

Invisible Influence

Artificial Intelligence and the Ethics of Adaptive Choice Architectures

Daniel Susser
daniel.susser@psu.edu
Penn State University
State College, PA

ABSTRACT

For several years, scholars have (for good reason) been largely preoccupied with worries about the use of artificial intelligence and machine learning (AI/ML) tools to make decisions *about us*. Only recently has significant attention turned to a potentially more alarming problem: the use of AI/ML to influence *our decision-making*. The contexts in which we make decisions—what behavioral economists call our *choice architectures*—are increasingly technologically-laden. Which is to say: algorithms increasingly determine, in a wide variety of contexts, both the sets of options we choose from and the way those options are framed. Moreover, artificial intelligence and machine learning (AI/ML) makes it possible for those options and their framings—the choice architectures—to be tailored to the individual chooser. They are constructed based on information collected about our individual preferences, interests, aspirations, and vulnerabilities, with the goal of influencing our decisions. At the same time, because we are habituated to these technologies we pay them little notice. They are, as philosophers of technology put it, *transparent to us*—effectively invisible. I argue that this invisible layer of technological mediation, which structures and influences our decision-making, renders us deeply susceptible to manipulation. Absent a guarantee that these technologies are not being used to manipulate and exploit, individuals will have little reason to trust them.

CCS CONCEPTS

• Security and privacy → Social aspects of security and privacy.

KEYWORDS

data ethics; influence; manipulation; choice architecture; transparency; accountability

ACM Reference Format:

Daniel Susser. 2019. Invisible Influence: Artificial Intelligence and the Ethics of Adaptive Choice Architectures. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES '19)*, January 27–28, 2019, Honolulu, HI, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3306618.3314286>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

AIES '19, January 27–28, 2019, Honolulu, HI, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6324-2/19/01...\$15.00
<https://doi.org/10.1145/3306618.3314286>

1 INTRODUCTION

For several years, scholars have (for good reason) been largely preoccupied with worries about the use of artificial intelligence and machine learning (AI/ML) systems to make decisions *about us*. The algorithms mediating much of our lives are impenetrable black boxes [30] that perpetuate discrimination [2], evade accountability, and deprive us of due process [8, 9]. Institutional decision-making processes once guided by flawed and oftentimes biased human decision-makers are now being replaced by automated systems that learn from and reproduce those flaws and biases, and, indeed, accelerate their negative effects [28]. In response, these scholars have called for increased fairness, accountability, and transparency ("FAT*") around algorithmic decision-making, and have begun to develop technical and policy frameworks for achieving it.¹

This work, which takes what I call a *structural approach* to AI/ML ethics, is deeply important. But it only tells part of the story. In this paper, I take a more *individual approach*, describing a set of problems that flows from the introduction of AI/ML systems into people's everyday lives. Specifically, instead of attending to the ways algorithms are used to make decisions *about us*, I focus here on the use of AI/ML systems to influence *our decision-making*.² Information technologies increasingly mediate our decision-making processes in a wide range of contexts, from consumer decisions to political ones. And such technologies—driven, more and more, by AI/ML tools—are not neutral arbiters of such choices. Rather, these "adaptive choice architectures," as I call them, work to subtly guide us toward certain ends [10, 38, 44].

What's more, given the way we tend to interact with and experience technology, these influences are largely hidden. As philosophers of technology and science & technology studies (STS) scholars argue, once we become adept at using technologies they become "transparent" to us. Here, transparent is meant not in the sense in which it is often used in policy contexts (i.e., as insight into otherwise obscure practices), but rather more literally: once we are habituated to technologies we stop looking at them and instead look *through* them to the information and activities we use them to facilitate [15, 42, 43]. For the sake of clarity, I will refer to transparency understood in this way as "invisibility." The problems I describe in what follows stem from the fact that we do not consciously attend to—we don't *see*—many of the technologies now mediating our day-to-day lives. They are, for practical purposes, invisible.

¹See, for example, <http://www.fatml.org/resources/relevant-scholarship>

²Of course, many of the problems I discuss in what follows have structural causes. By "individual approach" I mean only that the object of my analysis in this paper is the effects of a certain class of AI/ML tools on individual experience.

