

# EPIC2020

Ethnographic Praxis in Industry Conference

CATALYST

## Where Can We Find an Ethics for Scale?

### How to Define an Ethical Infrastructure for the Development of Future Technologies at Global Scale

IAN DULL, *ReD Associates*

FANI NTAVELOU BAUM, *ReD Associates*

THOMAS HUGHES, *ReD Associates*

*Despite companies facing real consequences for getting ethics wrong, basic ethical questions in emerging technologies remain unresolved. Companies have begun trying to answer these tough questions, but their techniques are often hindered by the classical approach of moral philosophy and ethics – namely normative philosophy – which prescribe an approach to resolving ethical dilemmas from the outset, based on assumed moral truths. In contrast, we propose that a key foundation for ‘getting ethics right’ is to do the opposite: to discover them, by going out into the world to study how relevant people resolve similar ethical dilemmas in their daily lives – a project we term ‘grounded ethics’. Building from Durkheim’s theory of moral facts and more recent developments in the anthropology of morals and ethics, this paper explores the methods and theory useful to such a mission – synthesizing these into a framework to guide future ‘grounded ethics’ practice.*

*Keywords: Ethics, Technology, Methodology, Moral Facts*

## INTRODUCTION: THE ETHICS FRONTIER IN TECHNOLOGY

Debates around the moral permissibility and the ethical implications of technology have long been ongoing in global fora (Moss and Metcalf 2019). In recent years, as technology has become a foundation for everyday life globally, such debate has become commonplace in public, academic, and practitioner fora as part of broader ‘tech-lash,’ leading to its questioning in political discourse and hearings, frequent publications in major news organizations, academic fora including EPIC (see Tamminen and Holmgren 2016; Moss and Schüür 2019), and discourses within the technology industry and organizations – what technology news outlet *The Verge* has termed, “an ethics explosion” (Vincent 2019).

This renewed attention has centered on a growing number of issues often in response to scandals around new technology practices, including but not limited to: algorithmic bias (i.e. in mortgage reviews or prison sentencing, see Eubanks 2020; Lee et al. 2019; Anderson et al. 2019), data collection (i.e. surrounding the Cambridge Analytica scandal or location tracking), privacy (i.e. the ‘right to be forgotten’), freedom of speech, and exploitative practices (i.e. surrounding Uber and the ‘gig economy’). Furthermore, momentum sparked by such debates has begun to spill over into discussions of future applications of new technologies, including augmented and virtual reality (AR/VR), more generalized artificial intelligence, self-driving technology, and brain-machine interfaces (i.e. Neuralink). With many major technology players becoming embroiled in unforeseen ethical controversies as a result, the industry is now increasingly concerned about the ethical implications of product

decisions and future technology investments, in terms of future product usage, adoption of new products & technologies, brand equity, and even employee satisfaction.<sup>1</sup> Recent commentators have noted how this has even begun to enter the language of corporate financial risk: Metcalf et al. (2019, 451) note how Alphabet, parent company of Google and Waymo, in 2018 reported to investors that AI products could “raise new or exacerbate existing ethical, technological, legal and other challenges, which may ... adversely affect our revenues and operating results.” That many of these companies and their technologies operate at (and require) enormous, often global, scale has also meant an amplification of concern around their potential moral effects. The bigger the impact, the higher the ethical stakes.

While concerns about the moral permissibility and ethics of technology development are notable, thus far, calls for more ethical practices have too often stopped there. Though this no doubt bears reference to the scale of change being demanded of technology players, it nevertheless leaves a gap for technology practitioners who seek to heed calls for change and implement more ethical practices. As such, this paper aims to argue not *if* ethics are needed, but *where and how* they can be studied. This paper will present a framework for a ‘grounded’ approach, employing anthropological theory and methodology for this pursuit.

However, the pursuit of understanding the ethics of new technological development is by no means straightforward. To take just one example: when introducing new products and investing in new technologies and applications, technology companies often take guidance from the legal limits currently in force. Yet just as new technologies often engender new behaviors, entering new technological terrains also means entering new moral and ethical ones. Facebook’s Live product, for example – launched in 2016 to enable live streaming of content to Facebook ‘friends’ and a wider audience – quickly led to freedom of speech and censorship concerns, specifically around the broadcasting of violent and terrorist acts (Issac & Ember 2016). Facebook’s standards likely met legal expectations prior to Live’s launch – but what of the social limits around what *should* be shared freely and without censorship, and the legal limits that may be drawn going forward? Today’s legal (and moral) acceptance offers uncertain guidance for the boundaries of future moral outrage.

With a new generation of cultural-paradigm-shifting technologies on the horizon, technology companies need better tools to understand what moral standards consumers and society will hold them to – and where boundaries are likely to fall in the future. As a foray into this domain, we propose a project of ‘grounded ethics’ to help companies understand users’ moral intuitions: When do technologies risk overstepping people’s moral boundaries, their sense of right and wrong and result in public reproach, i.e. with Facebook’s Live or Google’s Glass? How do we locate people’s boundaries around technologies like voice recognition, in light of social processes of normalization? What ethics should be scaled into agentive products themselves, like algorithm-driven feeds? How much can moral boundaries truly be global, and, if not, where should the lines be drawn? While morality and ethics have traditionally been the domain of philosophers and deduction, we argue that it is precisely the anthropologist’s inductive approach that is needed to develop the problem-specific and practical frameworks needed to guide the executives, technologists, and designers who will make future decisions. Building on Durkheim’s concept of the moral fact and recent work in the anthropology of morals (see Fassin 2006, 2014), we aim to propose an initial framework for how to study the moral landscapes surrounding future technology applications to create the ethical infrastructures for future technology development.

