Exploitation in the Platform Age

Daniel Susser

In Beate Roessler and Valeries Steeves (ed.), *Being Human in the Digital World.* Cambridge University Press, forthcoming.

Being human in the digital age means confronting a range of disorienting normative challenges. Social problems, such as ubiquitous surveillance, algorithmic discrimination, and workplace automation feel at once familiar and wholly new. It is not immediately apparent whether the language and concepts we've traditionally used to describe and navigate ethical, political, and governance controversies, the distinctions we've drawn between acceptable and unacceptable relationships, practices, and exercises of power, or the intuitions we've relied on to weigh and balance difficult trade-offs adequately capture the difficult issues emerging technologies create. At some level of abstraction, there is nothing truly new under the sun. But for our language and concepts to be practically useful in the present moment we have to attend carefully to how they track—and what they illuminate about—the real-world challenges we face.

In this chapter I consider a common refrain among critics of digital platforms: big tech "exploits" us (Andrejevic 2012; Jordan 2015; Fuchs 2017; Cohen 2019; Zuboff 2019; Muldoon 2022). It gives voice to a shared sense that technology firms are somehow mistreating people—*taking advantage* of us, *extracting* from us—in a way that other data-driven harms, such as surveillance and algorithmic bias, fail to capture.

Take gig work, for example. Uber, Instacart, and other gig economy firms claim that their platforms strengthen worker autonomy by providing flexible schedules and greater control over when, where, and how people work. Yet many worry that gig economy—or what Ryan Calo and Alex Rosenblat call "taking economy"—platforms are, in fact, exploiting workers (Calo and Rosenblat 2017). Regulators warn that gig platforms set prices using "nontransparent algorithms," charge high fees, shift business risks onto workers, and require workers to pay for overhead expenses that companies normally cover (e.g., car insurance and maintenance costs),

allowing platforms to capture an unfair share of proceeds.[1] Workers are subjected to opaque, even deceptive, terms of employment, "algorithmic labour management" enables fine-grained, potentially manipulative control over work practices (U.S. FTC 2022; Rosenblat and Stark 2016; Susser et al. 2019), and high market concentration leaves workers with few alternative options (U.S. FTC 2022). Especially worrying, some forms of gig work—most notably "crowdwork," where work assignments are divided into micro-tasks and distributed online, which commonly drives content moderation and the labelling of training data for artificial intelligence (AI)—are reproducing familiar patterns of racial exploitation, with the global north extracting labor, digitally, from workers in the global south. Tech workers in Kenya have recently described these practices as "modern day slavery" and called on the US government to stop big tech firms from "systemically abusing and exploiting African workers."[2]

Now consider a very different example: the increasingly common practice of algorithmic pricing. Price adjustment is a central feature of market exchange—the primary mechanism through which markets (ideally) optimize economic activity. Sellers set prices in response to—amongst other things—overall economic conditions, competitor offerings, the cost of inputs, and buyers' willingness to pay. Today, many sellers rely on algorithms to do the work of price-setting, and these new pricing technologies have sparked a number of concerns. Economists worry, in general, that algorithmic pricing drives prices upward for consumers, in some cases by enabling new forms of collusion between firms, and in others simply as a result of feedback dynamics between multiple pricing algorithms (Mackay and Weinstein 2020). But these technologies don't simply automate price-setting, they can "personalize" it, tailoring prices to individual buyers (Acquisti et al. 2016). "Personalized" (or "customized") pricing, as industry firms euphemistically call it, is opaque—buyers rarely know when and how prices are personalized, making comparison shopping difficult. And the information used to set prices can include personal information about individual buyers (Seele et al. 2021), leading to concerns that

[1] According to the U.S. Federal Trade Commission (2022, 5): "[G]ig companies may use nontransparent algorithms to capture more revenue from customer payments for workers' services than customers or workers understand."

[2] "Open Letter to President Biden from Tech Workers in Kenya" (May 22, 2024), https://www.foxglove.org.uk/open-letter-to-president-biden-from-tech-workers-in-kenya/. For context, see Caroline Haskins, "The Low-Paid Humans Behind AI's Smarts Ask Bident to Free Them from 'Modern Day Slavery,'" in *Wired* (May 22, 2024), https://www.wired.com/story/low-paid-humans-ai-biden-modern-day-slavery/

algorithmic pricing helps firms "extract wealth" from consumers and "shift it to themselves" (Mackay and Weinstein 2020, 1).