This invisible layer of technological mediation, which structures and influences our decision-making, renders us radically susceptible to manipulation. At an individual level, it threatens our autonomy, and at a collective level, it stands to undermine social values, like democracy, which are premised on individuals being capable of independent decision-making. Given the risks these technologies pose, and the relative powerlessness of those affected to challenge them, I argue (1) that designers of AI/ML-driven systems sometimes ought to sacrifice the seamlessness and invisibility they usually strive for in the construction of user experiences, for seamless, visible interventions that promote user awareness and autonomy; (2) that the firms and other organizations designing and deploying adaptive choice architectures ought to provide *radical transparency* about both the means and the ends of the systems they create; and (3) the design and management of algorithmic influence ought to be governed by a strict duty of care. Absent a guarantee that these technologies are not being used to manipulate and exploit, individuals will have little reason to trust them.

2 ADAPTIVE CHOICE ARCHITECTURES

In their 2008 book *Nudge*, Richard Thaler (a behavioral economist) and Cass Sunstein (a legal theorist) put forward the now well-known idea that the contexts in which we make decisions inexorably shape the decisions we reach [39]. At bottom, the reason for this is that we do not, for the most part, give our decisions a great deal of conscious, reflective thought. Rather, as the psychologists Daniel Kahneman and Amos Tversky argue, we make most decisions very quickly, using fairly unreliable cognitive shortcuts or heuristics [17]. We lean on examples that come quickly to mind, even if they are not representative [39]. We assume that if our friends are doing something then perhaps we ought to do it too [39]. This means that we rely heavily on contextual cues when making everyday decisions, often without realizing it. The way in which our choices are situated and framed—the options we have and how they are described, who offers them to us, in what way, and when—all help determine the decisions we eventually make. Thaler and Sunstein call these decision-making contexts *choice architectures*.

Given this basic fact about human deciders, the question for Thaler and Sunstein is: why not design specific choice architectures to incline people toward individually and socially optimal outcomes? If contexts make such a big difference, why not design them with intention and for good? One of their famous examples of how to do this—how to “nudge” people into making better decisions—describes how food items are arranged in a cafeteria. People are more inclined, they suggest, to choose whichever options are situated at eye-level than they are to select something arranged above or below it [39]. Again, for Thaler and Sunstein, this mainly raises a question about policy—shouldn’t we arrange cafeteria food items in a way that inclines most people to select the healthiest option? Here, I want to raise a different question: what if the cafeteria were arranged differently for every person who walked in the door?

While that would be difficult to pull off in a physical cafeteria, in digital environments it is commonplace. The websites we visit

and apps we use are dynamic; unlike Thaler and Sunstein’s cafeteria, the options they present to us and the way those options are framed constantly change, producing different choice architectures for each user, and for the same user in different instances. To some extent, this has been true for web-based technologies since the late 1990s, when the notion of “adaptive user interfaces” first appeared (see, e.g., [18, 19]). What has changed in the intervening period is the amount of information available about each of us and the emergence of tools capable of analyzing it [10]. Today, as Karen Yeung argues, we are confronted with “highly personalised choice environment[s]”—digitally-mediated user experiences structured to match our particular cognitive idiosyncrasies [44]. As a result, we are increasingly subject to more than the gentle nudges Thaler and Sunstein recommend. Vast amounts of data about both individuals and groups, combined with digital platforms that can dynamically alter the way they present information to maximally impact each user means we are today subject to something much more powerful—what Yeung terms “hypernudges” [44]. Because our choice architectures are increasingly *adaptive*, our choice-making is increasingly susceptible to outside influence.

These developments arrive at the same time as our lives become generally more technologically-laden. Increasingly, we decide what to buy by browsing retail websites like Amazon. We decide where to eat by asking recommender services like Yelp. We turn to Monster.com to find a job and Tindr or OkCupid to find a date. Which is to say: the number of decisions we face in unmediated, purely physical settings is shrinking. Algorithms determine more and more, in a wide variety of contexts, both the sets of options we choose from and the way those options are framed. In and of itself, this is not a bad thing. The problem is that these adaptive choice architectures are largely hidden from view.