## **BUT WHICH ETHICS? A THEORETICAL CHALLENGE FOR THE PRACTICE OF SCALABLE ETHICS**

Calls for more ethical approaches to the development and deployment of new technologies have largely focused on what technology companies *should* and *should not* do. This is to say that the primary calls for more ethical action have largely been – as has long been standard in moral philosophy and ethics, and thus, society at large – *normative*: they tend to prescribe how the future behavior of technology companies should look, and suggest the rules or ‘tests’ according to which they should function, or consequences which they should not permit. General business ethics have tended to follow the same path, with the additional step of aligning business goals (in this case, often considered the ‘norms’) to the individual behaviors of employees (see Sims 1991).

To their credit, major technology firms like those of Silicon Valley have followed suit – and have begun to respond in myriad ways. As Metcalf et al. (2019) have documented, ‘ethics owners’ have proliferated in Silicon Valley companies – ranging from individuals responsible for developing ethical procedures and supervising teams to ensure ethical practices to even making corporate ethics a personal mission. Others, like Google, have famously hired ‘in-house philosophers’ to handle philosophical quandaries that may come up in the development of its products and technologies (VentureBeat 2011). These personnel-driven solutions have come in addition to more traditional and longer-standing practices, like the following of local and international regulations for safety and ethical testing, and the now common practice of consumer surveys, A/B testing, and other large-scale quantitative instruments to determine consumer interest and drive more engagement or use. Work has continued outside of corporations themselves, with universities mandating ethics for software engineers and computer scientists (Fiesler 2018), and groups of practitioners and scientists calling for codes of ethics, statements of principles, or bans on certain practices, notably the development of autonomous weapons (see Sample 2018; Future of Life 2015).

Such reinvigorated practices are relatively new within major technology firms, making their efficacy uncertain in the short-term. Nevertheless, the continued appearance of new ethical challenges to technology companies’ products and practices – most recently surrounding the censorship of ‘fake news’ and misinformation during the 2020 COVID-19 pandemic and US Presidential Election (see Frenkel et al. 2020; Warzel 2020a) – points to significant opportunities for improvement. These discussions are all the more urgent against the backdrop of Silicon Valley’s increasingly outsize influence in shaping the public sphere – the gatekeepers to how society accesses and experiences information – and even more so, when contrasted with increasingly limited (or ineffective) government and civil society institutions. Even when major tech companies have taken proactive stances on ethical issues – like Facebook’s announcements around political ads in the 2020 US Presidential Election (see Isaac 2020; Warzel 2020b) – many have observed a tacit signaling of their increasing control. (And, we might add, a tacit acknowledgement that their products themselves adopt an ethical stance one way or the other). As one journalist wrote of a similar Facebook pledge in Germany’s 2017 elections: “It’s a declaration that Facebook is assuming a level of power at once of the state and beyond it” (Read 2017). The combination of technology players’ growing necessity, power, and inconsistent performance has only reinforced initial calls for more ethical actions.

Beyond continued demands for ethical accountability, commentators have also observed many challenges facing the practice of ethics within technology companies. While ethics remains a buzz word throughout Silicon Valley, Metcalf et al. note in their study of ‘ethics owners’ that many everyday practitioners – the designers, engineers, managers, executives, and more who make or drive product decisions – “are not yet moved by ethics” (2019, 453). In more plain terms, this is to say that ethics do not enter consciously into the day-to-day practices of product development. Moreover, the authors note that, in such climates, ethics owners’ mandates (*vis-à-vis* compliance, CSR, and others) and their roles organizationally (i.e. who they report to, how they can influence projects) are unclear, often leaving ethical concerns dangling within organizations. Even when voiced, ethical qualms are furthermore drowned out by common Silicon Valley discourses like ‘technological solutionism’ – the belief that *better* technology will resolve ethical problems – or ‘market fundamentalism’ – the belief that market indicators trump ethical decisions or that consumer demand (i.e. continued use) proves moral acceptability – have a tendency to downplay or entirely undermine otherwise legitimate concerns. In this context, weighing complex ethical decisions becomes “doing ethics” (2019, 453) – yet another task to be performed in the course of product development. That, they argue, points to a challenge as to if ethics can coexist within the current structures and internal logics driving firms.

We agree; yet these challenges represent only the organizational dimensions of the practice of ethics in technology companies. Conspicuously lacking from discussions has been the question of *which ethics* technology companies should abide by to begin with. This is to say: when technology companies bring in corporate philosophers, appoint ethics owners, or create ethics boards, which systems of ethics should they bring with them or judge proposed projects and products against?

From that perspective, calls for ethical accountability in technology companies have been quite unspecific – and it is here where a *normative* approach to defining ethics can fall short. The choice between normative ethical systems – e.g. between utilitarianism, Kantian ethics, care ethics, virtue ethics – ironically leaves open how ethical quandaries are to be interpreted and resolved, and does so in the absence of input from the people ethical decisions will affect. That challenge is not only theoretical, but empirical: normative standards for defining ethics have failed to deliver meaningful guidance on moral permissibility and ethical action, notably on three fronts:

1. *Lack of consistency*: While individual corporations have attempted to define their own normative ethics to guide corporate behavior, when looking across the technology sector, these individualized approaches to normative ethics have created different and competing systems – yielding anarchy, rather than a consistent approach to ethics. Recent research into 84 AI ethics guidelines from companies and organizations around the world found that “no single ethical principle appeared to be common to the entire corpus of documents” (Jobin et al. 2019). When each company selects their own normative moral foundations and ethical principles, as opposed to deriving them from prevailing moral and ethical tides, they are contributing to an overall climate of ambiguity that ultimately undermines the project of an ethics of technology development in the first place (see D’Ignazio and Klein 2019).

