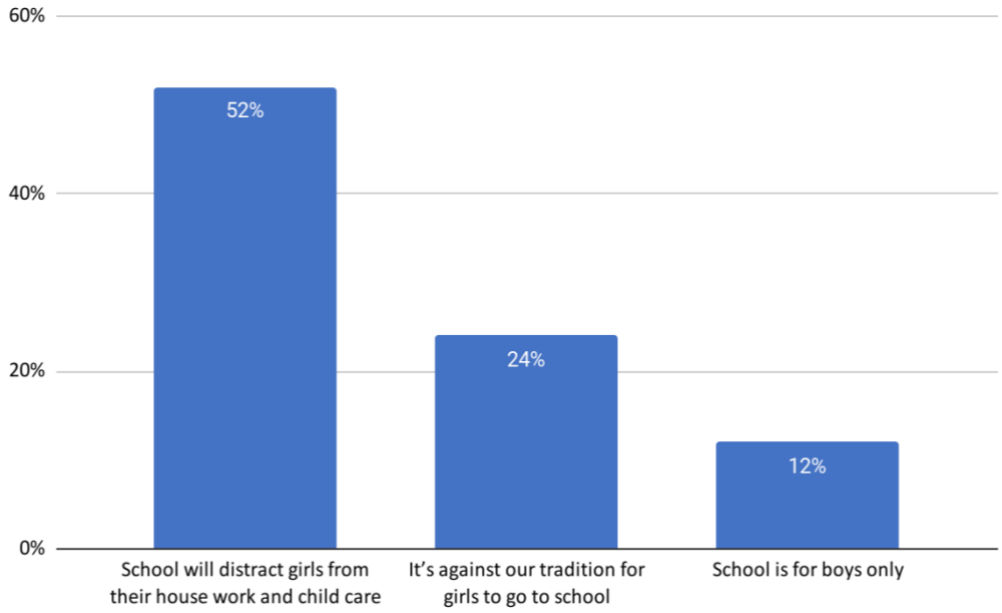

Fig 4: Why do you think school is not important for girls? From 25 responses.

Perspectives of Community Stakeholders in Northern Nigeria on Education for Adolescent Girls

Religious and cultural influences significantly shape how parents and caregivers raise their children, especially concerning gender roles and responsibilities. Islamic values and cultural norms lead to certain restrictions, such as limiting interactions between boys and girls. This limitation often prevents girls from attending schools. Furthermore, many community and religious leaders hold the belief that gender equality is fundamentally incompatible with Islamic teachings. They argue that the traditional roles and responsibilities prescribed by their interpretation of Islam do not align with contemporary notions of gender equality. As a result, these leaders often discourage the education of girls and their involvement in leadership roles within society. They emphasize that a woman's primary responsibilities should be centered around the home and family, advocating for a more limited public role for women and girls. Additionally, education is often viewed as a gateway to girls becoming wayward or straying from traditional and cultural values. This perspective is rooted in the belief that exposure to formal education, especially Western-style education, can lead to increased independence and empowerment, which some community

members fear may result in girls challenging established social norms and familial authority. This concern reflects a deep-seated apprehension about the potential for education to disrupt the traditional roles and expectations for girls within the community. This stance not only hinders the educational and professional opportunities available to girls but also reinforces existing gender norms that restrict their participation in broader societal activities.

> "I believe that modern education is a shortcut to girls being wayward"
> —Community leader, 42yrs, Rimi, Katsina
> "Girls should never go out without their parents' consent…but boys can go out whenever they want.. It's a religious practice"
> —Civil Servant, 50yrs, Rimi, Katsina

Fear and parental vigilance also plays a crucial role in guiding parents' decisions on raising their daughters. Many parents, during interviews, expressed concern about potential harm their daughters might face, leading them to isolate their daughters from society and restrict their freedom. Interviews revealed a pronounced gender disparity in this parental vigilance. Parents showed heightened concern and protectiveness toward their daughters compared to their sons. This disparity was particularly evident in monitoring their children's social circles, focusing on girls' friendships. This level of oversight often had detrimental effects on the girls, with some parents even admitting to preventing their daughters from attending school. Their rationale for such actions was firstly to shield their daughters from harmful influences and to restrict their interactions with boys. Additionally, there is a general perception among parents, community, and religious leaders that exposure to education could lead to immorality and bad behavior among girls by overly exposing them to western culture, serving as a barrier that prevents them from attending school. Also, just like some of the girls, parents view formal education as a diversion from their pursuit of Islamic and Quranic studies, leading them to delay formal schooling until after completing religious education.

> "Boys run errands while girls stay at home and do domestic chores. This is to protect the girls from danger"
> —Business woman, 43yrs, Rimi, Katsina
> "If not closely observed, children may be influenced by bad friends, especially girls"
> —Civil servant, 40yrs, Sokoto North, Sokoto
> "My male children attend Islamiyya and formal schools but my female children attend only Islamiyya, that's a rule in the family"
> —Farmer, 51yrs, Dange-Shuni, Sokoto

During the interviews, the male siblings shared that they often find themselves in a challenging position when confronted with gender inequality affecting their sisters. Many acknowledge the mistreatment and inequalities their sisters face and recognize these issues as significant. However, despite their awareness and concern, they often feel powerless to intervene effectively. This sense of powerlessness stems from deeply ingrained traditional norms and societal expectations that discourage men from challenging gender roles or advocating for women's rights. These norms create barriers for male siblings, making them hesitant to speak out or take action for fear of facing social stigma, ridicule, or backlash. As a result, despite their desire to support their sisters, they often find themselves constrained by societal pressures that undermine their ability to actively promote gender equality within their families and communities.

> "The problem with challenging some of these gender inequality is that parents will hate you for speaking up"
> —Male Sibling, 37yrs, Mani, Katsina

Government officials and members of civil society emphasized that the lack of education is a significant challenge for girls in these communities. Empowering them through education is vital to breaking harmful cycles and enabling informed decision-making. The interviews highlighted the further need for comprehensive education, including topics like menstrual hygiene and sexual health, to protect girls from exploitation and abuse. The officials acknowledged that changing deeply ingrained norms and mindsets would be challenging. However, they noted that various government ministries have initiated programs aimed at improving education for adolescent girls in Northern Nigeria.

Improving Education Access and Retention for Adolescent Girls in Northern Nigeria: The Girls' Perspective

In response to inquiries regarding methods to encourage girls to attend and remain in school during the interviews, a majority of the girls expressed that the initial step would involve raising awareness among their parents about the significance of education. Additional strategies mentioned by the girls during the interviews included financial sponsorship. Interestingly, some girls from Sokoto specifically pointed out that since many of the girls are compelled by their parents to engage in street vending instead of pursuing an education due to financial constraints, an effective approach to motivating girls would entail exploring

alternative means of supporting their parents without sacrificing their educational opportunities, thereby mitigating the need for street hawking.

Developing Intervention Strategies

Prioritizing "How Might We" Questions for Brainstorming

Following the synthesis of data collected during the formative, the various challenges that hinder girls' empowerment were identified. Each challenge was accompanied by corresponding "How Might We" (HMW) questions aimed at sparking innovative solutions. Afterwards, a collaborative prioritization exercise was undertaken by both the YUX and UNICEF teams to identify the most critical HMW questions. The criteria for prioritization included assessing the number of girls impacted, the potential level of impact, feasibility of solutions, and clarity of HMW questions for effective brainstorming sessions with girls during workshops. Through this process, 12 HMW questions emerged as top priorities. These questions were then presented to adolescent girls and stakeholders to guide collaborative efforts and solution development during subsequent workshops. The prioritized HMW questions were:

1. HMW address caregivers' concerns about bad company/ influence at school?
2. HMW promote the importance of educating daughters to parents who currently don't believe in it?
3. HMW establish support systems and counseling services for girls who have experienced sexual violence, ensuring they receive proper care and encouragement to reintegrate into education and society?
4. HMW find alternative ways to support girls' families financially that do not require them to sacrifice their education?
5. HMW address the stigma associated with adolescent pregnancy, providing support systems and resources to help girls continue their education and prevent exclusion and shame?
6. HMW provide comprehensive financial education programs specifically tailored for girls?
7. HMW educate girls on the importance and uses of legal identification?
8. HMW provide affordable or subsidized devices to ensure that girls can access technology despite financial constraints?
9. HMW design age-appropriate technology access programs that cater to younger girls, providing them with suitable devices and content while ensuring their safety in the digital world?

10. HMW introduce the concept of digital cash transfers and encourage it's use amongst girls?
11. HMW introduce girls to the concept of the internet, making them more comfortable and confident in using digital tools?
12. HMW encourage girls to seek medical care and allow them to be more comfortable doing so?

Enhancing Access to Education for Adolescent Girls: Interventions Suggested by the Girls

Below are some of the ideas generated by the girls during the workshop. As part of the activities, the girls sketched their ideas on cardboard, which are also presented here. Additionally, the second part of the workshop involved gathering feedback from other stakeholders on the girls' presented and sketched ideas. The feedback received from these stakeholders is included below.

1. Advocate the Importance of Educating Daughters Specifically to Fathers

The girls in both Sokoto and Katsina conveyed the idea that in promoting the significance of educating daughters, the focus should be on engaging fathers. They recommended involving additional intermediaries like uncles, sons, and respected neighbors to provide counsel to fathers. According to the girls, since fathers typically hold leadership roles within their households, they wield the authority to permit their daughters' school attendance. The girls also shared their perspective that fathers frequently undermine the value of educating their daughters. They also are aware of the fact that this intervention needs to be consistent as one conversation is not likely to change the fathers' mindset. The girls also proposed that the intermediaries present the advantages fathers can gain from educating their daughters, along with a few instances of successful stories involving educated girls.

In the storyboard, the girls share a story about a young girl named Fatima who dreams of attending school. However, her father does not permit her to pursue her education. To change her father's mind, Fatima seeks the help of her neighbors, her father's older siblings, and the first sons of the family. These stakeholders join together to convince Fatima's father by highlighting the benefits of education not only for Fatima but also for him. For example, one might say, "If Fatima goes to school, she can become successful and wealthy. Just look at Aisha, Baba Kamal's daughter, who now works in London!" Another might add, "If Fatima gets an education, she can become a nurse and take care of you and our mother. Healthcare

costs are rising, and as you grow older, her support will be invaluable." They emphasize that this approach requires ongoing advocacy and consistent intervention, rather than being a one-time effort.