One more case: "surveillance advertising." The contemporary digital economy is driven by targeted advertising.[3] Rather than charge consumers for the services they offer, such as search and social media, companies like Google and Facebook infuse their products with ads. Some argue that this business model is a win-win: users get access to valuable digital services for free, while technology firms earn huge profits monetizing users' attention.[4] But many have come to view the ad-based digital economy as a grave threat to privacy, autonomy, and democracy. Because targeted advertising relies on personal information—data about individual beliefs, desires, habits, and circumstances—to place ads in front of the people most likely receptive to them, digital platforms have become, effectively, instruments of mass surveillance (Tufekci 2018). And because targeted ads can influence people in ways they don't understand and endorse, they challenge important values like autonomy and democracy (Susser et al. 2019). Beyond these concerns, however, others argue that the surveillance economy involves an insidious form of *extraction.* Julie Cohen describes the market for personal information as the enclosure of a "biopolitical public domain," which "facilitates new and unprecedented surplus extraction strategies within which data flows extracted from people—and, by extension, people themselves—are commodity inputs, valuable only insofar as their choices and behaviours can be monetized" (Cohen 2019, 71).[5]

The goal of what follows is to unpack the claims that these platform-mediated practices are exploitative. What does exploitation entail, exactly, and how do platforms perpetrate it? Is exploitation in the platform economy a new *kind* of exploitation, or are these old problems

[3] As Tim Hwang (2020, 5) writes, "From the biggest technology giants to the smallest startups, advertising remains the critical economic engine underwriting many of the core services that we depend on every day. In 2017, advertising constituted 87 percent of Google's total revenue and 98 percent of Facebook's total revenue."

[4] For example, the Interactive Advertising Bureau (IAB), a trade association for the online marketing industry, argued in a recent comment in response to the U.S. Federal Trade Commission's Notice of Proposed Rulemaking on commercial surveillance: "there is substantial evidence that data-driven advertising actually benefits consumers in immense ways. As explained below, not only does data-driven advertising support a significant portion of the competitive U.S. economy and millions of American jobs, but data-driven advertising is also the linchpin that enables consumers to enjoy free and low-cost content, products, and services online" (2022, 10).

[5] Or, as Shoshana Zuboff (2019, 94) puts it, "the essence of the exploitation [typical of 'surveillance capitalism'] is the rendering of our lives as behavioural data for the sake of others'' improved control of us," the "self-authorized extraction of human experience for others' profit" (Zuboff 2019, 19).

dressed up as new ones? What would a theory of digital exploitation add to our understanding of the platform age? First, I define exploitation and argue that critics are justified in describing many platform practices as wrongfully exploitative. Next, I focus on platforms themselves—both as businesses and technologies—in order to understand what is and isn't new about the kinds of exploitation we are witnessing. In some cases, digital platforms perpetuate familiar forms of exploitation by extending the ability of exploiters to reach and control exploitees. In other cases, they enable new exploitative arrangements by creating or exposing vulnerabilities that powerful actors couldn't previously leverage. On the whole, I argue, the language of exploitation helps express forms of injustice overlooked or only partially captured by dominant concerns about, e.g., surveillance, discrimination, and related platform abuses, and it provides valuable conceptual and normative resources for challenging efforts by platforms to obscure or legitimate them.

I. Defining Exploitation

What exploitation is and what makes it wrong have been the subject of significant philosophical debate. In its modern usage, the term has a Marxist vintage: the engine and the injustice of capitalism, Marx argued, is the exploitation of workers by the capitalist class. For Marx, labour is unique in its ability to generate value; lacking ownership and control over the means of production, workers are coerced to give over to their bosses most of the value they create. This, in Marx's view, is the sense in which workers are exploited: value they produce is taken, *extracted* from them, and claimed, unjustly, by others.[6]

Some media studies and communications scholars have adopted this Marxian framework and applied it in the digital context, arguing that online activity can be understood as a form of labour and platform exploitation as appropriation of the value such labour creates.[7] For example, pioneering work by Dallas Smythe on the "audience commodity"—the packaging and selling of consumer attention by advertisers—which focused primarily on radio and television, has been extended by theorists such as Christian Fuchs and Mark Andrejevic to understand the internet's