2. *Lack of nuance & context-specificity*: As an approach founded on *a priori* truths, normative ethics tends to categorical assertions, and technology is no exception – whether aspiring to full transparency with consumers, asserting or denying the primacy of privacy, or defining what tasks machines should and should not take on. In practice, few morals operate in such black-and-white terms. Recent ReD technology studies have explored the boundaries of what types of data collection can be acceptable. Many informants were unaffected by their data being collected – surprisingly even for what one might consider ‘sensitive’ data, like home addresses. Yet when faced with unexpected voice or video data collection – like a dubious beep from an Amazon Echo during a private, political conversation at home, or unexpected filming in public – reactions were visceral, and anger immediate. In that case, ‘privacy’ was not so much an absolute value, but a contextual one. Without the right qualifiers in place, normative principles can be controversial or counterproductive to commercial aims. Finding the right context and execution for a technology can drastically modify its moral and ethical dynamics.
3. *Lack of future-proofing*: In asserting one way to understand the morality of the world they occupy, technology firms’ normative assessments of ethics fail to capture the shifting nature of moral systems, or account for how the technology they produce can shape moral systems. This can work to both the benefit and harm of companies’ ambitions. To take a positive example, the past decades have created a major shift in public intuition around ‘strangers online’ – from dangerous to 50 million people on Tinder. Had the creators of Tinder only followed dominant moral codes surrounding the early internet, they might not have found the same success. Yet cautionary tales also abound: While most photography was accepted and prevalent in the early smartphone era, Google’s Glass overstepped these boundaries by turning glasses – and by extension, the body – into a camera, therein reimagining norms around privacy. A meaningful picture of moral and ethical future action does not necessarily emerge from the standards in front of us today.

Taken together, these challenges point to the societal and corporate risks which a plethora of *normative* assessments of ethics in technology development can create. So why do companies still run these risks, especially after investing time and resources to develop their products and new moonshot technologies? This is not due to a lack of effort, but due to a methodological fallacy. There is a centuries-long tradition of armchair, top-down ethics: philosophers – and now corporate philosophers – have sat around, thinking about the right and wrong ways to live, based on virtues, the consequences of our actions, or the deontological imperative. But they tell only half of the story. Normative, top-down ethics has given a multitude of rulebooks for how one *should live*, but it does not say much about *how we do live*. Just because we know that lying is wrong, it does not mean that we do not lie. And as cases like the Milgram experiments clearly have proven, people rarely meet even their own standards of virtuous life. Moral philosophers have been long baffled and divided on how to trace these moral facts – and the many others – in society and how seriously to take them. As the above goes to show, the complexity and stakes of ethics decisions are too high for individual stances on morality.

Yet with the right tools, the picture can become simpler: rather than eliding the realities of moral facts in society to describe what we *ought to do* (as moral philosophers have), we suggest a knowledge of *how* people live morally and what they *will* and *will not* accept as the basis for defining an ethics for technology development. Just as moral philosophers likely should not look into absolute moral truths in lived daily life, corporate ethicists should not look for moral facts in the theoretical realm.

## FROM THE ARMCHAIR TO THE BAZAAR: ‘GROUNDING’ ETHICS IN LIVED MORAL FACTS

We have traced the cause for these risks down to the method of defining ethics itself. Companies tend to *theorize* what people think is ethical instead of *discovering* how ethics are navigated. In this paper, we propose an alternative approach that avoids these risks – by defining a ‘grounded ethics,’ designed to study and understand the nuances around ‘moral facts’ that govern the aspects of life a technology could change.

Durkheim still looms large in any discussion of moral facts. We embrace his view of ethics as grounded in social life, facts to be discovered through how people think and behave. While temporally far from current debates around the ethics of technology, Durkheim’s theory arose in the context of the social upheaval accompanying the industrial revolution (Laidlaw 2017). As such, his theory is attuned to understanding what is moral and ethical as both 1) defined by the realities of the social world (as uncoupled from normative, religious mores), and 2) flexible with regards to social changes, for example, those shaped by new technologies. As Durkheim proposed, we suggest ‘grounding’ a development of ethics in uncovering ‘moral facts’: the pillars people use to shape a sense of living a morally good life – which may be observed and studied in culture (Durkheim [1924] 2010). In broad strokes, in his view, ethics are observable through sanctions – social consequences to rule-violating actions. The actions that would trigger these sanctions delineate what is morally acceptable and what is not. The upside of adopting Durkheim’s view is that it clearly points to a domain of study and observation – the social rules that are observed and the sanctions to which they give rise.<sup>2</sup>

Anthropology has turned away from Durkheim’s moral facts in the past, due to some commentators’ interpretations of Durkheim’s focus on social sanctions as representing overly fixed norms (Laidlaw 2014), or norms *simpliciter*. Didier Fassin, however, offers a helpful argument to reposition Durkheim’s understanding of social sanctions to also include complex individual negotiations:

Durkheim himself had a more sophisticated and somewhat ambiguous theory than what is often simplified by commentators, including Laidlaw (2014, 21), who writes that the French sociologist “ended up with a conception of morality as thorough law-like as Kant’s, but with obedience to the law naturalized into the smooth functioning of a well-engineered mechanical system,” thus ignoring what Durkheim ([1924] 2010, 17) clearly asserts: “In opposition to Kant, we shall show that the notion of duty does not exhaust the concept of morality,” since “to become the agents of an act it must interest our sensibility to a certain extent and appear to us, in some way, desirable.” Such an act “cannot be accomplished without effort and self-constraint” and “is not achieved without difficulty and inner-conflict” — a

language not so far removed from the contemporary anthropology of ethics.  
(Fassin 2014, 430)

In this light, Durkheim's notion of a moral fact asserts not only a distinction of social norms from Kant's norms *simpliciter* – thereby creating space for cultural, individual, and temporal variation – but also locates discovering variations, and their future directions, in individual experience. The grounded approach to ethics we propose is built on this Durkheimian proposition that moral facts are to be discovered in the lived reality of human life: in the daily behaviors and choices of individuals, the symbols they respond to, and sanctions they recognize, as they navigate towards the right or wrong side of virtue. We furthermore believe that identifying these moral facts is at its most feasible and productive when it is focused on the individual experience of choice and conflict. Ideally, this approach would be supplemented with a larger understanding of the historical factors that give rise to the norms that constrain ethical action by creating implicit and explicit sanctions.<sup>3</sup> We believe, however, that a focus on the individual experience of moral decision-making is more valuable to building an understanding of the moral boundaries that are likely to shape future technology products. In the same way that learning about traffic laws does not teach us everything we need to know about acceptable driving conduct, understanding social norms does not tell the full story about acceptable moral behavior.