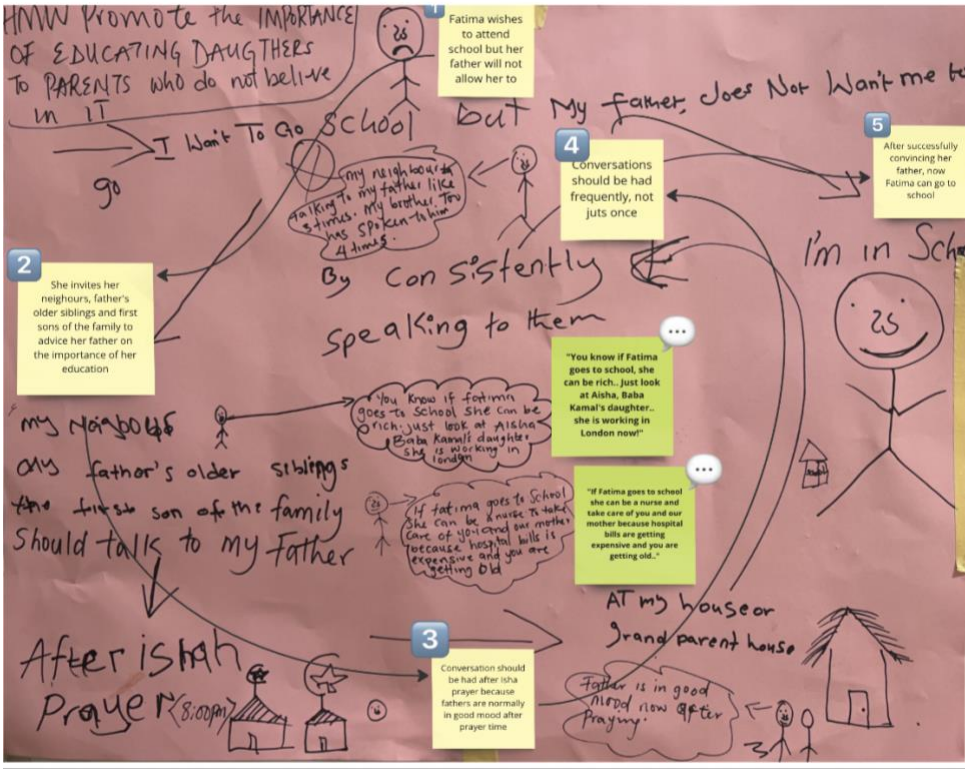

Fig 5: Storyboard Created by the Adolescent Girls During the Workshop

The male caregivers rate this intervention 4 out of 5, acknowledging that relatives greatly influence a man's decision, making this approach relevant. However, they expressed concerns that some fathers might resist advice on sending their daughters to school due to entrenched attitudes and lack of education. Despite this, they committed to encouraging fellow fathers to enroll their daughters in school and emphasized leading by example. They also suggested using religious leaders to reach fathers and educating them on the personal benefits of their daughters' education.

Religious leaders noted that some parents are set in their ways, making it hard to change their minds. They agreed with the girls that consistency is key. Leaders also identified a significant challenge: many parents are deeply entrenched in their beliefs. They suggested alternative approaches, such as consistent community engagements and incorporating the promotion of girls' education into mosque teachings.

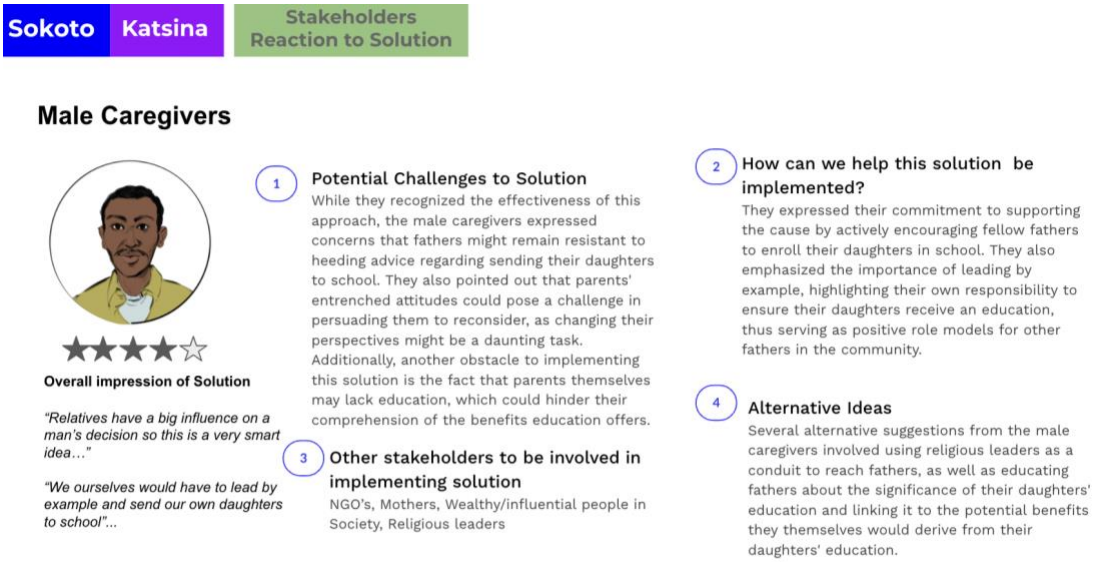

Fig 6: Perspective of Male Caregivers on the Idea provided by the Girls

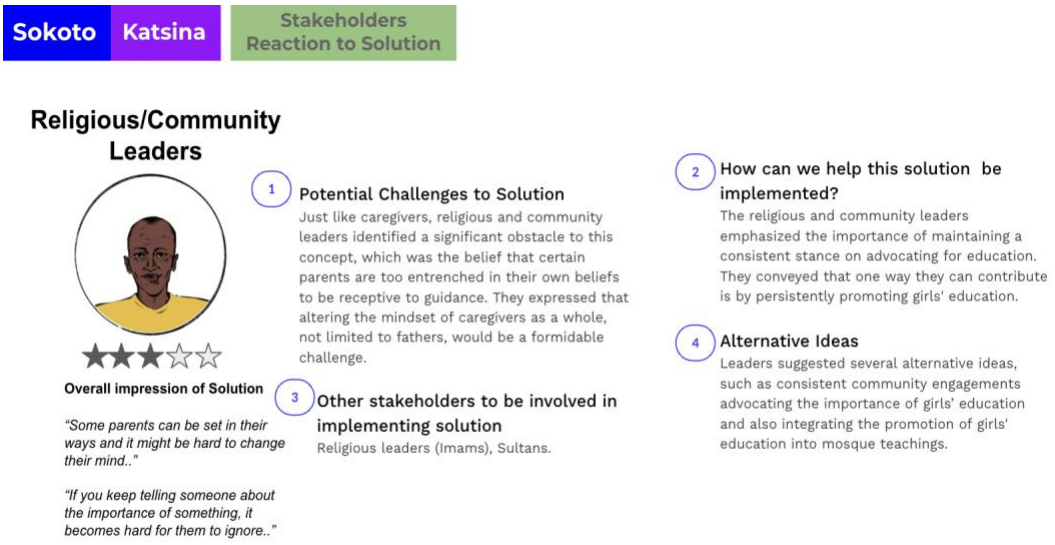

Fig 7: Perspective of Religious and Community Leaders on the Idea provided by the Girls

2. Ensure Access to Free Education

The girls conveyed that economic limitations frequently lead parents to refrain from enrolling their children in school. According to them, due to conflicting financial demands, parents sometimes underestimate the significance of providing education to their daughters. As a result, the girls suggest that offering cost-free education could serve as a remedy for this issue. This approach would eliminate the need for parents to make a difficult choice between their daughters' education and other endeavors.

In the storyboard, the girls continue the story of Fatima, whose father initially prevented her from attending school. Now, Fatima's father has been convinced to let her go to school, but he lacks the funds to pay for her tuition. Fortunately, Fatima receives a scholarship from NGOs and international development programs. These scholarships cover not only tuition but also school supplies and pocket money. The girls highlight that while tuition might be covered, other factors can affect a girl's sustained education and continued enrollment, such as the lack of pocket money for transportation and food or the absence of funds for school supplies like uniforms. Recognizing this additional barrier, the girls emphasize the need for interventions to cover all aspects of a girl's educational needs, ensuring they have everything necessary to stay in school and succeed.

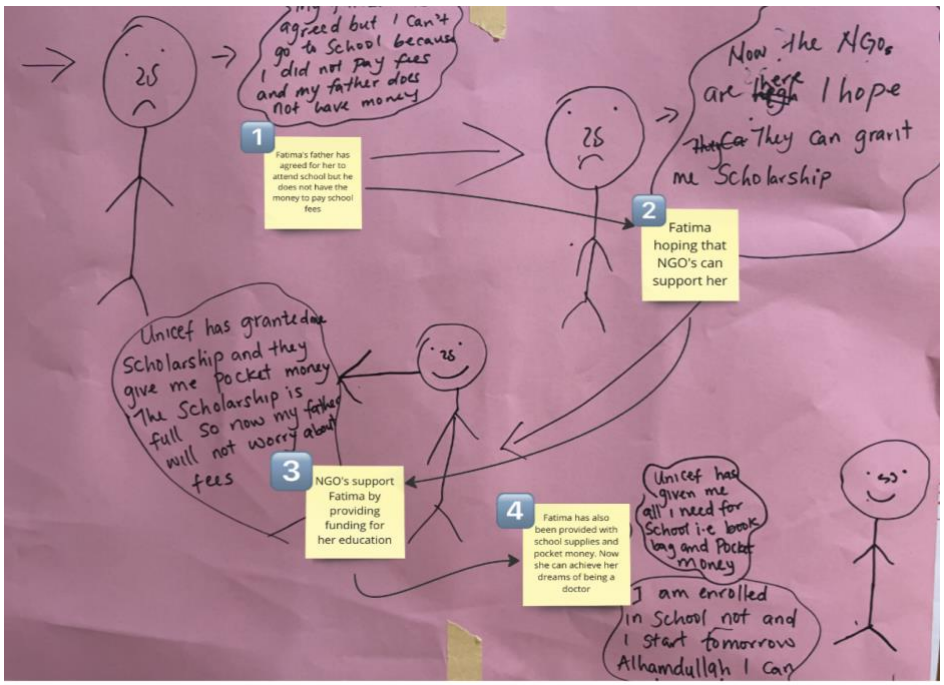

Fig 8: Storyboard Created By the Adolescent Girls During the Workshop

The female caregivers rated this intervention provided by the girls a 2 out of 5 and showed resistance to the concept, primarily due to doubts about its effectiveness. They voiced concerns that, despite offering tuition-free education, some parents might still refuse to send their daughters to school due to deeply entrenched beliefs. They emphasized that financial constraints were not the only obstacle to girls' education. Acknowledging that the undervaluation of girls' education is a significant barrier, the female caregivers expressed their willingness to help by consistently promoting its importance to their community members. They also suggested involving community leaders to help change the mindset of parents who currently do not see the value in educating their daughters.

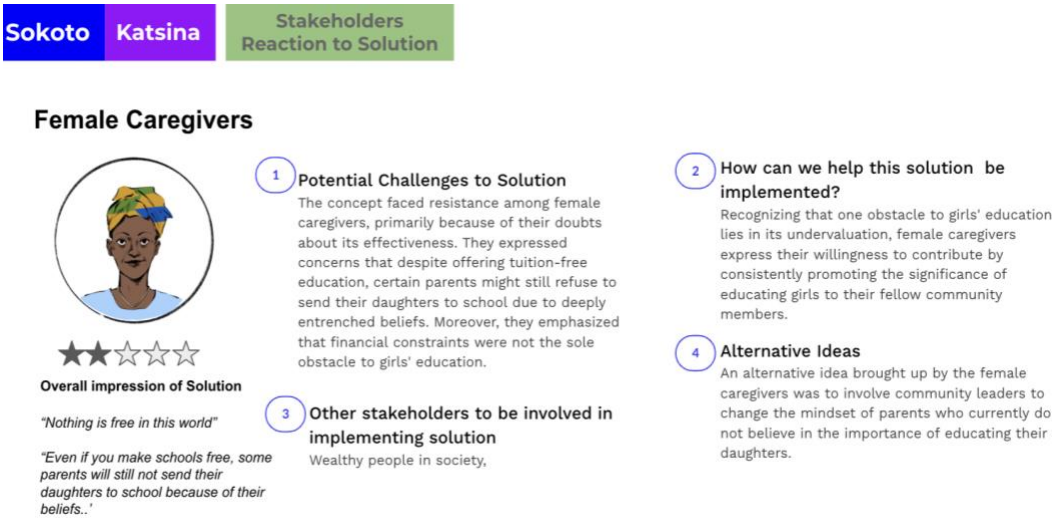

Fig 9: Perspective of Female Caregiver on the Idea provided by the Girls

Developing Social Behavioural Change (SBC) Strategy Drafts

After concluding the workshops, the YUX team formulated drafts of Social and Behavioral Change (SBC) strategies aimed at addressing and transforming harmful social norms. These strategies were crafted using evidence-based techniques to shift collective attitudes and behaviors within the community. The strategy drafts were informed by an in-depth analysis of the social and gender norms identified through formative research. The YUX team incorporated insights gained from the workshops, where girls proposed various interventions to overcome the obstacles they faced. This participatory approach ensured that the strategies were grounded in the real experiences and suggestions of the community members.