[6] For a more complex picture of the relationship between exploitation and capitalist appropriation, especially focusing on its racialized character, see Nancy Fraser (2016) in *Expropriation and Exploitation in Racialized Capitalism: A Reply to Michael Dawson.*

[7] See, e.g., Tiziana Terranova, "Free Labour: Producing Culture for the Digital Economy," *Social Text* 18(2), 2000.

political economy through a constellation of Marxist concepts, including exploitation, commodification, and alienation.[8] As Andrejevic argues, this work adds a crucial element to critical theories of the digital economy, missing from approaches focused entirely on data collection and privacy (2012, 73).

While these accounts offer important insights, I depart from them somewhat in conceptualizing platform exploitation, for several reasons. Many—including many Marxist theorists—dispute the details of Marx's account. Specifically, critics have demonstrated that the "labour theory of value" (the idea that value is generated exclusively by labour, that it is more or less homogeneous, and that it can be measured in "socially necessary labour time"), upon which Marx builds his notion of exploitation, is implausible (Cohen 1979; Wertheimer 1996, x). So, the particulars of the orthodox Marxist story about exploitation are probably wrong, and building a theory of digital exploitation on top of it would mean placing that theory on a questionable foundation. Still, the normative intuition motivating the theory—that workers are often subject to unjust extraction, that something of theirs is taken, wrongfully, to benefit others—is widely shared, and efforts have been made to put that intuition on firmer theoretical ground (Cohen 1979; Roemer 1985; Reiman 1987).

Moreover, the concept of exploitation is more capacious than the Marxist account suggests. Beyond concerns about capitalist exploitation, we might find and worry about exploitation more broadly, in some cases outside of economic life altogether (Goodin 1987). Feminist theorists, for example, have identified exploitation in sexual and marital relationships (Sample 2003), bringing a wider range of potential harms into view. And while the exploitation of workers—central to Marxist accounts—continues to be vitally important, as we will see, the incorporation of digital platforms into virtually all aspects of our lives opens the door to forms of exploitation Marxist accounts underemphasize or ignore.

Contemporary theorists define exploitation as *taking advantage of someone*—using them to benefit oneself. Paradigm cases motivating the philosophical literature include worries about sweatshop labour, commercial surrogacy, and sexual exploitation.[9] Of course, taking advantage

[8] For example, see Fuchs (2010) in *Labour in informational capitalism and on the Internet*. For a helpful intellectual history of related work on the political economy of media and communication technology, see Lee McGuigan (2014).

[9] On sweatshop labour, see e.g., Jeremy Snyder (2010) in *Exploitation and Sweatshop Labour: Perspectives and Issues*; and Matt Zwolinski (2012) in *Structural Exploitation*. On commercial surrogacy, see e.g., Wertheimer (1996). On sexual exploitation, see Sample (2003).

is not always wrong—one can innocently take advantage of opportunities or rightly take advantage of an opponent's misstep in a game. Much of the debate in exploitation theory has thus centered on its "wrong-making features," i.e., what makes taking advantage of someone morally unacceptable. There are two main proposals: one explains wrongful exploitation in terms of unfairness, the other in terms of disrespect or degradation.

A. Exploitation as Unfairness

Taking advantage of someone can be unfair either for procedural or substantive reasons. An interaction or exchange is procedurally unfair if the process is defective—for example, if one party deceives the other about the terms of their agreement or manipulates them into accepting disadvantageous terms. Substantive unfairness, by contrast, is a feature of outcomes. Even if the process of reaching an agreement is defect-free, the terms agreed to might be unacceptable in and of themselves. Consider sweatshop labour: a factory owner could be entirely forthright about wages, working conditions, and the difficult nature of the job, and likewise workers could reflect on, understand, and—given few alternative options—decide to accept them. The process is above-board, yet in many cases of sweatshop labour the terms themselves strike people as obviously unfair.