To put our stance succinctly, we are describing a grounded ethics framework with three necessary features:

1. *It is bottom-up.* We are interested in understanding the nuances, shortcuts, trade-offs and irregularities in how people experience the moral systems they inherit and create.
2. *It assumes a scale of flexibility of moral facts.* We believe that moral facts are malleable and subject to change by the same forces that forged them in the first place, be it social, political and religious factors, tradition or biases of moral psychology.
3. *It is application-dependent.* Finally, some housekeeping. This framework is not developed with the intention to be applied to the moral character of a society or group at large. This is too big of a project and not helpful for the purposes we have in mind. Rather, we are operating under the assumption that we can secure depth and nuance by focusing on the social phenomena a given technology has the potential to transform.

While this latter proposition could be questioned on the grounds of being too narrow – how could we understand moral obligations around privacy, without understanding the context of morality more generally? – given the resource constraints placed on practitioners outside the academy, we view undertaking such complete studies of morals and ethics to be too ambitious to be practical. Rather, as we will detail in the process of outlining a research approach for developing a 'grounded ethics,' existing ethnographies of relevant societies need to suffice to provide the moral backgrounds against which more focused questions of technology applications can be studied.

## DISCOVERING A ‘GROUNDED ETHICS’ IN PRACTICE – A FRAMEWORK

A ‘grounded ethics’ is then 1) the most productive approach to identifying the moral foundations to guide the development and deployment of technologies at scale, and 2) clearly grounded in a society’s moral facts, but especially in the daily behaviors, choices, and trade-offs faced by individuals living in those societies. How then should we as practitioners working with technology companies practically seek out and *discover* a ‘grounded ethics,’ for the real-world technology problems we are likely to face?

At this phase in applied social science research and ‘UX,’ a range of tools – from product-centered ethnography to usability observations and attitudinal/behavioral surveys – would normally be seen as the defaults for exploring new product and technology innovation challenges. Yet, as Amirebrahimi (2016) has already discussed at length in the EPIC forum, while these methods have proven successful at identifying new commercial opportunities through observed and emerging behaviors, promising attitudes, and a willingness to adopt or pay, such approaches have come to be co-opted and oversimplified in practice – too much so to address the “difficulty and inner-conflict” (Fassin 2014, 430) that accompany moral negotiations. To quote one of Amirebrahimi’s ‘UXer’ informants, these methods too often “don’t get at the very real issues” (Amirebrahimi 2016, 87) and by Amirebrahimi’s account reduces lives into “only [a person’s] moment of use” (2016, 89). To combine this critique with Metcalf et al.’s critique of ‘doing ethics’ (2019) would suggest that using UX approaches reduces the complex moral choices of individuals and their societies to a simple review of their “moment of acceptance” of a new product or technology – devoid of the context(s) in which such acceptance may occur, the moral ‘costs’ or ‘burdens’ of such decisions, and how flexible such moral facts are for people. That leaves the nuances of moral facts quickly reduced to binary – yes / no – permissibility.

In line with this approach, we believe that any study of ‘grounded ethics’ for technology development at scale must deeply explore several layers: (a) the cultural foundations of the targeted societies, (b) the ‘virtuous’ phenomena likely affected by the new technology, (c) the ethical interests of different social groups, and who is the moral ‘user’ in each case and (d) contexts relevant to *where* ethics may be applied, e.g. physical sites, varied social groups and (e) moral notions around monetization. Across these, we suggest that the foundations for defining a ‘grounded ethics’ for new technologies lie in understanding the social systems, moral intuitions and dilemmas, and visceral reactions around the underlying social phenomena a given technology has the power to shape. We suggest that uncovering these foundations will require incorporating methods beyond the conventional applied social science toolbox, like social experiments and experimental philosophy. Our hope is that this framework will be useful in rendering a prescriptive picture of the moral landscapes in which companies balance ethical trade-offs.

### Cultural Foundations

As many anthropological studies of morality and ethics (see Laidlaw 2014; Widlok 2004) have made clear, the morals and virtues of different societies can radically differ. That extends deeply into the fundamental assumptions about ‘who’ can make moral judgments and how they can be negotiated – as Kenneth Read (1955) noted of Gahuku-Gamu morality,



where the lack of personal individuality changes the types of moral relations in place from individual to distributive. Indeed, such claims have been foundational in relativism in anthropology. While many major tech firms will likely not consider Papua New Guinea to be a leading market for new technologies, understanding the ‘playing field’ for what is permissible with new technologies should be grounded in an understanding of such ‘ontological’ differences across relevant global markets’ spectrum. Many commentators have noted meaningful, if less extreme differences, between individualistic Western societies that champion free choice and those of former Soviet nations or collectivist nations of East Asia (see Widlok 2004; Hefner 1998). Exploring these differences as a minimum are not only to avoid allegation of purely ‘Western’ notions of morality, but to identify the different ontologies and processes that govern moral decision-making in each. Given the resource constraints often placed on similar studies, we would suggest that such fundamental exploration can be guided by existing recent ethnographies of different cultures.

### **‘Virtuous’ Phenomena**

Rather than focusing inquiry on the technology itself, in line with Widlok’s (2004) framing of an anthropology of virtue, any study of ‘grounded ethics’ must explore precisely the ‘virtuous’ moments where moral dilemmas play out. In the case of understanding future hardware like AR wearables, for example, that might mean exploring moments of dilemma and negotiation regarding a range of social phenomena, including privacy, presence, agency, equality of information, and representation. Such topics could be explored through observations of moments where these ‘virtues’ are negotiated, like the sharing and visibility of space in a home or between neighbors (i.e. privacy), or how friends, families, and colleagues delineate expectations for presence in the context of smartphone ubiquity.