3 MEDIATING TECHNOLOGIES AS INVISIBLE INFRASTRUCTURE

We like to think of technology as something we *use*—devices or tools we engage with focus and attention. Our language reflects this: we talk about technology “users” and “end-users,” “user interfaces,” and “human-computer interaction.” As philosophers, psychologists, and media theorists have shown, however, this picture of how we relate to technology is partial at best (see [42]). In most cases, once we become sufficiently adept at using technologies we stop focusing on the technologies themselves and direct our attention instead to the things we are able to do *through* them. For example, we rarely think consciously about the material features of our smartphone screens or computer monitors; rather, we attend to the texts, images, videos, and other information they offer up to us. This is often referred to as *technological transparency*, pointing to the fact that we generally experience the world through technology, rather than experiencing technology itself, directly [15, 33, 37].

This notion of technological transparency should not be confused with the ways “transparency” is discussed in policy contexts. For instance, transparency is often used to describe *institutional openness*. Transparency as institutional openness means giving individuals insight into how otherwise obscure organizational processes work. It is transparency as institutional openness that people have in mind when they refer to the old quip that “sunlight is the

best disinfectant” [3]. There is also talk of transparency in discussions about algorithms. Algorithmic transparency is analogous to institutional transparency and means offering insight into how automated systems reach the decisions they do [30]. Elsewhere, Kiel Brennan-Marquez and I refer to this kind of transparency as “a view-under-the-hood” [4].

Philosophers of technology (especially phenomenologists and postphenomenologists) use the notion of transparency in a more literal way. They argue that once we become adept at using technologies, the technologies themselves recede from conscious attention and perception. This allows us to focus on what we can accomplish *using* the tools, rather than focusing on the tools themselves. Technologies become transparent in the sense that users see right through them. To take the most famous example, the early 20th Century German philosopher Martin Heidegger described the experience of using a hammer. When we hammer things, he pointed out, we do not attend to the hammer itself. Rather, we attend to the thing we are hammering (say, a nail). It is only when the technology breaks down—when it fails to drive the nail in—that we are prompted to examine the tool itself in order to diagnose the problem [13].

Don Norman brought this insight into the realm of technology design when he argued that digital technologies ought to be “invisible” to users. “Tools should be noticed only when they break,” he writes in *The Invisible Computer* [27]. For the sake of clarity, and following Norman’s convention, I will refer to this notion of technological transparency as “invisibility.” As Yoni Van Den Eede has shown, this idea—that when technologies work they form an invisible layer of mediation between human users and the world—has been raised time and again by technology scholars in a range of fields. In addition to Norman, media theorist Marshall McLuhan [24], sociologist Bruno Latour [21], social psychologist Sherry Turkle [40, 41], and philosophers like Don Ihde [15] and Peter-Paul Verbeek [43] have all pointed, in one way or another, to the fact that technologies become invisible to the people who use them [42].

As Norman argues, from the perspective of user experience this invisibility is all to the good. People use technologies to accomplish their goals. What is important to users is not the technologies themselves, but rather the activities they are using the technologies *for*. Well-functioning tools get out of the way and let us focus on the things we have taken up the tools to do. “Technology is our friend when it is inconspicuous, working smoothly and invisibly in the background, like a proper infrastructure should,” Norman writes, “Technology is a pest when it is in the way, when it is intrusive” [27]. Indeed, as Robert Rosenberger points out, digital technologies need not even break down or fail completely to irritate us—when websites simply load too slowly we experience a “drop in transparency,” forcing us out of the flow of our work and directing our attention to the computer or smartphone itself, and to the rest of the technological apparatus that might be responsible for the problem, like the modem, router, and so on [32].

The fact that we have incorporated information technologies into nearly all aspects of our lives means that for many people this invisible infrastructure—this layer of technological mediation that facilitates and supports all of the meaningful activities we engage in—is now pervasive. In 2016, the marketing research firm Nielsen Company estimated that the average American spends nearly 11

hours per day looking at screens (including smartphones, computers, televisions, etc.) [14]. Given that we also spend, on average, nearly 7 hours per day sleeping, this suggests that many people spend almost twice as many of their waking hours engaging with screens than not [16]. And not all digital technologies have screens. We also interface with a wide variety of “smart” objects—tools that have been outfitted with sensors and computer processors and connected to the internet. Recent (conservative) estimates suggest that there are already at least 6 billion such devices deployed around the world, and perhaps as many as 17 billion [26]. The so-called “Internet of Things” (IoT) (i.e., the totality of all those sensor-laden and internet-connected devices) promises to make the invisible infrastructure of information technologies we experience the world through even bigger, more pervasive, and—given that these devices are often designed precisely to operate in the background—less visible.³

Again, from the perspective of user experience, this pervasive, invisible layer of digital mediation promises to improve things. Norman is surely right that most of us care little about the devices we use, in and of themselves. What we care about are the people they connect us with, the work they help us carry out, the entertainment they provide, and the conveniences they offer. There are, however, issues beyond user experience. As we’ve seen, the technologies mediating our everyday lives shape our choice architectures. They enable “hypernudging”—dynamically altering the contexts in which we make decisions in order to influence their outcomes. The invisible infrastructure of digital mediation is a platform for manipulation.