The Socio-ecological model was also applied to the development of these strategy drafts. This model recognizes that behavior is influenced by factors at different levels, including the individual, relationship, community, societal, and policy levels. SBC strategies take this multi-level perspective into account and target interventions at various levels to address the determinants of behavior change. Understanding the broader context in which behavior occurs is crucial for the design of SBC strategies, as it ensures that interventions are comprehensive and address all relevant factors that influence behavior. Hence, each intervention has been tagged based on it's applicable influence levels i.e interpersonal, community, services and institutions influence levels and so on.

By integrating these principles of SBC, the YUX team aimed to create a comprehensive and sustainable framework for shifting harmful norms and promoting the education and empowerment of girls. The strategies were designed to be adaptable and responsive to the evolving needs and feedback of the community, ensuring their long-term effectiveness and impact. The draft outlines the specific social norms and their underlying drivers, detailing how these norms negatively affect access to education and the continued schooling of girls. It then defines the specific SBC objectives designed to counteract these harmful norms. Lastly, it includes impact indicators to measure the success of these interventions, such as increased school enrollment rates, improved attendance, and changes in community attitudes towards girls' education.

For education, the main norms affecting the education of girls in Northern Nigeria identified include:

1. The devaluation of girls' education
2. The burdening of adolescent girls with the exclusive duty of managing household chores and caregiving
3. The practice of child marriage and prevalence of adolescent pregnancy
4. The belief that girls require extra protection and shielding from potential harm and social influence
5. The shaming and isolation of girls who have experienced sexual violence

The girls proposed several intervention strategies to enhance their enrollment and retention in education, including:

1. **Targeted advocacy [Interpersonal]**
   Advocate the importance of educating girls to fathers and other male caregivers. Involve intermediaries like uncles, sons, and respected neighbors to provide counsel to and enlighten fathers and other male caregivers about the importance of educating their daughters.

2. **Involve community and religious leaders in the advocacy [Community]**
   Involve respected stakeholders such as religious and community leaders in promoting gender equality in education through relevant religious teachings.
3. **Empower mothers and other caregivers [Interpersonal]**
   Provide skills training to mothers so as to empower them to be financially stable and independent. According to the girls, empowering their mothers would reduce the financial constraints within the family hence giving them the opportunity to attend school.
4. **Make education more accessible to girls [Services and institution]**
   Due to conflicting financial demands, parents sometimes underestimate the significance of providing education to their daughters. As a result, the girls suggest that offering cost-free education could serve as a remedy for this issue.
5. **Identify and highlight role models [Community]**
   Leverage the media (i.e radio announcements, TV ads, etc,) to showcase success stories of educated girls who have positively impacted their communities.
6. **Assign trusted female teachers who will serve as "school parents" [Services and institution]**
   These designated educators should maintain a certain level of connection with the girls' parents and provide periodic updates on the girls' conduct within the school environment. This would thereby lead parents to be comfortable enough to send their daughters to school.

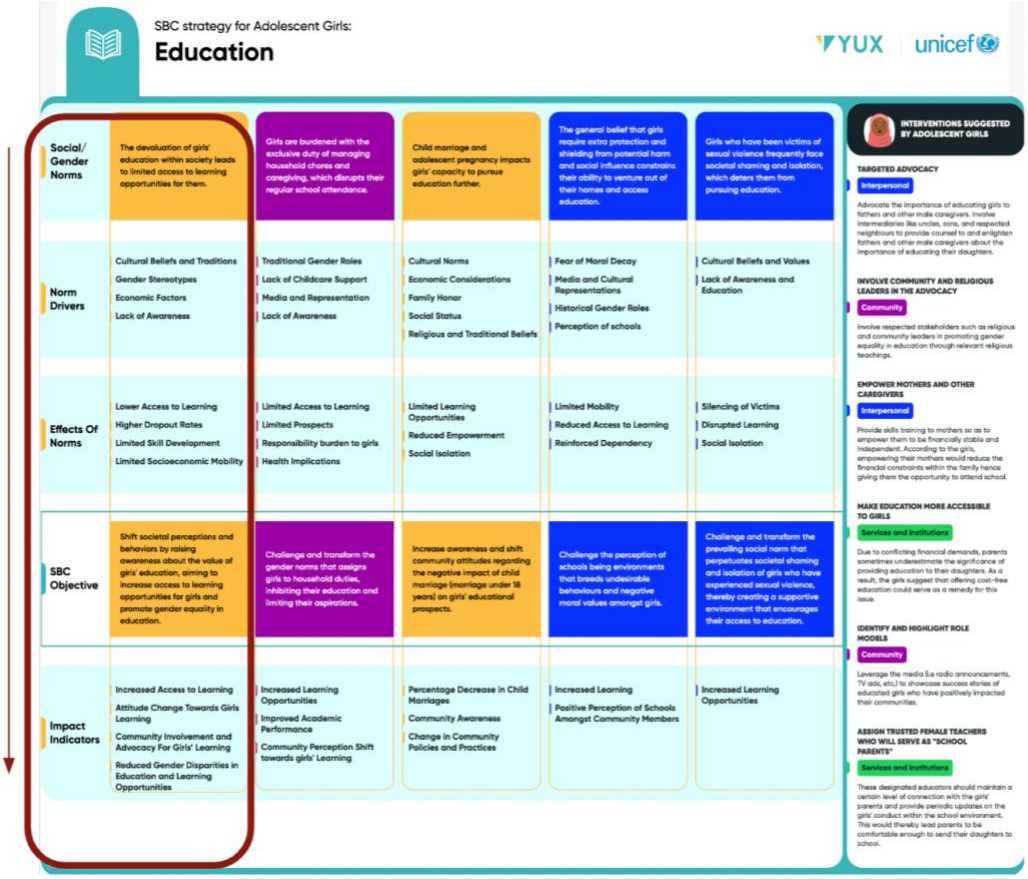

Fig 10: Social and Behaviour Change Strategy Draft Aimed at Enhancing Access to Education for Adolescent Girls

General Learnings from a Participatory Approach

Experience and Learnings from the Field

It is important to share some of the challenges we faced as well as our general learnings from the field research. Recognizing the benefits of an iterative approach, we chose to initiate the field research in Katsina before progressing to Sokoto. The challenges and corresponding learnings from the Katsina field research informed our subsequent work in Sokoto. Some of these learnings include:

Planning Logistics Ahead of Time

While a participatory approach adds significant value to the research process, it does need a lot of planning. Coordinating the logistics on transportation and accommodation for the 20 young women who were involved in the research team was quite challenging. We were committed to ensuring the safety of these young

women in the field and their secure return to their respective homes. In addition to the already nuanced situation, logistics planning was not done early due to a misunderstanding on who would be handling what. Hence, we had to figure out a lot of things just a few days before the field research was scheduled to start. Despite the stressful period for both myself and the team, we successfully managed the logistics (many thanks to the YUX local researchers: Rukayat Abdulwaheed, Balkisu Fasani, Yusuf Jamiu, and Lukman Muhammad, who were able to navigate this situation effectively and act fast!). Moving ahead to the Sokoto research, careful planning was carried out well in advance to avoid any further stress.

Involving Local Female Guides

An important lesson we learnt from the Katsina field research was the significance of including local female guides, who were well-acquainted with the community and its culture, to accompany the team in the field. The absence of these guides in Katsina posed challenges, particularly during household interviews with adolescent girls and other stakeholders. Despite the inclusion of local young women in the research team and female YUX researchers fluent in the local language, some families were hesitant to allow the research team into their homes. As a result, for the Sokoto field research, we made sure to engage local female guides, such as the daughters of village heads or community leaders, who, possessing a deep understanding of local dynamics, played a crucial role in establishing trust and rapport within households.

More Elaborate Consent Training for the Young Women on the Research Team

The young women on the research team demonstrated exceptional skills in conducting interviews and fostering comfortable environments for the adolescent girls. However, we observed a less effective consent process. Having established rapport with many of the girls being interviewed due to their familiarity with the community, the young women researchers sometimes bypassed the formal consent process. They often started conversations with the research participants before reviewing the consent material provided. To address this in the Sokoto research, we emphasized the critical importance of obtaining consent during our training, despite the researchers' familiarity with participants. We provided practical guidance on how to effectively collect consent, ensuring that all ethical protocols were followed. Given the sensitive nature of the topics discussed, we also ensured that the young women

fully grasped the interviewees' right to withdraw consent at any point, even after the interview had started.

Choosing When to Employ a Participatory Approach

Employing a participatory approach for this project was immensely valuable. It enabled us to gather insights that would have otherwise been inaccessible through conventional methods. The young women involved in the research played a crucial role in ensuring that their perspectives and experiences were authentically represented in the findings. Their active participation not only enriched the depth of our understanding but also empowered them to advocate for meaningful change within their communities.

While this approach was valuable for this specific project context, we recognize that participatory research may not always be suitable for certain types of projects. For instance, technology assessment projects often involve specialized technical knowledge that may not be readily available within the general community participating in the research. Assessing complex technologies requires expertise that typically lies within specialized fields, making participatory methods less practical for evaluating technical specifications and functionalities.

Similarly, topics such as reproductive health or gender-based violence necessitate handling delicate conversations with sensitivity and expertise. These subjects often involve personal experiences and sensitive information that require trained professionals to ensure ethical considerations and confidentiality. Participatory research, while valuable for community engagement, may not always provide the necessary expertise and structured support required for these sensitive topics.

Furthermore, in large-scale data collection efforts aimed at national or regional levels, standardized methodologies are often preferred to ensure consistency and comparability of data across different regions or populations. Participatory approaches, which emphasize local context and community involvement, may not always align with the rigorous standards and uniformity needed for national-level data aggregation and analysis.

When engaging with important decision-makers, particularly in formal settings, the focus typically shifts towards structured presentations and formal discussions rather than participatory methods. Decision-makers often require clear, concise information presented in a manner that supports informed decision-making and policy formulation. Formal presentations and structured discussions provide a platform for detailed analysis and strategic planning, which may be more effective in influencing policy and decision-making processes compared to participatory

approaches that emphasize community engagement and collaborative decision-making.

Considerations When Employing a Participatory Approach

Provide Elaborate Training to the Stakeholders

For this research, it was crucial to provide the young women researchers with comprehensive training to fully equip them with the necessary knowledge for fieldwork and workshop facilitation. During participatory research, stakeholders may lack expertise in research methods or subject matter, which can impact the quality of data collection, analysis, and interpretation. Therefore, it is essential to ensure adequate, detailed, and well-contextualized training before starting research. Additionally, we conducted pilot research rounds before the main fieldwork to ensure the young women researchers were thoroughly prepared and equipped with the necessary skills and knowledge.