One way to understand what has gone wrong here is via the notion of "social surplus."[10] Often when people interact or exchange the outcome is positive-sum: cooperation can leave everyone better off than they started. In economics, the surplus created through exchange is divided (sometimes equally, sometimes unequally) between sellers and buyers. But the concept of a social surplus need not be expressed exclusively in monetary terms. The idea is simply that when people interact they often increase total welfare. If I spend my Saturday helping a friend move, he benefits from (and I lose) the labour I've provided for free. But we both enjoy each other's company, feel secure in knowing we're deepening our relationship, and I derive satisfaction from doing someone a favor.

[10] For an overview of competing accounts, see Matt Zwolinski Benjamin Ferguson, and Alan Wertheimer (2022) in "Exploitation," in *The Stanford Encyclopedia of Philosophy,* https://plato.stanford.edu/entries/exploitation

Exploitation enters the picture when the social surplus is divided unfairly.[11] Returning to the sweatshop case, for example, the exchange is unfair—despite the absence of procedural issues—because the factory owner claims more than his fair share of the value created. He could afford to pay the factory workers more (by collecting a smaller profit) but chooses not to.[12] Likewise, we sometimes use the language of exploitation to describe similar dynamics within personal relationships: if one friend always relies on another for help but rarely reciprocates, we say that the first is exploiting the second.

B. Exploitation as Degradation

Not all exploitation, however, can be explained in terms of unfairness. Take price gouging, another standard example of exploitation: imagine, say, a thirsty hiker, lost in the desert, encounters a fellow traveler who offers to part with their extra bottle of water for $1000.[13] The seller is perfectly forthright about their product, its condition, and the terms of sale, and the buyer reflects on, understands, and decides to accept them. In other words, there is no procedural unfairness involved. Moreover, if buying the water will save the hiker's life, he is—in one sense—getting a pretty good deal. Most people value their life at a lot more than $1000. Indeed, as Zwolinski points out, in such cases there is reason to believe that the hiker is getting far more of the surplus created through the exchange than the greedy seller (the former gets his life, the latter $1000). So substantive unfairness—unevenly distributing the social surplus—can't explain the problem here either.

For some theorists, cases like this demonstrate another possible wrong-making feature of exploitation: degradation, or the failure to treat people with dignity and respect. Allen Wood (1995) argues that using another person's vulnerability to one's own advantage is instrumentalizing, demeaning. "Proper respect for others is violated when we treat their

[11] Determining what counts as an unfair division of the social surplus is, unsurprisingly, a matter of some controversy. Hillel Steiner (1984) in *A Liberal Theory of Exploitation* argues that the distribution is unfair when it's the product of historical injustice, while for John Roemer (1985), the unfairness derives from background conditions of inequality. On Alan Wertheimer's (1996) account, the distribution is unfair when one party pays more than a hypothetical "fair market price".

[12] This is another way of framing the normative intuition that motivates Marxist accounts of exploitation: the capitalist class claims an unfair share of the surplus created by the working class. See Roemer (1985) and Reiman (1987).

[13] This example is borrowed from Zwolinski (2022).

vulnerabilities as opportunities to advance our own interests or projects. It is degrading to have your weaknesses taken advantage of, and dishonorable to use the weaknesses of others for your ends" (Wood 1995, 150-51). Indeed, for Wood (1995, 154), even in cases like sweatshops, which—as we've just seen—can plausibly be explained in terms of unfairness, this kind of degradation is the deeper, underlying evil.

Some argue that exploitation is wrong solely in virtue of one or another of these moral considerations—at bottom, it is either unfair *or* degrading—and such theorists have worked to show that certain cases intuitively cast in one moral frame can be explained equally well or better through another. For present purposes, I follow theorists who adopt a more pluralistic approach and define wrongful exploitation as Matt Zwolinski (2012) does: *taking advantage of someone in an unfair or degrading way.*[14] In some cases, exploitation is wrong because it involves unfairness, in other cases because it involves degradation. Oftentimes more than one wrong-making feature is at play, and digital platforms potentially raise all of these concerns.

II. Platform Exploitation?

A first question, then, is whether the kinds of practices I described at the start reflect the normative problems discussed above. Are platforms exploiting people?

If exploitation is taking advantage of someone in an unfair or degrading way, and what enables exploitation—what induces someone to accept unfair terms of exchange, or what makes taking advantage of such terms degrading—is the exploitee's vulnerability (the fact that they lack decent alternatives), then identifying exploitation is partly an empirical exercise. It requires asking, on a case-by-cases basis: Are people vulnerable? What are their options? Are platforms taking advantage of them?