4 ONLINE MANIPULATION

To manipulate someone is to impose a hidden influence on their decision-making process [38]. The fact that our lives are increasingly mediated by a hidden infrastructure of technologies designed to influence the way we choose, suggests, then, that we have rendered ourselves deeply susceptible to online manipulation. Tal Zarsky has warned about this for some time. Detailing the ways advertisers can use information collected about us to craft highly influential ads, he writes that “the use of personal information to provide tailored, manipulative content should be considered a detrimental outcome of today’s enhanced flow of personal information” [45]. Similarly, Ryan Calo describes a wide range of manipulative online practices that commercial firms use to induce customers to buy their products, and at the highest possible prices [5].

Consider the case of mobile health (or “mHealth”) apps—smartphone applications that work in tandem with fitness tracking devices, like Fitbits or Apple Watches, to continuously (or semi-continuously) gather information about users’ physical activity, physiological metrics (heart rate, etc.), and sleep patterns, and offer suggestions for how to lose weight, sleep better, or generally improve physical health and well-being. As Sax et al. point out, many (if not most) mHealth apps attempt to influence more than people’s health; they nudge users to buy health-related commodities and

³The technologies that comprise our invisible technological infrastructure are not all digital either. All of the physical and mechanical and analogue technologies we engage with follow the same general patterns. Though a relevant and important consequence of IoT technology is that the boundary between the digital and non-digital is beginning to blur. That being said, my focus in this paper is information technology.

services [36]. Using the highly specific information gathered about individual users from their fitness trackers, these apps create choice contexts designed to exert maximal influence. For example, when the Garmin Connect app detects that a user has taken fewer than average steps one day, it presents a graph showing the decrease alongside suggestions for “10 Healthy Work Habits,” one of which is a native advertisement (i.e., an advertisement disguised as non-sponsored content) for a \$40 sleep program. In cases such as this, Sax et al. write, “users do not know that they are being targeted with personalized content that intentionally obfuscates commercial intent by embedding it in health content” [36]. And the Garmin Connect app is only the beginning: as fitness trackers and the apps that come bundled with them become increasingly sophisticated, they will enable the construction of increasingly adaptive choice architectures, which deliver increasingly effective ads [20, 36].

It would be a mistake, though, to see the problem of invisible influence as limited to the commercial sphere. After all, it is not only consumer decision-making that is mediated by adaptive choice architectures. The so-called “gig economy” or “sharing economy” has introduced algorithmic mediation into the domain of labor relations [34]. Martin et al. explore the algorithmic tools Amazon’s Mechanical Turk system uses to structure the work of its crowd-sourced laborers [23]. And taking the ride-hailing app Uber as a case study, sociologists Alex Rosenblat and Luke Stark describe the myriad ways Uber’s app “algorithmically manages” (i.e., influences the decision-making of) its drivers—from customer ratings to data visualizations to notification nudges [35]. Drawing on this work, Calo and Rosenblat describe a particularly devious strategy:

while Uber originally showed drivers precise surge premiums in a given area in association with heat maps that display varying levels of surge through color schemes—yellow means demand is rising, orange means surge may appear soon, and red means it is surging—it changed the design of its app in October 2015 to show heat maps with those color schemes but without precise prices. In effect, the app encourages drivers to believe in surge and travel to receive surge rides, but it fails to provide a precise indicator or a guarantee of what that price is. Heat maps thus function as a behavioral engagement tool but can effectively operate as a bait-and-switch mechanism similar to the use of phantom cars to entice ride-hailers [6].

In other words, Uber habituated its drivers to one choice architecture, and then changed it in a way that leverages the old habits to influence drivers to behave as Uber wants.