Avoid the Perception of "Tokenism"

There is a risk that participatory research processes may fall into tokenism, where the involvement of community members or stakeholders appears inclusive but fails to genuinely value their contributions or address power imbalances. Tokenism occurs when community members are superficially included to meet participation requirements without meaningfully integrating their input into decision-making. This can lead to frustration and disillusionment among participants, undermining the research goals and perpetuating existing inequalities.

To avoid tokenism, it is crucial to engage stakeholders as equal collaborators throughout all phases of the research process, from project design to dissemination. This involves actively seeking their input, respecting their perspectives, and incorporating their suggestions into the research framework. For instance, in this project, girls and young women were instrumental in curating the research guide, formulating survey questions, and designing workshop activities. Their insights and experiences were essential in shaping the research approach and ensuring its relevance and benefit to their communities. Additionally, during the workshops, the girls determined UNICEF's priorities by voting on the "How Might We" questions they deemed most critical based on their lived realities. The final decisions reflected the issues the girls identified as significant, ensuring that their voices were not only heard but also influential in guiding the project's direction.

Moving Forward: The Impact

Moving forward, the UNICEF Nigeria team is actively considering the findings from the field research and workshops, preparing to implement the interventions developed with and for the girls. A key theme that emerged from the girls' suggestions was the critical importance of community engagement and securing the support of influential stakeholders, including community and religious leaders. For instance, one of the workshop's "How Might We" questions focused on encouraging fathers to prioritize their daughters' education. The primary suggestion from the girls was to sensitize community and religious leaders who can influence parental attitudes. Consequently, efforts are now concentrated on engaging these key stakeholders to ensure the successful implementation of the strategies.

The participatory approach in NGO programming is transformative as it positions community members as active contributors rather than mere beneficiaries. By involving them directly in decision-making, NGOs ensure that their programs are tailored to the specific needs and aspirations of the people they serve. This approach fosters a sense of ownership, empowering communities with practical skills and positioning them as partners in designing solutions that impact their lives. Furthermore, participatory methods enhance transparency and accountability, building trust and mutual respect between NGOs and communities. By creating a comfortable environment that mitigates traditional power imbalances, participatory approaches encourage openness and elicit deeper insights, leading to more relevant and effective interventions. This inclusive strategy not only improves program outcomes but also strengthens the social fabric, promoting collaboration and resilience among all stakeholders. For instance, the girls involved in the research team were not just subjects of study but partners in the process. This involvement provided them with valuable skills and a sense of agency, fostering a deeper understanding of the issues they face and empowering them to contribute to solutions. Their participation also ensured that the interventions developed were more reflective of their realities and aspirations, enhancing the overall impact and sustainability of the programs.

Fig 11: Adolescent Girls in School in Katsina State

About the Author

Oluchi Audu is a design researcher at YUX design, a Pan African research and design company with a mission to create digital products and services localized to the African context. She holds a bachelor's degree in Psychology and a masters' degree in Human Computer Interaction from Carleton University. Her interests lie in inclusive design and finding ways to make technology usable and accessible to as many people as possible. She also focuses on projects around gender norms and social behaviors that impact marginalized communities, particularly women and girls. Through her work, Oluchi aims to develop innovative strategies that promote gender equality, empower underrepresented groups, and ensure that sustainable solutions address diverse needs.

Notes

I would like to express my deepest gratitude to my exceptional EPIC curator, Anni Ojajärvi, for her invaluable feedback in shaping this case study. I would also like to acknowledge the UNICEF Nigeria country office for initiating this critical project. I am very grateful to my outstanding managers, Camille Kramer Courbariaux and Yann Le Beux, for their unwavering support, and to my incredible team members—Yusuf Jamiu, Rukayat Abdulwaheed, Lukman Salihu, Azeezah Adekunle, Rajay Shah,

Alexandra Aho, Adama Mbengue, Yash Chudasama, and Cumi Oyemike—whose significant contributions were instrumental in bringing this work to fruition.

References

Bolarinwa, O.A., Ahinkorah, B.O., Okyere, J. et al. (May, 2022). A Multilevel Analysis of Prevalence and Factors Associated with Female Child Marriage in Nigeria Using the 2018 Nigeria Demographic and Health Survey Data. *BMC Women's Health 22*(158). https://doi.org/10.1186/s12905-022-01733-x

British Council (2012). Gender in Nigeria Report 2012: Improving the Lives of Girls and Women in Nigeria, Issues, Policies, Action. Second Edition. https://www.britishcouncil.org/sites/default/files/british-council-gender-nigeria2012.pdf

British Council (2014). Girls' Education in Nigeria Report 2014: Issues, Influencers and Actions. https://www.britishcouncil.org/sites/default/files/british-council-girls-education-nigeria-report.pdf

Chukwu, O. (April, 2021). New Dangers to Girls' Education in Northern Nigeria. https://inee.org/blog/new-dangers-girls-education-northern-nigeria

Eweniyi, G.B. & Usman, I.G. (July 2013). Perception of Parents on the Socio-Cultural, Religious and Economic Factors Affecting Girl-Child Education in the Northern Parts of Nigeria. *African Research Review 7*(3). https://doi.org/10.4314/afrrev.v7i3.5

Global Coalition to Protect Education from Attack (October 10, 2018). The Impact of Attack on Education for Nigerian Women and Girls. https://reliefweb.int/report/nigeria/i-will-never-go-back-school-impact-attacks-education-nigerian-women-and-girls-summary

James, G. (2010). Socio-Cultural Context of Adolescents' Motivation for Marriage and Childbearing in North-Western Nigeria. *Current Research Journal of Social Sciences* 2(5): 269-275.

Mawusi Amos, P. (December 18, 2013). Parenting and Culture – Evidence from some African Communities. In Parenting in South American and African Contexts, M.L. Seild-de-Moura ed. IntechOpen. https://doi.org/10.5772/56967

National Bureau of Statistics and United Nations Children's Fund (August, 2022). Multiple Indicator Cluster Survey 2021, Statistical Snapshot Report. Abuja, Nigeria: National Bureau of Statistics and United Nations Children's Fund. https://mics-surveys-prod.s3.amazonaws.com/MICS6/West%20and%20Central%20Africa/Nigeria/2021/Snapshots/Nigeria%20MICS%202021%20Statistical%20Snapshots_English.pdf

National Population Commission (2019). Nigeria Demographic and Health Survey 2018. https://dhsprogram.com/pubs/pdf/FR359/FR359.pdf

Nielsen, H. (2021, September 27). *Literacy in Nigeria (SDG Target 4.6)*. Education Articles, AWCC Scotland. education2@fawco.org

Odimegwu, C., Somefun, O.D., & Akinyemi, J. (2017). Gender Differences in the Effect of Family Structure on Educational Outcomes Among Nigerian Youth. *Sage Open*, *7*(4). https://doi.org/10.1177/2158244017739948

UNICEF Nigeria (2019). "Education." https://www.unicef.org/nigeria/education

Yewande, T., & Olawunmi, A. (2023). Exploring the Impact of Cultural Beliefs and Practices on Women's Education in Northern Nigeria. *Journal of Education Review Provision*, *3*(1), 1-10. https://doi.org/10.55885/jerp.v3i1.191

Yusuf, H. (January, 2014). Purdah: A Religious Practice or an Instrument of Seclusion. *American International Journal of Contemporary Research* 4:1, 239–245. https://www.aijcrnet.com/journals/Vol_4_No_1_January_2014/23.pdf

Unveiling the Human Experience of Medicaid Renewals: Using a Data-Driven and Human Centered Design Approach to Improve Public Service Delivery

ALYSSA KROPP, *United States Digital Service*
EMILY MANN, *United States Digital Service*
HEATHER MYERS, *formerly United States Digital Service*
IZZIE ZAHORIAN, *United States Digital Service*

This case study demonstrates how a transdisciplinary, human-centered approach to research and process improvement can increase efficiency, enrollment, and customer experience across a complex ecosystem of organizations, people, processes, and technologies. In the wake of an executive order to improve customer experience and service delivery, a team at the United States Digital Service (USDS) worked with Centers for Medicaid and the Children's Health Insurance Program Services and state program teams to improve Medicaid eligibility and enrollment. USDS used a mixed-methods research approach that centered lived experience, and mobilized small strike teams to systematically review software and processes. The ROI was substantial: USDS estimates that this work facilitated over 5 million automatic renewals and saved state employees and the public approximately 2 million hours of burden by the end of 2024.

Introduction

This case study explores the work of the United States Digital Service (USDS) in partnership with Centers for Medicaid and CHIP Services (CMCS) and state Medicaid offices to streamline and improve the Medicaid redeterminations process for enrollees, application assistors, and state staff. This work is part of an interagency program supporting President Biden's Executive Order on Transforming Federal Customer Experience and Service Delivery to Rebuild Trust in Government. The Presidential Management Council formed a cross-government effort and interagency team to tackle the designated life experience of facing a financial shock and becoming newly eligible for critical supports.

In response to the COVID-19 Public Health Emergency (PHE) in 2020, Congress enacted several policies to protect Americans. Among these was a continuous enrollment provision that allowed states to maintain Medicaid and CHIP coverage for enrollees throughout the duration of the PHE—a policy that protected

Kropp, Alyssa, Emily Mann, Heather Myers, and Izzie Zahorian. "Unveiling the Human Experience of Medicaid Renewals: Using a Data-Driven and Human Centered Design Approach to Improve Public Service Delivery." *EPIC Proceedings* (2024): 465–484. https://www.epicpeople.org/unveiling-human-experience-of-medicaid-renewals. ISSN 1559-8918.

health care coverage for, at its peak, more than 90 million eligible low-income adults, children, pregnant women, elderly adults and people with disabilities (Medicaid, 2023). This provision expired on March 31, 2023, and states were given 12–14 months to resume annual redeterminations: the process of verifying enrollees' eligibility to either renew or terminate their coverage. The period following the end of the continuous coverage provision in 2023 and return to normal operations is often referred to as the "unwinding" period.

Medicaid agencies around the country faced a massive challenge: to conduct redeterminations for a record number of enrollees with fewer staff, rebooting systems that sat dormant during the pandemic, and finding up-to-date enrollee contact information for the first time in three years. Many enrollees had never experienced the redetermination process before, having received Medicaid for the first time during the pandemic, and thus did not know what was required of them.

Estimates from 2022 projected that over 7 million people were at risk of losing their health care coverage—not because they were no longer eligible—but because of technological and administrative barriers including undelivered mail, burdensome renewal packets, and/or system errors (Office of the Assistant Secretary for Planning and Evaluation, 2022).

This case study describes how some of the challenges faced by states during this time were met with a transdisciplinary human-centered design (HCD) approach. USDS, a component of the Office of Management and Budget, specializes in technology and design. USDS, in partnership with CMCS, joined the all-of-government response to unwinding to bridge the gap between policy and the challenges of implementation through a series of intensive sprints with 8 states from March 2023 to April 2024. Applying a mixed-methods research approach and leveraging the power of a cross-functional technology team, USDS's efforts resulted in the reduction of millions of hours of burden time and improved customer experience of Medicaid renewals for enrollees, application assistors, and eligibility workers.