While software platforms and algorithmic products, as are common on social media, may initially appear further from observable ‘virtuous’ acts – when precisely does the harm on social media happen, for example? – we nevertheless see the same approach as being relevant for the development of these products (in addition to content moderation, rules, and more). To surface the moral facts that govern many of the challenges faced by social media in the age of populism – like the spread of misinformation, the incitement of hatred or violence, and more – one might, for example, study real-life negotiations of facticity, free speech, or content curation, in addition to experiences of encountering the ‘other’ or escalating/de-escalating conflict. Following the Durkheimian thread, the moral facts guiding their future boundaries lie less in law and formal debate, and more in observations of how these experiences unfold in fora on- and offline: in confrontational sub-Reddits, live protests, and conspiratorial YouTube channels, but also in mixed office cultures, parent-teacher conferences, and content selection at home. Understanding how these moments of virtue are negotiated points to the underlying standards and mechanics at work.

Yet this only covers *what* to look for, not *how*. While a foundational understanding of the moral systems affected may be explored through traditional ethnographic techniques, the challenge of understanding how and in which ways these moral systems could change – and ensuring these are anchored to the capabilities of a future technology – are more uncertain. So-called ‘experimental philosophers’ have explored other routes to testing the ‘boundaries’ of certain virtuous actions. In attempting to resolve the famed trolley problem – wherein an out-of-control train can kill one group or another based on the train conductor’s choice of

track – experimental philosophy researchers (Copp 2012) asked a large sample to answer the problem in a range of different permutations, e.g. that one group had the conductor’s mother, that one group was older or overweight. By testing ‘real’ resolutions of the problem with a representative sample of the population, researchers were able to identify many of the contours and nuances which shape the resolution of the problem. While not ‘real’ in the sense of remaining a thought experiment – and more quantitative in nature than capturing the qualitative dilemmas – such experimental design points to additional ways to systematically explore the boundaries of moral intuitions, beyond how foundational dilemmas are experienced today.

In recent years, ReD has attempted to explore new ways to bring such experiments into real, lived experience. We have in recent ReD work explored designing ‘social experiments,’ combining approaches of experimental design from social psychology (see Isen and Levin 1972; Darley and Batson 1973) and the situated interactions of ethnomethodology, through methods like breaching (see Goffman 1971; Garfinkel [1967] 1991). This was tested most recently in a study on the related topic of the social acceptability of AR wearables. ReD researchers designed a ‘trivia night’ experiment to test the acceptability of the unequal distribution of information via AR glasses. Researchers used a live, planned trivia night as the setting, providing one pre-selected team with high-tech looking glasses and the answers to the night’s trivia questions. Through gestures and other artefacts, participants ‘simulated’ receiving information through the glasses as they consistently answered correctly. While no complaints were lodged by other, uninformed teams during the trivia rounds, when the final results were counted – and prizes were to be awarded – the other teams reacted with uproar. Such dynamics revealed both the tension with social norms of calling out abusers – pointing to moral intuitions that people were less likely to expose – as well as the importance of the ‘stakes’ in creating a context in which those intuitions were broken.

In order to balance a baseline understanding of these ‘virtuous phenomena,’ and how they could change, we would advocate a balance of both ethnographic research into foundational instances of these dilemmas and similar veins of experiments to contextualize them within the technology, and better tease out the nuances, boundaries, and ‘flexibility’ of such elements – notably through three additional variables of *User, Context & Monetization*.

## **Moral ‘Users’ of New Tech**

Studies of new technology often focus on lead users using comparable technologies (i.e. for future AR products, heavy users of smartphone-based AR or wearables) or expected early adopters (i.e. frequent gadget buyers). While such users no doubt provide valuable insight into new applications for technologies or adoption drivers, their status precisely as ‘lead’ does little to evidence the more general moral facts and ethical decision-making that will eventually drive the reception of these technologies. One need only consider the initial excitement about Google’s Glass from some groups to recognize that such a disconnect can be fatal. Instead, the target should be more representative ‘mainstream’ users – and, importantly, not only the ‘users’ themselves. Many moral quandaries – like the aforementioned ‘trolley problem’ – force a balancing of individual, group, and societal interests. We believe that a future ethics will need to understand how to balance the interests of purchasers, close and distant social groups, as well as unacquainted bystanders.

## Application Contexts

Coming in tandem with the need to explore the broad set of ‘users’ affected by the ethical decisions around a technology is an attention to the social and spatial distinctions which may also be inherent in deploying new technologies. As the trivia night example elucidates, bringing new technology into a space that is 1) shared, and 2) where information is viewed as ‘valuable,’ can drastically change the dynamics around what is and is not ethical. Were the same experiment repeated at home over a friendly game of Trivial Pursuit, the stakes might – although not always – be lower. Similarly, there are no shortage of examples – take watching pornography, for example – where behavior that is appropriate or acceptable changes widely from public spaces to the office to the home. Understanding, at a bare minimum, the difference between private (e.g. home alone), shared-private (e.g. friends’ homes), shared-public (e.g. offices), and public spaces (e.g. malls, parks) will likely be relevant for many technologies.

## Monetization

While perhaps seeming more focused in scope, recent attention to the monetization of personal data and renewed criticism of the exploitative practices of companies – and technology companies, in particular (see Zuboff 2015, 2019) – suggests a particular sensitivity to pricing, data monetization, and related business model questions as altering the ‘stakes’ around a certain issue. This also accompanies a shift from attention to the user as ‘purchaser’ of services to technology companies aiming to deliver a continued ‘experience’ or ‘relationship’ with the user (see Amirebrahimi 2016) as a driver of revenue – implying a growing relationship between even broader engagement with a technology and the notion of its monetization. And all of this rests on top of long-recognized moral issues surrounding the role of money and broader forms of exchange in societies (see Parry and Bloch 1989). As a result, different strategies of monetization have become intertwined in what counts as ethical action. A future ethics will need to understand how ‘financial stakes’ of a buying or even engaging with a product impacts its moral role in society.