Since 2016 the public has also become all too aware that our political life too is mediated by algorithms. Whereas worries about the effects of social media on the political system were once limited to questions about “filter bubbles” [29], “media manipulation”⁴ and voter microtargeting are now foremost concerns. Data analytics firm Cambridge Analytica reportedly used information collected about hundreds of thousands of Facebook users, combined with aggregate data purchased from third-party data collectors, to create “personality profiles” on voters in the United States, the United Kingdom, France, Germany, Kenya, and elsewhere around the world.

⁴See, for example: <https://datasociety.net/research/media-manipulation/>

According to a Cambridge Analytica whistleblower, those profiles allowed the company to deliver highly personalized political advertisements designed to influence voters by exploiting their “inner demons” [1]. To whatever extent these tactics were effective it was surely due in part to social media’s pervasive presence as a digital communications infrastructure running through the public sphere.

These examples of how adaptive choice architectures can invisibly influence decision-making across social contexts—from consumer contexts, to the workplace, to the political sphere—gives an indication of their ability to harm [38]. At an individual level, manipulation obviously threatens one’s autonomy. To act autonomously is to act independently, motivated by one’s own reasons. As Joseph Raz puts it, “the ideal of personal autonomy is that people should make their own lives” [31]. Invisible influence threatens this ideal by inducing people to act for reasons they don’t understand, and therefore can’t endorse. Just as clear, however, are the collective harms. Important social values, like democracy, are premised on the notion that the individuals taking part in the institutions they govern are themselves independent, autonomous, deciders. Diminishing the capacity of individuals to make meaningfully independent decisions in the political sphere undermines their ability to recognize political institutions as their own.

5 THE ETHICS OF INVISIBILITY

“If you could imagine any one obtaining this power of becoming invisible, and never doing any wrong or touching what was another’s, he would be thought by the lookers-on to be a most wretched idiot, although they would praise him to one another’s faces, and keep up appearances with one another from a fear that they too might suffer injustice.” (Glaucon to Socrates in Plato’s *Republic*, Book II)

So, what is to be done? In what remains, I briefly describe three normative implications of the preceding discussion.

First, the functionally invisible, adaptive choice architectures that mediate much of our experience render individuals radically vulnerable to the whims of others. More and more, as the Garmin, Uber, and Cambridge Analytica examples described in the previous section illustrate, the digital infrastructure that supports our everyday activities is designed to influence the directions those activities take—from using information about our sleep and exercise patterns to sell us weight loss plans to leveraging our subtle personality types to influence our votes. Such manipulative practices substitute the preferences and interests of the choice architects for the preferences and interests of choosers. Although they don’t coerce people into behaving the way advertisers, employers, or political campaigns would have them act, they nonetheless undermine people’s capacity to decide for themselves. The harms that flow from these hidden influences are twofold: at bottom, the fundamental harm, as discussed above, is that manipulation subverts autonomy—the individual’s capacity to act on the basis of reasons they themselves recognize and endorse. Furthermore, because autonomous individuals usually act in their own interest, the subversion of autonomy often leads to a second harm: diminished welfare. When our capacity to advocate for our own interests is compromised, we should not expect our interests to be served.

The second implication—correlative with the first—is that the creators of adaptive choice architectures wield an enormous amount of power, which carries with it special ethical responsibilities. In general, those who make decisions for others—e.g., parents, guardians ad litem, or fiduciaries—are understood as being duty-bound to make decisions in the interests of their charges. In fact, this is a special case of a broader principle pertaining to the relationship between vulnerability and responsibility. As Robert Goodin writes, “If A’s interests are vulnerable to B’s actions and choices, B has a special responsibility to protect A’s interests” [11]. Philosophically deep, this notion is also utterly commonsensical: when someone’s fate is in your hands, you ought to try to do right by them. T. H. Green states this idea even more directly: “There is no clearer ordinance of that supreme reason, often dark to us, which governs the course of man’s affairs, than that no body of men should [...] be able to strengthen itself at the cost of others’ weakness” [12].⁵ In the context of adaptive choice architectures, the meaning of Green’s observation is plain: it is wrong—and obviously so—to use the tools of invisible influence to induce people to act against their own interests.⁶ To do so is an abuse of power.