Project Context

The Impact of COVID on Medicaid

Program History and Foundational Policies

Medicaid is a public program that provides low-income Americans, children, pregnant women, elderly adults, and people with disabilities with health care coverage**.** Administered by states and funded jointly with the federal government,

Medicaid was established in 1965 within Title XIX of the Social Security Act. It was created to provide a healthcare coverage option for those who could not access private health insurance or other medical support (Medicaid.gov, n.d.).

The continuous enrollment provision that was enacted as part of the Families First Coronavirus Response Act (FFCRA) allowed eligible Medicaid enrollees to receive health coverage for the duration of the pandemic—for some, that meant 3 years of continuous coverage. Under most circumstances, Medicaid enrollees are only entitled to coverage for 12 months, after which, states must conduct a redetermination of eligibility to either renew or terminate the enrollee's coverage.

There are two possible renewal pathways an enrollee may experience: (a) "automatic" renewal wherein a state uses electronic data sources to verify eligibility, or (b) "manual" renewal wherein the enrollee provides eligibility information by completing a form and including any necessary documentation (such as proof of income). The two experiences vary dramatically for both the enrollee and state eligibility workers who process annual renewals.

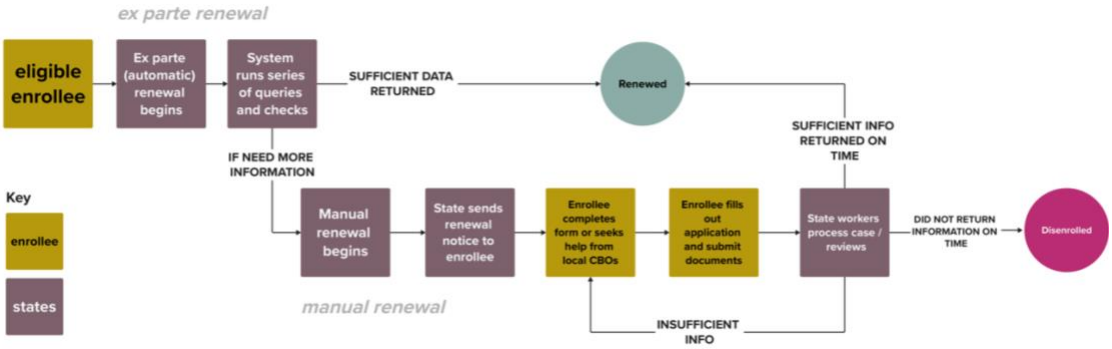

Figure 1: Simplified Medicaid Renewal Journey for an eligible individual. State implementations of Medicaid Renewal pathways vary. This is for illustrative purposes only.

In the first renewal pathway, if a state can verify eligibility using existing data, no action is required of the enrollee. Referred to as "ex parte," this automatic renewal process saves time for both eligibility workers and enrollees. According to 42 CFR 435.916(b), state must attempt ex parte renewal before requesting additional information from the enrollee.

If states are unable to verify eligibility on an ex parte basis, enrollees are required to manually complete and return a renewal form verifying that no information has changed from the time of application or updating information if it has. The manual renewal form can commonly be completed online, on-paper, or over-the-phone, and states must pre-populate the form with existing information (Electronic Code of

Federal Regulations, n.d.). Across most states, the end-to-end renewal form can be over a dozen pages long.

Framing the Problem: Policy, Meet Implementation

Both regulations mentioned above are intended to reduce administrative burden on enrollees, and to save eligibility workers time in processing renewals. With these time-saving policies on the books, one might wonder why reinstating annual renewals posed such a challenge for states, and why millions of Americans were at risk of losing their health care coverage during the unwinding period.

The first challenge states faced was the sheer volume of renewals that needed to be completed during the unwinding period. During COVID, because of the continuous enrollment condition, enrollment in Medicaid and CHIP skyrocketed, with roughly 94M people enrolled in the program by the end of the pandemic; this was a gain of close to 23M eligible people in the program in just 3 years (Medicaid.gov, n.d.). At the same time, state Medicaid agencies were facing staff shortages, many hurrying to hire and train new eligibility workers to complete redeterminations for the first time. CMCS described this period as, "the single largest health coverage transition event since the first open enrollment period of the Affordable Care Act." (Medicaid.gov, n.d.)

Given this historic volume of enrollees, states looked to better leverage automatic renewals in order to improve compliance, process redeterminations on time, and reduce the load on strained eligibility worker. However, many states struggled with low ex parte rates—at the start of redeterminations, states processed only 1 in 4 of their redeterminations on an ex parte basis (Medicaid.gov, n.d.).

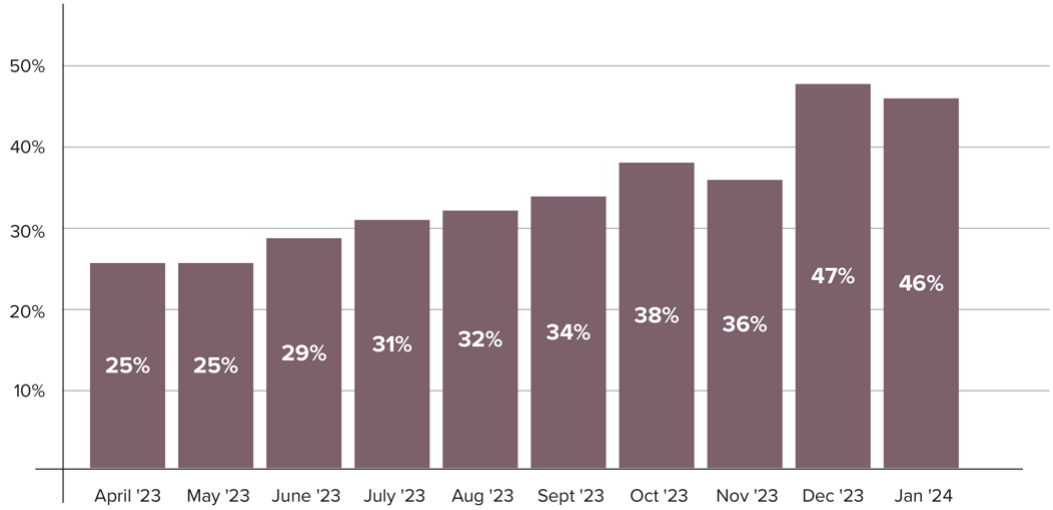

Figure 2. Average national ex parte rate broken down by month. Data from https://www.medicaid.gov/resources-for-states/coronavirus-disease-2019-covid-19/unwinding-and-returning-regular-operations-after-covid-19/data-reporting/monthly-data-reports/index.html.

While this may seem like a simple problem to solve, the reality is that federal policy makers must navigate a complex landscape of Medicaid and CHIP programs across 56 states and territories that each have their own operating constraints. Each Medicaid program has a bespoke eligibility system, different data integrations, distinct eligibility criteria, and different demographic compositions. Some states, but not all, have expanded eligibility under the Affordable Care Act; additionally, some Medicaid programs are integrated with other health and human services programs like Supplemental Nutrition Assistance Program (SNAP) or Temporary Assistance for Needy Families (TANF). This diverse ecosystem not only makes a "simple" task of setting performance benchmarks a challenge, but also complicates the process of creating implementation guidance, such as use-cases, that capture all of the variance in these programs.

For the remaining cases that could not be processed automatically, states were preparing to mail out and process historic numbers of manual renewals. Medicaid offices faced an entirely different set of challenges, such as undeliverable or delayed mail to recipients; the public's general lack of awareness and understanding of what the unwinding period was; and challenges in the experience of completing renewal forms due to poor UX on websites and the lack of plain language in the guidance of how to complete them.

Hypothesis

From early conversations with state Medicaid agencies, CMCS leadership, and USDS's desk research, the team hypothesized that a) helping state agencies optimize automatic renewals, and b) decreasing barriers throughout the manual renewal process would provide time and cost savings for the state while providing a streamlined and more equitable experience for the enrollee. The team proposed going on the ground with states with a small cross-functional team, could make a considerable impact to states during this historic transition.

Framing the Project Approach

CMCS and state Medicaid offices are not alone in this struggle between writing the policy and implementing it to produce its intended outcome. In 2008, the National Academy of Public Administration surveyed career policymakers in federal government and found only "16% consider government proficient at designing policies that can actually be implemented" (Pahlka, 2020).

Examining Structural Foundations of Government Service Delivery

Jennifer Pahlka, in her book *Recoding America*, explores "the temporal, organizational, structural, and cultural gaps between policy and tech teams, and between tech teams and the users of that tech," through dozens of examples across federal and state government (Pahlka, 2023:16). By examining these gaps and challenging the behaviors and beliefs that created them, Pahlka urges technologists and policymakers alike to join together. The structures that created these gaps, however, are strong. Pahlka notes:

> Chief among [government's dysfunctions] is a structure and way of thinking deeply rooted in American culture: hierarchy. …What's valued in government isn't the nuts and bolts of implementation but the rarefied work of policymaking. Digital work…in government is reduced to an afterthought. It's not what important people do, and important people don't do it. … At times it almost seems that status in government is dependent on how distant one can be from the implementation of policy (Pahlka, 2023:15).

The USDS team knew that to be successful it would require a cross-functional team that met states where they are. This approach allowed the team to bring implementation expertise in design, engineering, data science, and product directly to states, with the added benefit of having direct connections with the agency that knew

each policy's intention. The team's experience working at, with, or adjacent to vendors also meant that the team could understand the language and culture of state vendor teams and support the translation between policy ideals and goals to the nuts and bolts of implementation.

These approaches are unusual in government. While the civic technology field is growing, many states and government agencies do not have access to this type of expertise. By engaging with states directly and helping define the problem and work together to operationalize ways to solve it, the team aimed to encourage a reimagining and reworking of how technology could be implemented at the state level.

An Introduction to USDS Values and Process

USDS holds six values that drive the approach to delivering better government services. These include (1) Find the truth, tell the truth, (2) design with users, not for them, (3) go where the work is, (4) optimize for results, not optics, (5) create momentum, and (6) hire and empower great people (usds.gov, n.d.). While all tenets influenced the team's approach to the work, the first three helped underscore how the team made decisions around methodologies and implementation. Namely, this informed the decision for the team to go on the ground to meet with states directly, to engage community-based organizations and staff working on the front lines early and often, and to bring up hard discussions for the betterment of the service for the public beneficiaries and state staff overall.

Similarly, USDS strives to engage multi-disciplinary teams to ensure that projects have cross-functional specialties that can meet any challenge. For this project, the USDS team consisted of 4 human centered designers, 4 engineers, and 2 product managers. The cross-cutting nature of benefit services meant that data science and engineering were a vital skill needed to understand the technical systems used in implementing the policies and eligibility determinations, while designers led research of surrounding support services, public facing workshops, and redesigns of visual touchpoints and operational practices. The cohesion of the team meant that each practitioner brought their own lens to challenge, champion, and support the various workstreams to strengthen the final results.

Focusing on Addressing Administrative Burden

As noted by Herd and Moynihan in their book *Administrative Burden*, changing foundational ways of working can be deeply challenging for many reasons. While policy makers work to develop programs, structures, and rules to support American

livelihoods, unintended consequences can come about during implementation, such as costs, access difficulties, and user experience challenges (Herd and Moynihan, 2018). This friction in completing tasks that allow someone to receive results is known administrative burden. The stress, frustration, and lack of ability to complete tasks for these programs can end up impeding access for the very people the programs are trying to serve. Don Moynihan sums this up well in the following example:

> We go through life with a set of expectations about how our interactions with administrative services work. And over time, if we are dependent upon certain services, if there's a change in those, it can have these fairly dramatic effects. And so there are other examples in, for example, Arkansas when they introduce work requirements for Medicaid. A paper in the New England Journal of Medicine founded about 95% of those who lost their Medicaid coverage due to the work requirements were actually working or should have been exempted because of a disability, but where they really struggled was with the paperwork requirements. It was really the reporting part that they found difficult. It wasn't that they weren't actually completing the eligibility guidelines. (Grossman et. al., 2023).