Taken together, we see these variables as a framework for identifying the key objects of study necessary in defining a ‘grounded ethics’ for a given technology – as well as the broader toolbox of methods needed to discover these ethics in practice.

## CONCLUSION: TOWARDS A PRACTICE OF ‘GROUNDED ETHICS’

In this paper, we have aimed to address the much-discussed challenge of defining an ethics for developing and deploying new technologies and technology products globally – by shifting *where* such an ethics should come from. We have argued that the classical, *normative* approach to developing ethical frameworks – now guiding much of the approaches of major technology firms and related practitioners – does not sufficiently solve this problem, given that it leads to incoherent and inconsistent responses to the same problem, remains too open to interpretation in practice, and lacks the nuance necessary to guide practitioners as they make decisions. Rather, we have argued that a different epistemological approach – that of *discovery* – is needed in order to create a reliable system of ethics. Building of the growing field of the anthropology of ethics, we have located that discovery in the moral facts of

societies, but especially in the individual dilemmas and moral conflicts that elucidate the processes, systems, and practices by which ethics are developed – and what these systems suggest about the state of moral permissibility and its future flexibility & evolution. Finally, based on both theoretical and empirical examples, we have tried to synthesize the approach for a ‘grounded ethics’ into a framework to guide the research design of future explorations, notably: (a) the cultural foundations of the targeted societies, (b) the ‘virtuous’ phenomena likely affected by the new technology, (c) the ethical interests of different social groups, and who is the moral ‘user’ in each case, (d) contexts relevant to *where* ethics may be applied, e.g. physical sites, varied social groups and (e) moral notions around monetization. While this framework remains incompletely tested in full, in this paper’s role as a ‘catalyst’ for the EPIC community and wider practitioners, we envision a future vein for research and praxis to activate this framework in order to refine it and better explore how to integrate it into contemporary technology practice.

Let us stop for a moment to explore that last word – practice. While we have discussed the practical challenges of past ethical approaches from major technology companies, we have yet to discuss what ‘practice’ could look like in a ‘grounded ethics.’ This inevitably touches on the more often discussed question of ethics for anthropologists and other social scientists: that of our own role, practices, and positions relative to the people we study and represent to others. Given the challenges that Metcalf et al. (2019) raise surrounding the practice of ethics within technology firms, and Amirebrahimi’s (2016) concerns about the ‘flattening’ of ethnographic research, there are significant practical challenges to a grounded ethics, most notably: How does the ethnographers’ study of moral facts and ethical processes not become an ethical ‘rubber-stamp’ for technology products or projects? And how can the toolkit of ‘grounded ethics’ not become an over-simplification of complex moral negotiations?

These are, of course, complex questions worthy of extensive theoretical reflections, original research, and practical experimentation. As a starting point, we take inspiration from two of Laidlaw’s reflections on the practice of an anthropology of ethics:

Ethics, as self-formation, intrinsically includes a practice of inquiry, and presupposes ... an initial disjunction or difference between the self and one’s teacher or exemplar. (2014, 216)

...

[A]nthropological thought, in particular the exercise of the ethnographic imagination, can be a mode of reflective self-formation, a form of spiritual exercise, and since it necessarily involves not only ironic detachment tempering whatever degree of understanding ‘from the inside’ we are able to achieve, but also necessarily a certain suspension and detachment from one’s own knowledge and standpoint, it is an intrinsically sceptical one. (2014, 224)

In line with the framing of the negotiation of ethics that we had outlined in this paper, we see in both of these instances of ‘self-formation’ – the one, engaging with the ‘ethics’ of someone studied, the other, engaging with the ‘ethnographic imagination’ – a powerful foundation for engaging executives, technologists, and designers into the complexities of the moral and ethical negotiations they face. In that light, the ethnographer’s responsibility becomes to ensure that the recipient of a ‘grounded ethics’ also engages deeply with the experience of ‘self-formation’ that both the content and medium for communicating

intrinsically should enable. Put in its simplest terms, that boils down to a question of format: the production of a ‘grounded ethics’ should be communicated in ways that ensure deep engagement with the same moral negotiations that future users will face.

With ethnographic practitioners working for decades in technology, there are no doubt no shortage of immersive presentation, workshop, and communication formats that could support such a practice – many of which have surely been discussed at length within the EPIC community. For the needs of a ‘grounded ethics,’ as a starting point we would highlight one form of knowledge production which, we believe, will be well-suited to the function of ‘self-formation’ while engaging with ethics: *the trade-off*.

Rather than representing moral and ethical findings as static facts – thereby reducing them to binary guidance to be followed or rejected – describing moral and ethical systems as a set of trade-offs has the advantage of immersing decision-makers in the same balancing and weighing of virtues and moral costs which informants themselves are likely to face. In practice, this amounts to representing ‘grounded ethics’ in terms of the lived experience of:

- *Competing Virtues*: The virtues the members of society must balance in a given dilemma, how much they weigh or ‘pull’, and the underlying factors shaping their relevance and weight
- *Costs*: The moral ‘costs’ that individuals in a society would experience as a result of one decision or another – pointing also to the costs that a company would incur in the same decision
- *Processes*: The dominant logics, negotiation processes, and actors considered and/or involved in engaging with such decisions

In representing findings as a nexus of these virtues, costs, and processes, a ‘grounded ethics’ can thus represent not only a binary indication of ethical-unethical, but come closer to forcing users to engage more deeply with the dilemmas faced from the perspective of the people likely to be impacted in the future. That such a ‘grounded ethics’ is not a static snapshot thus also increases its applicability long-term: by including the underlying factors shaping how society engages with morality, the framework can be adjusted to account for evolutions in society. This is only more relevant when one considers the often circuitous and winding long-term path that guides the development of major technologies, in terms of business model, customer, and applications or use cases: the balancing of trade-offs can equally shift to match the changing realities within an innovation process. There is good reason to believe that, as such, a ‘grounded ethics’ can become an integrated tool in the innovation process – just as ethnography has become for problem or opportunity definition.