However, that need not prevent us from generalizing a little. Returning to the alleged abuses by gig economy companies, we can now recast them in this frame. Recall the FTC's concern that gig platforms set prices using "nontransparent algorithms." Reporting on ethnographic work in California's gig-based ride hail industry, legal scholar Veena Dubal describes drivers struggling to understand how the prices they're paid for individual rides are set, why different drivers are paid different rates for similar rides, or how to increase their earnings.

[14] For an overview and argument in favor of a pluralist approach, see Snyder (2010).

Not only because the algorithms powering ride-hail apps are opaque, but because they set prices dynamically: "You've got it figured out, and then it all changes," one driver recounts (Dubal 2023, 1964).[15] Using the language developed above, we can describe this opacity and dynamism as sources of procedural unfairness—whether the terms of exchange reached are fair or not, the process of reaching them is one in which drivers are disempowered relative to the gig platforms they are "negotiating" with.[16]

There is also reason to worry that the terms reached are often substantively unfair, with platforms siphoning off more than their fair share of profits—an unfair distribution of the social surplus. Beyond concerns about how gig apps set prices, or about the ability of drivers to understand and exert agency in the process, the FTC complaint points out that ride hail apps charge drivers high fees, shift risks of doing business—usually absorbed by firms—onto drivers, and require them to pay for overhead expenses that companies normally cover, such as car insurance and maintenance costs. Similarly, crowdworkers in the global content moderation industry describe doing essential but "mentally and emotionally draining work" for little pay and without access to adequate mental health support: "Our work involves watching murder and beheadings, child abuse and rape, pornography and bestiality, often for more than 8 hours a day. Many of us do this work for less than $2 per hour."[17]

While charges of exploitation may be unwarranted in cases where, for example, ride hail drivers really are just driving for a little bit of extra cash on the side, in the mine run of cases, where gig workers lack other job options and depend on the income earned through gig app work, the charges seem fitting. Moreover, there is reason to believe that gig companies like Uber actively work to create the very vulnerabilities they exploit, by using venture capital funding to underprice competition, pushing incumbents out of the market and consolidating their own position. One reason ride hail drivers often lack alternative options is Uber has put them out of business.

Algorithmic pricing in consumer contexts also raises procedural and substantive fairness concerns. Like ride hail drivers navigating opaque, dynamic fare setting systems, consumers are

[15] Veena Dubal (2023) in *On Algorithmic Wage Discrimination*. See also, Zephyr Teachout (2023) in *Algorithmic Personalized Wages.*

[16] Even describing the process as a negotiation is perhaps too generous—drivers simply have the option of accepting a ride and the designated fare or not.

[17] "Open Letter to President Biden from Tech Workers in Kenya" (May 22, 2024), https://www.foxglove.org.uk/open-letter-to-president-biden-from-tech-workers-in-kenya/

increasingly presented with inconsistent prices for the same goods and services, making it difficult to understand why one is offered a particular price or how it compares to the prices others are offered (Seele et al. 2021). And because the algorithms determining prices are inscrutable (as in the gig app case), there is an informational asymmetry between buyers and sellers that puts the former at a significant disadvantage, potentially creating procedural fairness problems. How can a buyer decide if prices are competitive without knowing (at least roughly) how they compare to prices others in the marketplace are paying, and how can they comparison shop when prices fluctuate unpredictably?[18]

Personalized pricing makes things even worse. In addition to issues stemming from algorithmic opacity and dynamism, price personalization—or what economists call "first-degree" or "perfect" price discrimination (i.e., the tailoring of prices to specific attributes of individual buyers)—raises the specter that sellers are preying on buyer vulnerabilities. On one hand, as Jeffrey Moriarty (2021, 497) argues, price discrimination is commonplace and generally considered acceptable.[19] Even highly personalized pricing might be unproblematic, provided buyers know about it and have the option to shop elsewhere.[20] From an economics perspective, first-degree price discrimination has traditionally been viewed as bad for consumers but good for overall market efficiency. If buyers pay exactly as much as they are hypothetically willing to (their "reservation price")—and not a cent less—then sellers capture all of the surplus but also eliminate deadweight loss (Bar-Gill 2019).