For Medicaid, the repercussions of administrative burden can be the difference between having lifesaving health coverage or not. As Wilke et. al. point out, much has been done to reduce burdens within the Medicaid program, like removing the need for in person interviews, but other burdens still exist, such as requiring signatures on applications (Wilke et. al., 2022).

For CMCS and the USDS team, reducing administrative burden meant focusing on ways the team could decrease the touchpoints the public had in their renewal experience. And when they could not, the focus would be improving the experience of these touchpoints when it was essential to have that interaction.

Transitioning from Executive Order to Implementation

The Executive Order on Transforming Federal Customer Experience and Service Delivery to Rebuild Trust in Government declares that government should return to focusing on a government of, by, and for the people, and that government programs and services should be held accountable for designing and delivering services with a focus on the actual experience of the people whom it is meant to serve. It particularly calls out the use of HCD methodologies the federal government's management of its customer experience and service delivery should be

driven fundamentally by the voice of the customer through human-centered design methodologies.

While the Executive Order was issued in 2021, these methods are still fairly new in government settings. When starting this project, USDS was aware that CMCS and states had probably heard some design terminology, but had most likely not worked in the applied design and technology practice that the team are experts in. This meant that many approaches that were familiar to USDS, such as interviewing beneficiaries or their supporters, prototyping, and more could seem unfamiliar and polarizing to start. That, in tandem with our focus on moving quickly, meant that USDS had to ensure stakeholders were aware and engaged in each step of the process.

Doing so meant encouraging states to rethink the foundations of how program implementation had been carried out in many federal benefit programs and directly challenging "the durability of social conventions and cultural traditions and thus resistance to organizational transformation" (The White House, 2021). In the past, this has looked like implementing programs without feedback from the public or siloing policy and technology teams away from each other. USDS knew that building trust with the state teams would be vital to be successful. As such, the USDS team engaged key principles of empathy, stakeholder engagement and participation, and visualization to ensure states and CMCS were on board and consistently aware of the processes.

Research Design and Methodology

State Selection Criteria

USDS assessed the need, capacity, and potential for impact across all 56 states and territories using key criteria such as Medicaid population size and composition, state processing times, and recent ex parte rates. After reviewing the states against these criteria and working with CMCS, the team identified potential partner states. All states then had an opportunity to express their interest in receiving technical assistance. Based on state interest and the team's need analysis, they conducted interviews with a subset of states to identify and prioritize their state partners.

From these interviews USDS and CMCS selected 8 states to provide in-depth technical assistance. States were chosen based on the following factors: need, potential for impact, and capacity to accept help. Once chosen, USDS provided hands-on technical assistance, which is further detailed below.

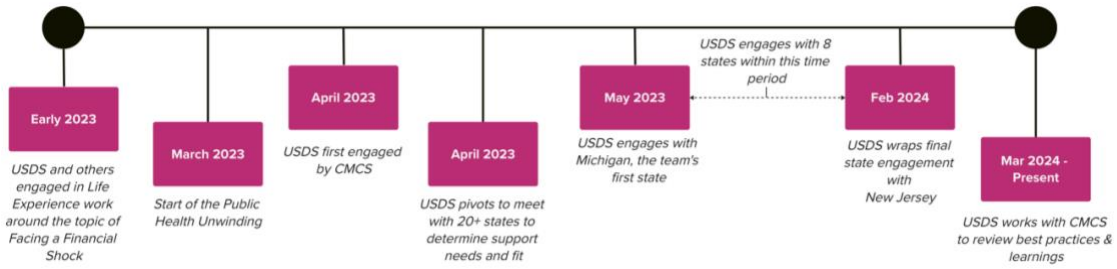

Figure 3: Timeline of USDS engagement with CMCS and states.

Rapid Response Framework

Implementation through Hands-on Technical Assistance

Central to the team's approach was to meet states where they were. This included not just going on site to visit states, but also working with states' policy, operations, and technical teams to understand their constraints and identify areas of impact. This work generally moved through three phases:

1. **Preparatory Research.** Activities included: conducting virtual kick offs with state leadership; collecting data from the state to aid in prioritization of impact areas; and conducting one-on-one semi-structured interviews with state subject matter experts and community-based organizations.
2. **One Week Onsite Sprint.** The team went on the ground in each state for one week and conducted an onsite visit, which included 4-5 days of intensive meetings centered around challenges, then collaboratively with stakeholders working towards identifying potential approaches and creating implementation plans for a few key ones prioritized by the state.
3. **Follow-on Implementation Support**. USDS provided implementation support, general project management, and technical advice for the 2-4 weeks following the onsite.

This multi-stage approach allowed the team to align with the states' needs in advance of onsite engagement, therefore creating the opportunity for more intense and focused work while in the state. It also allowed state actors to ease into the asks USDS and CMCS were making, and building trust and relationships through the process.

Creating Evidence-Driven Prioritization through Data Visualization

CMCS requires states to submit data about their Medicaid enrollments, renewals, and compliance to federal regulations. During the unwinding period, CMCS required

states to report on the number of renewals that were processed on an ex parte basis each month. This meant that many states had data already in hand that USDS was able to use to create an evidence-driven approach to prioritization.

While the ex parte data was useful, it was only one piece of the puzzle. To tell the full picture of where enrollees and states were facing administrative burden throughout the renewals process, the USDS team needed to gather data and understand processes from eligibility workers, technology vendors, and operations teams. USDS brought all of these actors together to collaboratively identify the greatest opportunities to streamline the renewals experience.

To help facilitate and communicate the data requests and analysis, the USDS team used data visualization to support the team and state's understanding of where beneficiaries were exiting the renewal process. Referred to as "the funnel" by the team, the visualization leveraged state data at key points in the customer journey to highlight where these drop off points were occurring. In particular, the funnel allowed the team to identify the gaps between the ex parte process and the multiple steps that happen during the manual renewal, and solidify the extent to which manual renewals should be a priority. The reframing of the data into an end-to-end view of the beneficiary experience allowed for USDS team to leverage their expertise in two workstreams: one focused on implementation of policy into state technological systems, and the other around improvements to the various steps involved in the manual renewal process.

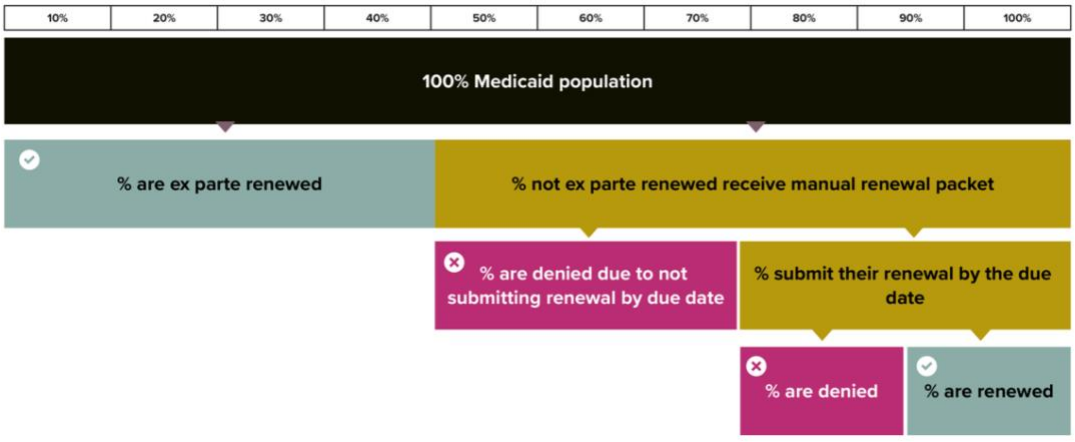

Figure 4. Example of USDS funnel data visualization. Starting from 100% of Medicaid population, the funnel continues down showcasing who is ex parte renewed and the steps involved in the manual renewal process.

Using this data-driven approach was novel for many states, who often approached their data piecemeal rather than holistically. Upon seeing the breakdown of their data in this new way, state leadership and implementors were able to clearly

grasp the extent to which each step in the beneficiary process was either working or needed attention. While the data was essential to understand each state's ex parte challenges, this data driven approach was particularly beneficial in allowing USDS and states to collaboratively align on where to prioritize for interventions within the manual renewal process.

Service Design Methodology

Secondary Research and Heuristic Analyses

The USDS team spent between 2-4 weeks reviewing existing materials related to each state's context before going onsite. This included reviewing news articles, combing through social media, reading state-specific policy documents, analyzing search engine results, in combination with other materials to understand the state's current context.

In parallel to reviewing documentation, USDS would conduct a heuristic analysis of the state's beneficiary-facing materials, including paper forms, websites with applications, and mobile apps. The team would note any potential challenging areas they encountered, such as plain language concerns and with user experience issues.

Finally, using a mystery shopping approach, USDS called into the state's call center to evaluate how renewals were presented to callers. The team reviewed the system's interactive voice response phone trees for plain language, timing, and ease of use. When the team was able to reach a customer representative, USDS tried to get clear answers around how to apply during the renewal period.

This broad and quick review provided general insight into each states' specific Medicaid ecosystem and allowed the team to get a high-level understanding of challenges beneficiaries were experiencing.

Qualitative Research

The USDS team engaged a range of qualitative methods to help inform their strategic approach to implementation. One of the main approaches was the use of semi-structured interviews, particularly with community-based organizations (CBOs). Working collaboratively with state leadership, USDS held conversations with key CBO groups in lieu of talking directly with beneficiaries. USDS chose to do this because of the rapid timeline of implementation and the expert knowledge CBOs have gathered from their front-line work; their breadth and depth of experience was gauged to be a much deeper synthesis than what USDS could understand with only a few interviews done during a week of research.

The team also leveraged CBOs in applied focus groups or listening sessions. These sessions often created opportunities to share USDS created customer journey maps and to work with CBOs to pin point critical pain points. In these sessions, USDS engaged front-line staff, including application assisters and case workers, to confirm the challenges they were seeing beneficiaries have, and sometimes engaged them to speak to solutions that might mitigate administrative burden within the Medicaid program.

Synthesis and Service Design Mapping

USDS created customer journey maps, service blueprints, and data visualizations to support the synthesis, analysis, and communication of information gained within the team, to CMCS, and with the states. These materials often created the basis for conversations and group convenings.

While common with design and technology practitioners in the private sector, government does not always have the capacity to attempt these types of synthesis methods. USDS found that visualizing the outputs of synthesis encouraged those unfamiliar with our methodologies to understand the information faster, engage more directly, and ask more detailed questions than when presenting only in an auditory or text format.

Facilitation and Group Convenings

Central to the team's approach was convening all stakeholders into one room for the duration of the onsite. This included state leadership, eligibility workers, and various implementation staff to review the research done to date and then work together to prioritize and initiate new approaches together. The USDS team leveraged collaborative problem-solving methods and facilitated design workshops to document end-to-end processes in real time as the group discussed potential implications from both technical and procedural standpoints. These conversations began to lay new foundations of how to collaborate across various departments and disciplines, and bring policy and implementation pillars closer together. They also captured decision making in real time and provided concrete references to refer back to over the week.