Third, despite the simple and obvious ethical imperatives just described, and given that the creators of adaptive choice architectures and the users who confront them often have competing interests, there is little reason to trust that these ethical responsibilities will be met and systems will be designed to the ultimate benefit of users, unless they are governed by stringent controls. This is true for several reasons. As the quote from Plato’s *Republic*, above, reminds us, the feeling of impunity carries with it a temptation that is difficult to resist. Moreover, the relationship between choice architect and chooser is not as straightforward as I have let on—modern algorithmic systems are rarely the product of one person’s or even a few people’s work. The apps and digital platforms described above are produced by huge teams of designers, project managers, programmers, quality assurance teams, and so on. As a result, there may be no single person who feels responsible for their effects. This is the so-called “problem of many hands” [25]. To be clear: the fact that there is no single choice architect—but rather sprawling teams of them—does not, morally, absolve those teams of responsibility. But as a practical matter it might mean that moral concerns are not always salient to the people building the tools.

So, again, the question is: what is to be done? I don’t pretend to have a complete answer. While the problem of manipulation is, one must assume, as old as we are, the problem of adaptive choice architectures—of an invisible infrastructure of technological mediation constantly collecting information about us and using that information to shape and influence our choices—is something new, and it will take a great deal more work to reckon fully with its consequences.

As a modest start, several things follow from my argument thus far. First, with respect to technology design, if the problem is, at bottom, hidden or invisible influence, then an obvious (if imperfect)

remedy is for the invisible infrastructure of technological mediation to reveal itself. Since, as I argued in section 3, invisibility (or “technological transparency”) is good for user experience, this means occasionally *compromising user experience in the name of user welfare and autonomy*. When creating adaptive systems, designers ought to consider how the value of a seamless user experience—the system itself receding from view—trades off against the user’s ability to understand how the system is influencing them. In some cases, seamlessness ought to be abandoned in favor of what Chalmers and Galani call “seamfulness” [7], or what I have called “strategic opacity” [37]—revealing to the user precisely what is going on beneath the surface.

But seamful user experiences will not be enough. Thus, second, in addition to requiring thoughtful, autonomy-promoting design decisions, adaptive choice architectures demand *radical transparency* on the part of the firms creating and deploying them. Here I mean transparency in the policy sense—we need a view under the hood. Those filling our everyday decision-making environments with adaptive choice architectures must be entirely forthcoming, not only about the *fact* that they are using such tools, but about their *purpose* (intended outcomes), and about the *mechanisms* by which they achieve it.

Finally, because the designers of these systems are highly incentivized to subvert user interests, and because the invisible nature of invisible influence renders it exceedingly difficult to detect, demands for radical transparency and seamful design must be accompanied by stringent accountability measures. Creators of adaptive systems ought to be held to the highest standard of care—they should be required, that is, to create systems that beyond any doubt advance users’ interests. Even then difficult questions would likely remain open, about the harms to autonomy that could accompany even the most beneficently manipulative systems. In tandem with radical transparency, though, and with user interface designs that promote user awareness and autonomy, it would represent a meaningful start.

6 CONCLUSION

My aim in this paper has been to shine light on a set of concerns stemming from advancements in artificial intelligence and machine learning—namely, concerns that adaptive choice architectures (using information collected about individuals to shape their decision-making contexts) facilitate online manipulation. In contrast with worries about how intelligent systems are used to make decisions *about us*, the issues I described stem from the use of AI/ML to influence *our decision-making*, raising fundamental questions about the integrity of user autonomy. To meet these ethical challenges, we must demand radical transparency from the creators of adaptive systems, encourage seamful user experiences that enable perceptive and deliberative human-computer interaction, and we must hold choice architects to the highest standard of care.

REFERENCES

- [1] 2018. Cambridge Analytica and Facebook: The Scandal so Far. *Al Jazeera News* (March 2018). <https://www.aljazeera.com/news/2018/03/cambridge-analytica-facebook-scandal-180327172353667.html>
- [2] Solon Barocas and Andrew D. Selbst. 2016. Big Data’s Disparate Impact. *California Law Review* 104 (2016), 671–732. <https://doi.org/10.2139/ssrn.2477899>

⁵Quoted in [11], p. 37.

⁶Arguably, it is also wrong to invisibly influence someone to act to *further* their own interests. Which is to say, due to its autonomy harms manipulation might always be wrong, even if it results in increased welfare. I leave that question to the side here. Goodin’s account of the relationship between vulnerability and responsibility understands vulnerabilities in entirely welfarist terms. For a discussion about these issues that understands vulnerability more broadly, see the introduction to [22].