Algorithmically personalized pricing changes things. First, as we have seen, it is often opaque and inscrutable—buyers do not know that they are being offered individualized prices, or if they do, how those prices are determined. Thus, even if they could shop elsewhere, they might not know that they should. Second, the above arguments assume that personalized pricing simply attempts to find and target the buyer's reservation price. But Oren Bar-Gill (2019) points out that the conception of "willingness to pay" underlying these traditional arguments, which imagines

[18] For a related discussion, see Ariel Ezrachi and Maurice Stucke, "The Rise of Behavioural Discrimination." *European Competition Law Review* 37(12), 2016.

[19] Indeed, offering different people different prices may, on balance, benefit the worst off. To use a well-known example, if pharmaceutical companies couldn't charge different prices to consumers in rich and poor countries, they would have to charge everyone (including those with the fewest resources) higher prices in order to recoup costs. See Jeffrey Moriarty (2021) in *Why Online Personalized Pricing is Unfair*.

[20] Moriarty explicitly argues that under these conditions price personalization is non-exploitative (2021, p. 498). Etye Steinberg (2020) disagrees, arguing that data-driven personalized pricing is unfair on account of concerns about relational equality.

the reservation price simply as a function of consumer preferences and budgets, misses an important input: how buyers perceive prices and a product or service's utility.

People are often mistaken about one or both, misjudging, for example, how much something will cost overall, how often they will use it, the value they will ultimately derive from it, and so on (one can think here of the cliché about gym memberships purchased on January 1). Personalized pricing algorithms can provoke and capitalize on these errors, encouraging people to over-value goods (increasing willingness to pay) and under-predict total cost—i.e., it can *change* their reservation price (Calo 2014). In such cases, Bar-Gill (2019, 221) argues, the traditional economics story is wrong—first-degree price discrimination harms consumers *and* diminishes overall efficiency, as "cost of production exceeds the actual benefit (but not the higher, perceived benefit)." The only benefit is to sellers, who capture the full surplus (and then some), raising substantive fairness concerns. Thus, the exploitation charge seems plausible in this case too. Though again, much depends on the details. If buyers know prices are being personalized, and they are able to comparison shop, it is less obvious that sellers are taking advantage of them.

Finally, behavioral advertising. Are data collectors and digital advertisers taking advantage of us? In the U.S., commercial data collection is virtually unconstrained, and data subjects have little choice in the matter. Companies are required only to present boilerplate terms of service agreements, indicating what data they will collect and how they plan to use it. Data subjects usually have only two options: accept the terms or forego the service. As many have argued, this rarely amounts to a real choice.[21] If, for example, one is required to use Microsoft Office or Google Docs as part of their job, are they meaningfully free to refuse the surveillance that comes with it? Put another way, many people are in a real sense dependent on digital technologies—for their jobs, at school, in their social lives—and surveillance advertisers, unfairly, take advantage of that dependency for their own gain.

Having said that, it is worth asking further questions about how those gains are distributed—who benefits from this system? Much of the value derived from surveillance advertising obviously flows directly into the industry's own coffers: revenue from online advertising accounts for the vast majority of profits at Google and Facebook, the two largest

[21] For an overview, see Susser (2019) in *Notice After Notice-and-Consent: Why Privacy Disclosures are Valuable Even if Consent Frameworks Aren't.*

industry players (Hwang 2020). But where does the surplus come *from*? On one view, elaborated most dramatically by Shoshana Zuboff, the surplus comes from us. It is a "behavioural surplus"—information about our individual desires, habits, and hang-ups, used to steer us toward buying stuff (Zuboff 2019). According to this argument, personal information and the predictions they make possible are merely conduits, carrying money from regular people's pockets into the hands of companies running ads (with the surveillance industry taking a cut along the way). In other words, data subjects are being exploited for the benefit of advertisers and sellers.

There is another view, however, according to which this whole system is a sham. Tim Hwang and others argue that behavioral advertising simply doesn't work—the predictions sold to sellers are largely wrong and the ads they direct rarely get us to buy anything (Hwang 2020).[22] But as Hwang points out, that does not mean people do not benefit from online advertising. *We* benefit from it, enjoying for free all of the services digital ads underwrite, which we would have to pay for if the ads went away. On this view, personal data is a conduit carrying money from the advertising budgets of sellers into the hands of app makers and producers of online content (with, again, the surveillance industry collecting its cut along the way). In other words, *the companies running ads are being exploited* for our benefit.