This represents only an initial foray into imagining how a ‘grounded ethics’ could look, and how it might, in practice, resolve some of the challenges faced by current approaches to ethics in the development of technologies. We challenge practitioners – from the closer world of ethnographers and design researchers, but also, from further afield, technologists and technology executives – to seek out and discover the ethics that will drive their future decisions. Considering the scale of societal change that technologies promise and technology companies aspire to, a more thoughtful route to defining the direction for that societal change remains *out there*, waiting.

**Ian Dull** is a Senior Manager at ReD Associates, and a lead in ReD's technology and mobility practices. He focuses on long-term industry changes driven by technological and societal shifts, and is interested in new tools to help companies decide what they *should* do, beyond economics. Ian holds an M.Phil in Archaeology from the University of Cambridge. [idu@redassociates.com](mailto:idu@redassociates.com)

**Fani Ntavelou Baum** is a Senior Consultant at ReD Associates, where since 2018, she has advised leaders in finance, FMCGs, medtech, pharmaceuticals and nonprofits on how to ground solutions for commercial problems in human insights. Fani holds an M.St. in Philosophy from the University of Oxford. [fnb@redassociates.com](mailto:fnb@redassociates.com)

**Thomas Hughes** is a Senior Consultant at ReD Associates. Thomas has helped world-leading Life Science companies improve their medical technologies by bringing them closer to the patient experience. Thomas is a medical anthropologist by training and holds a PhD in Anthropology from the University of Copenhagen. [thu@redassociates.com](mailto:thu@redassociates.com)

## NOTES

We would like to thank the many contributions of our colleagues at ReD Associates, whose methodological inspiration, empirical research, thoughtful reflections, and incisive questions created the foundation for this paper.

1. In a recent survey of UK tech workers, researchers found that 28% had seen decisions made about a technology that they believed would have a negative effect upon people or society. Among them, 20% went on to leave their companies as a result (Miller 2019).
2. Durkheim's focus on sanctions has given rise to a determinist view of ethics for some commentators, creating a norm-driven view of ethics. The argument goes that, if our ethical obligations are defined by our social duties, then there is no room for individual input and interpretation, the argument goes. Therefore, ethical facts are demoted to the status of norms *simpliciter*. James Laidlaw claims that ethics has been largely ignored by anthropologists with few exceptions due to the influence of Durkheim's deterministic vision of the moral fact (Laidlaw 2014). Laidlaw argues that until the early 2000s, most studies on morality and ethics in anthropology adopted more or less explicitly the so-called Durkheimian paradigm: ethnographic work consisted in the elucidation of a set of norms and values for a given group or society.
3. It is useful to understand social systems surrounding moral facts: they offer a stable reference point of acceptable norms and values, as well as an informed take of the context that has shaped these moral frameworks. This is indispensable knowledge for any understanding of moral facts.

## REFERENCES CITED

Amirebrahimi, Shaheen. 2016. "The Rise of the User and the Fall of People: Ethnographic Cooptation and a New Language of Globalization." In *2016 Ethnographic Praxis in Industry Conference Proceedings*, 71-103. <https://www.epicpeople.org>

- Anderson, Ken, Maria Bezaitis, Carl Disalvo & Susan Faulkner. 2019. "A.I. Among Us: Agency in a World of Cameras and Recognition Systems." In *2019 Ethnographic Praxis in Industry Conference Proceedings*, 38-64.
- Copp, David. 2012. "Experiments, Intuitions, and Methodology in Moral and Political Theory." In *Oxford Studies in Metaethics: Volume 7*, edited by Russ Shafer-Landau, 1-36. Oxford: Oxford University Press.
- Darley, John M., and C. Daniel Batson. 1973. "From Jerusalem to Jericho: A Study of Situational and Dispositional Variables in Helping Behavior." *Journal of Personality and Social Psychology* 27(1): 100–108.
- D'Ignazio, Catherine, and Lauren Klein. 2019. *Data Feminism*. MIT Open Press.
- Durkheim, Émile. [1924] 2010. *Sociology and Philosophy*. New York: Routledge.
- Durkheim, Émile. 1975. *Textes: II. Religion, Morale, Anomie*. Paris: Editions de Minuit.
- Eubanks, Virginia. 2018. *Automating Inequality: How High-Tech Tools Profile Police & Punish the Poor*. New York: St. Martin's Press.
- Fassin, Didier. 2006. "The End of Ethnography as Collateral Damage of Ethical Regulation?" *American Ethnologist* 33 (4): 522-524.
- Fassin, Didier. 2014. "The Ethical Turn in Anthropology: Promises and Uncertainties." *HAU: Journal of Ethnographic Theory* 4 (1): 429–435.
- Fiesler, Casey, Natalie Garret & Nathan Beard. 2020. "What Do We Teach When We Teach Ethics?: A Syllabi Analysis." In *The 51<sup>st</sup> ACM Technical Symposium on Computer Science Education (SIGSCE '20), March 11-14, 2020, Portland, OR, USA*. New York: ACM.
- Frenkel, Sheera, Davey Alba & Raymond Zhong. 2020. "Surge of Virus Misinformation Stumps Facebook and Twitter." *The New York Times*, 8 March 2020; updated 1 June 2020. <https://www.nytimes.com/2020/03/08/technology/coronavirus-misinformation-social-media.html>. Accessed 7 August 2020.
- Future of Life. 2015. "Autonomous Weapons: An Open Letter from AI & Robotics Researchers." *Future of Life Institute* website. 28 July 2015. <https://futureoflife.org/open-letter-autonomous-weapons/>. Accessed 7 August 2020.
- Garfinkel, Harold. [1967] 1991. *Studies in Ethnomethodology*. Cambridge: Polity Press.
- Goffman, Erving. 1971. *Relations in Public: Microstudies of the Public Order*. New York: Basic Books.
- Hefner, Robert W., ed. 1998. *Society and Morality in the New Asian Capitalisms*. New York: Routledge.
- Isaac, Mike, and Sydney Ember. 2016. "Live Footage of Shootings Forces Facebook to Confront New Role." *The New York Times*, 8 July 2016. <https://www.nytimes.com/2016/07/09/technology/facebook-dallas-live-video-breaking-news.html>. Accessed 6 August 2020.