USDS also led collaborative design workshops to ideate on solutions that emerged from the qualitative research with CBOs. Potential solutions that emerged included streamlining the renewal form design, new forms of communication with beneficiaries, and streamlining the call center experience. This meant constantly bringing those in the room back to the challenges beneficiaries were facing.

Implementation Support: Project Management, Applied Design, and More

USDS supported each state for up to 4 weeks in implementing the outputs from the onsite. Often, this meant supporting conversations between the state and their vendor, translating between policy and technological implementation. As the partner states were managing their normal workload on top of this engagement, USDS often assisted in keeping track of tasks, updating state leadership, and maintaining forward momentum.

Due to the multidisciplinary nature of the USDS team, implementation support also meant occasionally designing the drafts or final version of outputs. This included things like writing copy for new text messages or mocking up new designs for websites.

Impact of Direct Technical Assistance

The USDS team in partnership with CMCS and states were able to improve not only ex-parte results but the customer experience of Medicaid redeterminations in the states they visited.

Common Manual Renewal Challenges

A benefit of working across multiple states and in collaboration with CMCS was being able to see broadly across the federal ecosystem. By visiting 8 states for this work, USDS was able to identify common renewal challenges from a CX perspective.

The challenges the team found included: difficulty in communications and outreach to beneficiaries, particularly around the use of plain language; low access to and low usability of renewal forms; challenging call center access and navigability; lack of modern mobile development and user experience practices; a dearth of feedback loops from beneficiaries to state policy and technology teams; and a shortage of talent at the state level who brought capacity to oversee vendors.

While each state had many of these challenges broadly, the difference between the implementation needs of these were bespoke due to each state having individual eligibility systems, processes, and standard operating procedures. Outputs are therefore not easy to categorize.

Automatic Renewal Impact Numbers

Because the work for both manual renewals and ex parte go hand in hand in providing a full experience to beneficiaries, we consider our impact numbers

collaborative and shared. When combining the work of both approaches in this project, the entire engagement across all states supported an estimated 5 million additional Americans to be renewed for health coverage via ex parte processing in 2024. The project was also estimated to have saved 2 million hours of eligibility worker processing time in 2024 (Medicaid.gov).

For example, within New York, the work on both ex parte and manual renewals helped increase their ex parte rates to 44% based on data reported to Medicaid that compares the data pre-USDS visit and post-visit by April 2024 (Medicaid, n.d.). In California, rates increased to 63%, representing a 30% increase. For the entire engagement, the average change in ex parte rate was 21%, and the average reduction in procedural termination rate for states that engaged with us was a decrease of 10%. See Figure 5 for more details.

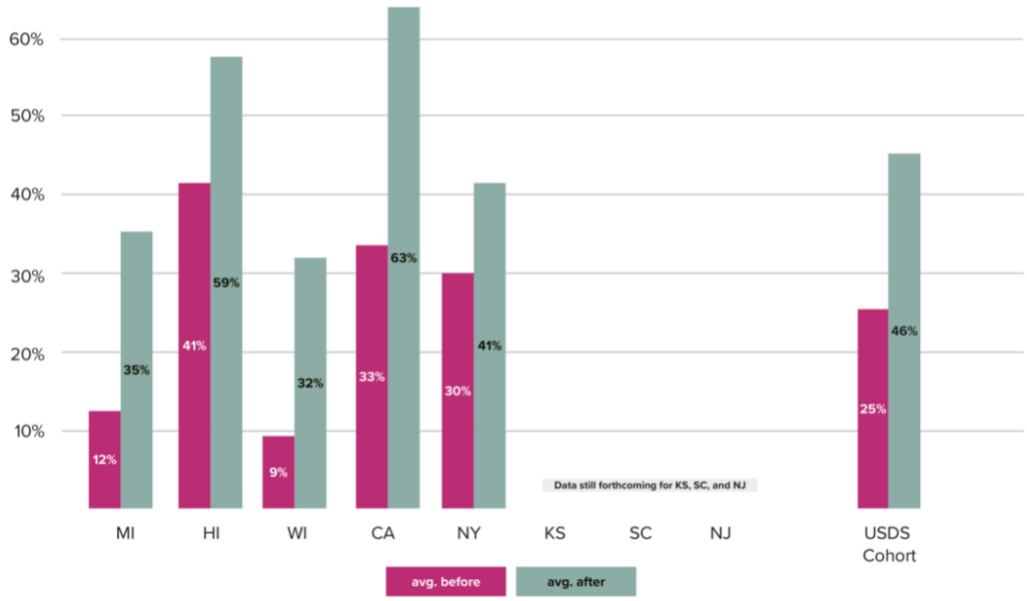

Figure 5. Average percentage of monthly renewals completed ex parte before and after USDS intervention. Average is the average for between May 2023 and April 2024, for which data is available. Data currently is not available for KS, SC, and NJ. HI data is only from June through Sept. from pausing redeterminations due to wildfires. Please note that this analysis is limited, as it does not take into account individual changes made to state systems or monthly population various unrelated to USDS interventions. Data source: https://www.medicaid.gov/resources-for-states/coronavirus-disease-2019-covid-19/unwin%5b…%5dfter-covid-19/data-reporting/monthly-data-reports/index.html.

Discussion

Changing Federal Foundations and Shifting Mental Paradigms

While this work had direct impact on Medicaid recipients and state systems issuing eligibility determinations, much of what the team learned is applicable across federal policies that require state implementation, no matter the program. This points to common learnings that can be leveraged by anyone working in civic tech and guide the space to shifting the foundations and paradigms of how we currently operationalize policy.

Improving State Coordination with Community Organizations

Community based organizations often partner closely with states to help guide enrollees or prospective members to the right information. In many cases, these CBOs often help beneficiaries apply directly. States often have good working relationships with their CBO partners, but often struggle to find ways in which to balance CBO's advocacy with implementation and policy.

Through this engagement, USDS leveraged its expertise in facilitation and design research to help bridge that divide. By hosting design workshops, USDS was able to uplift multiple CBO challenges to state partners in real time with tangible examples; it also encouraged CBOs to brainstorm solutions or approaches to these challenges, giving state leaders opportunities to identify which might be feasible and which were currently out of scope. States have mentioned that opening these lines of direct communication in a collaborative setting have changed the ways in which they think about receiving and prioritizing feedback. As implementation becomes more paramount to delivering on the policy goals, it may behoove federal programs to engage CBOs and other state partners in new ways.

Leveraging Technical Expertise to Shift Paradigms and Change Mental Models

State staff are deeply dedicated to their work. They however face a complex, dynamic, and opaque environment when operating Medicaid. While many are policy experts, wider state staff teams do not typically include technical experts to oversee or support their state vendor's implementation of policy regulations. As a result, policy intent may not match technical implementation, or inefficient and antiquated processes may get built into systems without pressure to initiate changes that would better benefit a technological approach.

USDS was able to help demonstrate and provide new ways of working by showcasing the value of cross-functional implementation support. This technical expertise in logic and engineering systems, service and user experience design, and product management allowed for a bridge between policy intent and technical build; by connecting various subject matter experts with complementary expertise to inform decision making, and clearly articulating and synthesizing the challenges to outline options for senior level staff, teams were able to collectively agree upon goals and plans to initiate work in vendor roadmaps.

These changes have helped to shift state and federal teams' mental models around the ways in which they communicate between stakeholders and how to better engage with vendors. By focusing on clear prioritization and technical implementation, states were able to improve experiences for staff and enrollees. States reported that these changes were lasting—after a USDS engagement, previous resistance to new changes were not as entrenched moving forward.

Institutionalizing Implementation for Long Term Impact

In partnership with CMCS and states, USDS worked to ensure that these approaches and ways of working could continue without their presence. The team also supported states in building their capacity, including giving guidance on how to start a Member Experience Advisory Council or hiring new talent with technical experience.

During this work, USDS was only able to engage with 8 states due to the nature of the rapid response approach. While learnings have been shared and championed by CMCS, there is a stark need for technical expertise across various fields to work with, for, and in support of federal programs and various levels of government. While the USDS team has advocated for hiring employees within programs in addition to working with excellent vendors, more talent is required. Civic technology organizations, vendors, and individual talent are all needed to help improve the experience the public has with federal and state systems.

Conclusion

By examining the foundations of the Medicaid renewal experience, the USDS team hopes that this case study demonstrates how federal programs can be redesigned by challenging long held assumptions, refocusing on systemic implementation, and centering the lived experience of Medicaid members, application assisters, and case workers, to transition towards a more service-driven benefit program.

By centering the lived experiences of case workers, application assisters, and Medicaid beneficiaries, government leaders can foster creating more efficient programs, decreased administrative burden, and better customer experiences for Medicaid beneficiaries and the front-line workers who support them.

Notes

Thank you to our state partners and CMCS colleagues for your dedication, trust, and willingness to partner with us. A special thank you to Max Mazzocchi, Chris Wren, Greg Novick, Navin Eluthesen, Luke Farrell, and Megan Cage for their support on the data in this study, as well as being great partners in the implementation of this project.

About the Authors

Alyssa Kropp is a service designer based in Brooklyn, NY. She has a love for participatory practices like co-design that she uses to help improve the design and delivery of government experiences. Currently she's working with USDS to help strengthen federal programs, policies, and products across the benefits spectrum.

Emily Mann is an applied anthropologist who uses her ethnographic toolkit to improve the delivery of critical government benefits. At USDS, Emily has had the opportunity to work alongside civil servants to improve access healthcare benefits at multiple federal agencies including the VA, CMS, and SSA.

Heather Myers is a product designer who loves using design and research for the public good. She previously worked at USDS where she used her skills to address big challenges such as responding to the COVID-19 pandemic and reducing administrative burden for low-income Americans on Medicaid.

Izzie Zahorian is a Design Lead at USDS where she partners with agencies across government to improve the accessibility and usability of their services. She currently leads design research and service design on the Facing a Financial Shock charter established pursuant to President Biden's Executive Order on customer experience (Executive Order 14058).

References Cited

42 U.S.C. §435.916(a)(2). https://www.ecfr.gov/current/title-42/part-435/subject-group-ECFR0717d3fdf4a090c#p-435.916(b)

42 U.S.C. §435.916(a)(3). https://www.ecfr.gov/current/title-42/part-435/subject-group-ECFR0717d3fdf4a090c#p-435.916(b)(2)

Corallo, Bradley and Moreno, Sophia. 2023. "Analysis of National Trends in Medicaid and CHIP Enrollment During the COVID-19 Pandemic." Kaiser Family Foundation. https://www.kff.org/coronavirus-covid-19/issue-brief/analysis-of-recent-national-trends-in-medicaid-and-chip-enrollment/

Grossmann, Matt, Donald Moynihan, Pamela Herd, and Niskanen Center. 2023. "How Administrative Burdens Undermine Public Programs." Niskanen Center. https://www.niskanencenter.org/how-administrative-burdens-undermine-public-programs/

Heard, Pamela, and Moynihan, Donald P. 2018. Administrative Burden: Policymaking by Other Means. Russell Sage Foundation.