- [3] Louis Brandeis. 1933. *Other Peoples Money and How the Bankers Use It*. Jacket Library.
- [4] Kiel Brennan-Marquez and Daniel Susser. 2016. Obstacles to Transparency in Privacy Engineering. In *2016 IEEE Security and Privacy Workshops (SPW)*. IEEE, San Jose, CA, 49–52. <https://doi.org/10.1109/SPW.2016.18>
- [5] M. Ryan Calo. 2014. Digital Market Manipulation. *The George Washington Law Review* 82, 4 (2014). <https://doi.org/10.2139/ssrn.2309703>
- [6] Ryan Calo and Alex Rosenblat. 2017. The Taking Economy: Uber, Information, and Power. *Columbia Law Review* 117 (2017). <https://doi.org/10.2139/ssrn.2929643>
- [7] Matthew Chalmers and Areti Galani. 2004. Seamful Interweaving: Heterogeneity in the Theory and Design of Interactive Systems. In *Proceedings of the 2004 conference on Designing interactive systems processes, practices, methods, and techniques - DIS '04*. ACM Press, Cambridge, MA, USA, 243. <https://doi.org/10.1145/1013115.1013149>
- [8] Danielle Keats Citron. 2008. Technological Due Process. *Washington University Law Review* 85, 6 (2008), 66.
- [9] Kate Crawford and Jason Schultz. 2014. Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms. *Boston College Law Review* 55, 1 (2014), 37.
- [10] Sara Degli Esposti. 2014. When big data meets dataveillance: the hidden side of analytics. *Surveillance & Society* 12, 2 (May 2014), 209–225. <https://doi.org/10.24908/ss.v12i2.5113>
- [11] Robert E. Goodin. 1985. *Protecting the Vulnerable: A Reanalysis of Our Social Responsibilities*. The University of Chicago Press, Chicago, IL.
- [12] T.H. Greene. 1881. Liberal Legislation and Freedom of Contract. In *Works 3: 3365–86*, R.L. Nettleship (Ed.). Longmans, London.
- [13] Martin Heidegger. 1962. *Being and Time*. Blackwell Publishing Ltd., Oxford.
- [14] Jacqueline Howard. 2016. Americans at More Than 10 Hours a Day on Screens. <https://www.cnn.com/2016/06/30/health/americans-screen-time-nielsen/index.html>
- [15] Don Ihde. 1990. *Technology and the Lifeworld: From Garden to Earth*. Indiana University Press, Bloomington.
- [16] Jeffrey M. Jones. 2013. In U.S., 40% Get Less Than Recommended Amount of Sleep. <https://news.gallup.com/poll/166553/less-recommended-amount-sleep.aspx>
- [17] Daniel Kahneman. 2013. *Thinking, Fast and Slow* (1st pbk. ed ed.). Farrar, Straus and Giroux, New York. OCLC: ocn834531418.
- [18] Pat Langley. 1997. Machine Learning for Adaptive User Interfaces. In *KI-97: Advances in Artificial Intelligence*, Jaime G. Carbonell, Jürg Siekmann, G. Goos, J. Hartmanis, J. Leeuwen, Gerhard Brewka, Christopher Habel, and Bernhard Nebel (Eds.). Vol. 1303. Springer Berlin Heidelberg, Berlin, Heidelberg, 53–62. https://doi.org/10.1007/3540634932_3
- [19] Pat Langley. 1999. User Modeling in Adaptive Interfaces. In *Proceedings of the Seventh International Conference on User Modeling*. Springer, 357–370.
- [20] Marjolein Lanzing. 2018. Strongly Recommended: Revisiting Decisional Privacy to Judge Hypernudging in Self-Tracking Technologies. *Philosophy & Technology* (June 2018). <https://doi.org/10.1007/s13347-018-0316-4>
- [21] Bruno Latour. 1992. Where Are the Missing Masses? The Sociology of a Few Mundane Artifacts. In *Shaping Technology/Building Society: Studies in Sociotechnical Change*, Wiebe E. Bijker and John Law (Eds.). MIT Press, Cambridge, MA.
- [22] Catriona Mackenzie, Wendy Rogers, and Susan Dodds (Eds.). 2014. *Vulnerability: new essays in ethics and feminist philosophy*. Oxford University Press, New York.
- [23] David Martin, Benjamin Hanrahan, Jacki O'Neill, and Neha Gupta. 2013. Being a Turker. In *CSCW '13*. ACM, 12.
- [24] Marshall McLuhan. 2003. *Understanding Media: The Extension of Man*. Ginkgo Press, Corte Madera.