III. What's Old is New Again

To this point, I have discussed platforms in general terms, focusing on what they do and whether we ought to accept it rather than on what they are and how they are able to treat people this way. I turn now to the latter: what platforms are, how they are able to engage in these different forms of exploitation, and what role digital technology specifically is playing in all of this.

The term "platform" is used in multiple registers. In some contexts, it is used to describe a set of companies—e.g., Amazon, ByteDance, Meta, or Google. In other contexts, the term is used to describe the heterogeneous set of digital technologies such companies build, deploy, and use to generate revenues—e.g., Amazon's marketplace, the TikTok or Instagram apps, or Google's digital advertising service. This ambiguity or multiplicity of meaning is neither a mistake nor an accident; platforms are both of these things simultaneously, businesses and technologies, and they must be understood both in economic and sociotechnical terms.

[22] For a more careful investigation into this question and its implications, see Daniel Susser and Vincent Grimaldi (2021) in *Measuring Automated Influence: Between Empirical Evidence and Ethical Values*.

Unlike ordinary service providers, platforms function primarily as social and technical infrastructure for interactions between other parties. TikTok, Instagram, and social media platforms more broadly find audiences for content creators and advertisers who will pay to reach them. Gig economy platforms, like Uber and Lyft, facilitate exchanges between workers and people in need of their labour. As Tarleton Gillespie (2010, 4) points out, the term "platform" misleadingly brings to mind a sense of neutrality: "platforms are typically flat, featureless, and open to all." In fact, digital platforms work tirelessly to shape the interactions they host and to influence the people involved. As we've seen, they do this by carefully designing technical affordances (such as opaque and personalized pricing algorithms) and by pressing economic advantages (when, for example, they leverage venture capital to underprice incumbents and eliminate competition).

So: platforms mediate and structure relationships. Some of these relationships have long existed and have often been sites of exploitation; when platforms enter the picture they perpetuate and profit from them. Other relationships are new—innovations in exploitation particular to the platform age.

A. Perpetuating Exploitation

Many platforms profit by creating new opportunities for old forms of exploitation. Platform-mediated work is a case in point: while not all employers exploit their employees, the labor/management relationship is frequently a place where worries about exploitation arise, and digital platforms breathe new life into these old concerns.

Indeed, platforms can increase the capacity of exploiters to take advantage of exploitees by enabling exploitation at scale, expanding the reach of exploitative firms and growing the pool of potential exploitees (Pfotenhauer et al. 2022).[23] Gig app firms, based in Silicon Valley and operated by a relatively small number of engineers, managers, and executives, profit from workers spread throughout the world—in 2022, for example, Uber had 5 million active drivers worldwide (Biron 2022). Moreover, as we've seen, these dynamics are visible in the broader

[23] Pfotenhauer et al. (2022) describe the inexorable march toward massive scale as "the uberization of everything," which introduces, they argue, "new patterns of exploitation."

phenomenon of "crowdwork," or what Dubal (2020) terms "digital piecework."[24] Platforms like Amazon Mechanical Turk (AMT) carve work (such as social media content moderation and labelling AI training data) into small, discrete, distributable chunks, which can be pushed out to workers sitting in their homes or in computer centers, new sites of so-called "digital sweatshops" (Zittrain 2009). As sociologist Tressie McMillan Cottom (2020) argues, these practices constitute a kind of "predatory inclusion"—one of many ways digital platforms have implicated themselves in broader patterns of racial capitalism.