- Isaac, Mike. 2020. "Facebook Moves to Limit Election Chaos in November." *The New York Times*, 3 September 2020, updated 22 September 2020. <https://www.nytimes.com/2020/09/03/technology/facebook-election-chaos-november.html>. Accessed 25 September 2020.
- Isen, Alice M., and Paula F. Levin. 1972. "Effect of Feeling Good on Helping: Cookies and Kindness." *Journal of Personality and Social Psychology* 21(3): 384–388.
- Jobin, Anna, Marcello Ienca & Effy Vayena. 2019. "The Global Landscape of AI Ethics Guidelines." *Nature Machine Intelligence* 1: 389-399.
- Laidlaw, James. 2002. "For an Anthropology of Ethics and Freedom." *Journal of the Royal Anthropological Institute* 8 (2): 311–32.
- Laidlaw, James. 2014. *The Subject of Virtue: An Anthropology of Ethics and Freedom*. Cambridge: Cambridge University Press.
- Laidlaw, James. 2017. "Ethics / Morality." In *The Cambridge Encyclopedia of Anthropology*, edited by F. Stein, S. Lazar, M. Candea, H. Diemberger, J. Robbins, A. Sanchez & R. Stasch. <https://www.anthroencyclopedia.com/entry/ethics-morality>. Accessed 26 September 2020.
- Lee, Nicol T., Paul Resnick & Genie Barton. 2019. "Algorithmic Bias Detection and Mitigation: Best Practices and Policies to Reduce Consumer Harms." Washington: The Brookings Institute. <https://www.brookings.edu/research/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/>. Accessed 6 August 2020.
- Metcalf, Jacob, Emanuel Moss & Danah Boyd. 2019. "Owning Ethics: Corporate Logics, Silicon Valley and the Institutionalization of Ethics." *Social Research: An International Quarterly* 82 (2): 449-476.
- Moss, Emanuel, and Metcalf, Jacob. 2019. "The Ethical Dilemma at the Heart of Big Tech Companies." *Harvard Business Review*, November 14, 2019. <https://hbr.org/2019/11/the-ethical-dilemma-at-the-heart-of-big-tech-companies>. Accessed September 17, 2020.
- Moss, Emanuel, and Schüür, Friederike. 2019. "Tutorial: Ethics in Data-Driven Industries." Tutorial conducted at EPIC2019 in Providence, Rhode Island.
- Parry, Jonathan, and Maurice Bloch, eds. [1989] 1996. *Money and the Morality of Exchange*. Cambridge: Cambridge University Press.
- Read, Kenneth E. 1955. "Morality and the concept of the person among the Gahuku-Gama." *Oceania* 25 (4): 233-282.
- Read, Max. 2017. "Does Even Mark Zuckerberg Know What Facebook Is?" *New York Magazine – Intelligencer*, 2 October 2017. <https://nymag.com/intelligencer/2017/10/does-even-mark-zuckerberg-know-what-facebook-is.html>. Accessed 25 September 2020.
- Sample, Ian. 2018. "Thousands of leading AI researchers sign pledge against killer robots." *The Guardian*, 18 July 2018. <https://www.theguardian.com/science/2018/jul/18/thousands-of-scientists-pledge-not-to-help-build-killer-ai-robots>. Accessed 7 August 2020.



- Simms, Ronald R. 1991. "The Institutionalization of Organizational Ethics." *Journal of Business Ethics* 10 (7): 493-506.
- Tamminen, Sakari and Holmgren, Elisabet. 2016. In *2016 Ethnographic Praxis in Industry Conference Proceedings*, 154–174. <https://www.epicpeople.org>
- VentureBeat. 2011. "Google's In-House Philosopher: Technologists Need a 'Moral Operating System.'" *VentureBeat*, 14 May 2011. <https://venturebeat.com/2011/05/14/damon-horowitz-moral-operating-system/>. Accessed 7 August 2020.
- Vincent, James. 2019. "The Problem with AI Ethics." *The Verge*, 3 April 2019. <https://www.theverge.com/2019/4/3/18293410/ai-artificial-intelligence-ethics-boards-charters-problem-big-tech>. Accessed 7 August 2020.
- Warzel, Charlie. 2020a. "What We Pretend to Know About the Coronavirus Could Kill Us." *The New York Times*, 3 April 2020. <https://www.nytimes.com/2020/04/03/opinion/sunday/coronavirus-fake-news.html>. Accessed 6 August 2020.
- Warzel, Charlite. 2020b. "Mark Zuckerberg is the Most Powerful Unelected Man in America." *The New York Times*, 3 September 2020. <https://www.nytimes.com/2020/09/03/opinion/facebook-zuckerberg-2020-election.html>. Accessed 25 September 2020.
- Wheeler, Melissa A., Melanie J. McGrath & Nick Haslam. 2019. "Twentieth century morality: The rise and fall of moral concepts from 1900 to 2007." *PLoS ONE* 14 (2), e0212267.
- Widlok, Thomas. 2004. "Sharing by Default? Outline of an Anthropology of Virtue." *Anthropological Theory* 4 (1): 53–70.
- Zuboff, Shoshana. 2015. "Big Other: Surveillance Capitalism and the Prospects of an Information Civilization." *Journal of Information Technology* 30: 75-89.
- Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: Public Affairs.