- [25] Helen Nissenbaum. 1996. Accountability in a Computerized Society. *Science and Engineering Ethics* 2 (1996), 25–42. <https://link.springer-com.ezaccess.libraries.psu.edu/content/pdf/10.1007%2FBF02639315.pdf>
- [26] Amy Nordrum. 2016. Popular Internet of Things Forecast of 50 Billion Devices by 2020 Is Outdated. <https://spectrum.ieee.org/tech-talk/telecom/internet/popular-internet-of-things-forecast-of-50-billion-devices-by-2020-is-outdated>
- [27] Donald Norman. 1999. *The Invisible Computer: Why Good Products Can Fail, the Personal Computer Is So Complex, and Information Appliances Are the Solution*. MIT Press, Cambridge, MA.
- [28] Cathy O'Neil. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (first edition ed.). Crown, New York.
- [29] Eli Pariser. 2011. *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Press, New York. <http://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=1118322> OCLC: 759841356.
- [30] Frank Pasquale. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press, Cambridge.
- [31] Joseph Raz. 1986. *The Morality of Freedom* (reprinted ed.). Clarendon Press, Oxford. OCLC: 838707273.
- [32] Robert Rosenberger. 2009. The Sudden Experience of the Computer. *AI & Society* 24 (2009).
- [33] Robert Rosenberger and Peter-Paul Verbeek. 2015. A Field Guide to Postphenomenology. In *Postphenomenological Investigations: Essays on Human-Technology Relations*, Robert Rosenberger and Peter-Paul Verbeek (Eds.). Lexington Books, Lanham, Maryland.
- [34] Alex Rosenblat. 2018. *Uberland: How Algorithms are Rewriting the Rules of Work*. University of California Press, Oakland, CA.
- [35] Alex Rosenblat and Luke Stark. 2016. Algorithmic Labor and Information Asymmetries: A Case Study of Uber's Drivers. *International Journal of Communication* 10 (2016), 27.
- [36] Marijn Sax, Natali Helberger, and Nadine Bol. 2018. Health as a Means Towards Profitable Ends: mHealth Apps, User Autonomy, and Unfair Commercial Practices. *Journal of Consumer Policy* 41, 2 (June 2018), 103–134. <https://doi.org/10.1007/s10603-018-9374-3>
- [37] Daniel Susser. 2017. Transparent Media and the Development of Digital Habits. In *Postphenomenology and Media: Essays on Human-Media-World Relations*, Yoni Van den Eede, Stacy O'Neal Irwin, and Galit Wellner (Eds.). Lexington Books, New York, 27–44.
- [38] Daniel Susser, Beate Roessler, and Helen Nissenbaum. [n. d.]. Online Manipulation: Hidden Influences in a Digital World. *Under Review* ([n. d.]). <https://papers.ssrn.com/abstract=3306006>
- [39] Richard H. Thaler and Cass R. Sunstein. 2008. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press, New Haven. OCLC: ocn181517463.
- [40] Sherry Turkle. 1995. *Life on the Screen: Identity in the Age of the Internet*. Simon and Schuster, New York, NY.
- [41] Sherry Turkle. 2005. *The Second Self: Computers and the Human Spirit*. MIT Press, Cambridge, MA.
- [42] Yoni Van Den Eede. 2011. In Between Us: On the Transparency and Opacity of Technological Mediation. *Foundations of Science* 16, 2/3 (May 2011), 139–159. <https://doi.org/10.1007/s10699-010-9190-y>
- [43] Peter-Paul Verbeek. 2005. *What Things Do: Philosophical Reflections on Technology, Agency, and Design*. Pennsylvania State University Press, University Park, Pa.
- [44] Karen Yeung. 2017. Hypernudge: Big Data as a Mode of Regulation by Design. *Information, Communication & Society* 20, 1 (Jan. 2017), 118–136. <https://doi.org/10.1080/1369118X.2016.1186713>
- [45] Tal Zarsky. 2006. Online Privacy, Tailoring, and Persuasion. In *Privacy and Technologies of Identity: A Cross-Disciplinary Conversation*, Katherine J. Strandburg and Daniela Stan Raicu (Eds.). Springer-Verlag, New York, 209–224. https://doi.org/10.1007/0-387-28222-X_12