Medicaid.gov. 2023. "January 2023 Medicaid and CHIP Enrollment Snapshot." Accessed August 1, 2024. https://www.medicaid.gov/medicaid/national-medicaid-chip-program-information/downloads/january-2023-medicaid-chip-enrollment-trend-snapshot.pdf

Medicaid.gov, n.d. "Medicaid and CHIP Enrollment Trend Snapshot." https://www.medicaid.gov/medicaid/program-information/medicaid-chip-enrollment-data/medicaid-and-chip-enrollment-trend-snapshot/index.html

Medicaid.gov. n.d. "Monthly Data Reports." Accessed August 1, 2024. https://www.medicaid.gov/resources-for-states/coronavirus-disease-2019-covid-19/unwinding-and-returning-regular-operations-after-covid-19/data-reporting/monthly-data-reports/index.html

Medicaid.gov. n.d. "Monthly Medicaid & CHIP Application, Eligibility Determination, and Enrollment Reports & Data." https://www.medicaid.gov/medicaid/national-medicaid-chip-program-information/medicaid-chip-enrollment-data/monthly-medicaid-chip-application-eligibility-determination-and-enrollment-reports-data/index.html

Medicaid.gov. n.d. "Unwinding and Returning to Regular Operations after COVID-19." Accessed August 1, 2024. https://www.medicaid.gov/resources-for-states/coronavirus-disease-2019-covid-19/unwinding-and-returning-regular-operations-after-covid-19/index.html

Office of the Assistant Secretary for Planning and Evaluation, U.S. Department of Health and Human Services. 2022. Issue Brief No. HP-2022-20. "Unwinding the Medicaid Continuous Enrollment Provision: Projected Enrollment Effects and Policy Approaches." https://aspe.hhs.gov/sites/default/files/documents/dc73e82abf7fc26b6a8e5cc52ae42d48/aspe-end-mcaid-continuous-coverage.pdf

Pahlka, Jennifer. 2023. Recoding America: Why Government Is Failing in the Digital Age and How We Can Do Better First edition., Metropolitan Books/Henry Holt and Company.

Pahlka, Jennifer. 2020. "Digital Is Not Always Better Than Paper." *Medium*. https://pahlkadot.medium.com/digital-is-not-always-better-than-paper-1f06dac2c900

The White House. 2021. "Executive Order on Transforming Federal Customer Experience and Service Delivery to Rebuild Trust in Government." www.whitehouse.gov/briefing-room/presidential-actions/2021/12/13/executive-order-on-transforming-federal-customer-experience-and-service-delivery-to-rebuild-trust-in-government/

Usds.gov. n.d. "Mission." Accessed August 1, 2024. https://www.usds.gov/mission

An Apprenticeship in Attention for a Time Between Worlds

INDIGO WELLER, *Stripe Partners*

PechaKucha Presentation Abstract

The aim of this Pecha Kucha is to describe the power and beauty of intergenerational sense-making at the end of life and what it asks of us as ethnographic researchers in times of upheaval and transition. Reminiscence work asks ethnographers to reexamine: 1. our conceptions of expertise; 2. The supposed 'depth' of in-depth interviews; 3. The need to welcome serendipity and silence, whereby the guiding measure of an ethnographic encounter is the quality of presence and play, not the extractive rush to 'actionable insight.'

Indigo Weller is a consultant and narrative researcher with a background in patient advocacy and medical education. Prior to joining Stripe Partners, he served as a development editor for patients writing illness memoirs and legacy projects. He is a facilitator for the European Reminiscence Network, an international dementia arts initiative, and holds degrees in Bioethics, Narrative Medicine, and Creative Writing.

Forget Nature, Forget Humans: Redefining Our Approach toward Biodiversity

MARC-ANTOINE MORIER, *frog*

PechaKucha Presentation Abstract

While human-centred approaches have enabled considerable progress to be made in the design of innovative and socially responsible solutions, they fail to take account of a decisive factor in maintaining our societies. Medicine, food, drinking water and breathable air... Biodiversity underpins many of the "services" that "nature" provides free of charge, and on which we are largely dependent. With biodiversity collapsing, how is it that we take so little account of it in our work? So how can we include living non-humans in our investigations and creations?

Marc-Antoine is a trained anthropologist working at the intersection of design, social sciences and fiction.

Glitch in the Matrix: What Broken Screens Reveal About Digital Matter

PIERRE BONNIER

PechaKucha Presentation Abstract

As digital screens increasingly populate our cities, we are more and more often confronted with their malfunctions. Through a poetic exploration of the visual effects produced by these broken screens, I offer an opportunity to reflect on the place of materiality and maintenance in the user experience.

With a background in social and cognitive sciences, Pierre has been a freelance ux researcher for almost 10 years. His two favorite fields are large-scale industry and public transport. Passionate about urban environments, he also co-founded an agency specializing in the application of behavioral sciences to architecture and urban planning. He works mainly in France, where he lives between Paris and Marseille.

A History of EPIC in 20 Unwritten Proposals

EVAN HANOVER, *Conifer Research*

Given we will be up the road from Hollywood, I explore the history of this conference – its themes and contexts – by pitching a series of unwritten submission proposals, one for each year. Reflection and humor are features throughout, and essential at moments of such achievement (20 years!) and anxiety (see: the world). This may be more performance than presentation. I sure hope it works.

Evan is linguistic anthropologist and strategist with Conifer Research in Chicago. He has conducted business and consumer research for 25 years in industries ranging from healthcare and financial services to automotive and boating. His passion for combining ethnographic insight with performance is perfectly encapsulated in the Pecha Kucha format. 2024 marks 10 years since he first stepped onto an EPIC stage.

“How Can I Help You?” Exploring Empathy Work in Customer Support

NICOLE LABORDE, *Sutherland Labs*
MARISE PHILLIPS, *Sutherland Labs*

PechaKucha Presentation Abstract

In this PechaKucha, we explore the role of human connection and empathy in customer support work – a setting that is often fraught with frustration. As ethnographers, we have the opportunity to work with companies to understand and improve employee and customer experience and needs, and to leverage the emotional labor provided by agents answering customer calls.

A cultural anthropologist, Nicole brings over 20 years of research experience in international and domestic settings, often among vulnerable populations and on uncomfortable topics. She is a senior design researcher and ethnographer at Sutherland Labs where she works with clients on strategic questions regarding a wide range of existing and developing products and technologies. Out of the office, she loves thrift shopping with her daughters, practicing karate, and caring for too many pets.

Marise leads the service design practice at Sutherland. Her experience in participatory design facilitation empowers cross functional teams to co-create empathy maps, ecosystems, journey maps and experience blueprints. At Wells Fargo, she managed the customer insights team shaping digital banking experiences whilst helping the bank become a more agile organization. In her spare time, she leads ethnographic initiatives to promote ecosystem restoration in the San Francisco Bay.

How to Avoid Path Dependency: Learnings from Neuroscience, Physics and Organizational Theory

CYRIL MAURY, *Stripe Partners*

PechaKucha Presentation Abstract

In 1814, the polymath Pierre Simon Laplace posited that "We may regard the present state of the universe as the effect of its past and the cause of its future". This presentation will examine this core idea – the fact that what will happen next is determined by what happened before – from different vantage points: neuroscience (Hebbian plasticity, Canalization), economic theory (Bayes's Theorem) and organisational theory. How and how much does legacy narrow down what is possible next? How can we, at the individual and collective levels, carve out the space for alternative, unexpected futures?

Cyril is a seasoned strategy and innovation practitioner, whose expertise centres on helping tech companies understand their users to develop better products. Having lived in Latin America, the Middle East and Europe, he particularly enjoys untangling the operational, organisational and cultural complexities inherent in adapting tech products to emerging markets. Cyril is a Science Po (PoliSci) and HEC Paris (MBA) graduate.

"If Today's Problems Are So Urgent, Why Are We So Stuck?"

MEG KINNEY, *Bad Babysitter Productions*

PechaKucha Presentation Abstract

We are surrounded by messages urging us to see the suffering, injustice, and corruption of both the natural and built world. But today, people seem to be gripped by passive despair. Why is that? Perhaps we are weighted by a loss of optimism. This presentation questions if optimism is misunderstood and if, as ethnographers, we are getting in our own way and underestimating it as a powerful tool of social transformation. Optimism requires us to examine our own emotional intelligence, capacity for change, and belief in our own human potential.

Meg is a partner at Bad Babysitter – a brand strategy consultancy focused on using an ethnographic approach and documentary storytelling style to humanize stakeholders and contextualize brand meaning across the value chain. Clients span Fortune 500 brands and retailers, Climate Tech startups, emerging sectors, and marketing agencies. She has a Masters in Natural Resources & Sustainability Leadership to help brands and agencies find their breakthrough narrative for this moment.

Kites Rise Against the Wind: 20 Years of Research

LEAH MCDOUGALD, *McDougald Research*
MACKENZIE SHAW, *McDougald Research*

PechaKucha Presentation Abstract

Across two decades, our journey as research practitioners has been undeniably dynamic, defined by moments of recognition and rejection. As we reflect on the evolution of our industry amidst shifting social landscapes, corporate upheavals, and evolving economic demands, it's crucial to reexamine our collective role. Finding harmony within our journey, this narrative portrays how, like the flight of a

kite, we remain steadfast against the gusts of change, anchored by our unyielding commitment to enhancing lived experiences through design and innovation.

With a degree in Sociology and Women's Studies from The Ohio State University, Leah's early career in the non-profit sector involved leading empowerment courses for survivors of abuse, shaping her commitment to understanding human needs. Her journey led her to the forefront of design research, conducting international ethnographic studies in Asia and Europe, providing a global perspective to her work. In 2014, Leah founded McDougald Research.

Mackenzie is an experienced strategist, researcher, design thinker, and problem solver with 20+ years of professional experience. Armed with an Industrial Design degree from The Ohio State University, her career spans both corporate and agency environments where she has conducted global research and led design, research, and strategy teams. With the ability to navigate the intersection of user needs and business objectives, she has an expertise in aligning stakeholders and influencing strategy.

Sankofa; Look Back, Side-eye the Future

SASHA OFORI, *YUX Design*

PechaKucha Presentation Abstract

This PechaKucha is an invitation into my journey, probing questions such as: What is the true meaning of empathy; self-empathy. How does it manifest, especially in the African context? How do I safeguard my well-being while being fully present for those I serve and work with? What lessons from my experience of being a full-time caregiver can inform my approach to "empathy" as researcher? This a practical inquiry of a researcher; researching from within using methodologies inherent to our work to initiate these conversations. Because, in looking back and reflecting within - we might have a glimpse of the future ahead.

Sasha Ofori is a Ghanaian-based UX Researcher at YUX Design, specializing in the application of design thinking and HCD methodologies across sectors; health, education, agriculture, and technology, leveraging research to create positive impact in Africa. As a young professional, these experiences has ignited a profound passion in her to delve into the nexus of empathy, self-discovery, and mental health within the realm of research in Ghana. Sasha actively documents her learnings on her Medium.

Saving (Tattoo) Clients from Themselves

DUSTIN KISKADDON, *JP Morgan Chase*

PechaKucha Presentation Abstract

Drawing on direct experience with professional tattooing, I explore the art and strategy of saving clients from their own bad ideas. Doing well for others, it turns out, means doing well for yourself.

Dustin Kiskaddon is an applied sociologist who studies the intersection of culture, bodies, and money. His ethnographic memoir, Blood and Lightning: On Becoming a Tattooer explains how tattooers think, feel, and act at work. It's based on Dustin's experience working as a professional tattooer, and it was published by Stanford University Press. Learn more at https://www.dustinkiskaddon.com