At a more granular level, digital platforms also facilitate worker exploitation by reconfiguring work, work conditions, and wage determination. A growing body of scholarship explores the nature and functioning of "algorithmic labour management": the use of digital platforms to control workers and organize work. In contrast with simplistic narratives about automation displacing workers, this research brings to light the myriad ways digital technologies are becoming insinuated in human labour, changing its character, shifting risks, and creating new pathways for discrimination and extraction. Pegah Moradi and Karen Levy (2020) argue, for example, that automation and platform intermediation often increase profits for firms not by producing new efficiencies, but rather by shifting the costs of inefficiencies onto workers. "Just-in-time" scheduling algorithms make it possible to employ workers at narrower intervals dynamically tailored to demand, reducing labour costs by rendering jobs more precarious and less financially dependable for workers (Moradi and Levy 2020). And algorithmic management lets employers "narrowly define work to include only very specific tasks and then pay workers for those tasks exclusively" (Moradi and Levy 2020, 281). Ride-hail drivers, for instance, are compensated only for active rides, not for the time they spend searching for new passengers.

From a law and policy perspective, platforms also make it easier to exploit workers through legal arbitrage. By creating the appearance of new forms of work, gig economy apps render workers illegible to the law, and in so doing, they allow firms to ignore worker rights and circumvent existing worker protections. For example, high profile political battles have recently been waged over whether gig workers should be legally classified as independent contractors or

[24] Others describe this as "ghost work." See Mary L. Gray and Siddharth Suri (2019) in *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*; and Veena Dubal (2020) in *The Time Politics of Home-Based Digital Piecework.*

as employees of gig economy companies.[25] Gig economy firms contend that all their platforms do is connect workers to paying customers; the workers don't work *for them*, but rather for app users. Gig workers and their advocates argue that firms carefully manage and directly profit from their labour, and as such they ought to be given the same rights, benefits, and protections other workers enjoy. As Dubal writes about app-based Amazon delivery drivers, "In this putative non-employment arrangement, Amazon does not provide workers' compensation, unemployment insurance, health insurance, or the protected right to organize. Nor does it guarantee individual DSPs [Delivery Service Providers] minimum wage or overtime compensation" (Dubal forthcoming, 4).

B. Innovations in Exploitation

Different dynamics are at work in cases like algorithmic pricing. Here, the relationship mediated by digital platforms—in the pricing case, the relationship between buyers and sellers—is not normally a site of exploitation.[26] The introduction of digital platforms transforms the relationship into an exploitative one, making one party vulnerable to the other in new ways, or giving the latter new tools for taking advantage of existing vulnerabilities they couldn't previously leverage.

As we've seen, sellers can use algorithmic pricing technologies to capture more and more of—and perhaps even raise—a buyer's reservation price, by engaging in increasingly sophisticated forms of first-degree price discrimination. In part, this means utilizing the particular affordances of digital platforms to take advantage of existing vulnerabilities sellers couldn't previously leverage. Specifically, platforms enable the collection of detailed personal information about each individual buyer, including information about their preferences, finances, and purchasing histories, which are highly relevant to decisions about pricing. And platforms can analyze that information to make predictions about buyer willingness to pay on-the-fly, dynamically adjusting prices in the moment for different buyers (Seele et al. 2020). Thus, while it has always been the case that some buyers were willing to pay more than others for certain

[25] Or perhaps some third thing. See Valerio De Stefano (2016) in *The Rise of the 'Just-in-Time Workforce': On-Demand Work, Crowdwork, and Labour Protection in the 'Gig-Economy'*; Orly Lobel (2019) in *The Debate Over How to Classify Gig Workers Is Missing the Bigger Picture*; Veena Dubal (2021) in *The New Racial Wage Code.*

[26] We often worry about sellers deceiving buyers or selling them unsafe products, and consumer protection law is designed to prevent such harms. But we don't normally worry that sellers will *exploit* buyers.

goods, sellers haven't always been able to tell them apart, or to use that information to take advantage of buyers at the point of sale.

The affordances of digital platforms also create new vulnerabilities, by making prices more inscrutable. Without knowing (or at least being able to make an educated guess about) why a seller has offered a particular price, and without being able to see what prices other buyers in the marketplace are paying, buyers are placed at a significant disadvantage when bargaining with sellers. And lest one think this is "merely" an issue when shopping online, think again: retailers have tested personalized pricing systems for physical stores, where cameras and other tracking technologies identify particular customers and electronic price tags vary prices accordingly (Seele et al. 2020). If sellers deploy such systems they will deprive buyers of access to information about even more of the marketplace, creating new vulnerabilities sellers can exploit.

Moreover, beyond transforming typically non-exploitative relationships into exploitative ones, platforms can create entirely new social relationships, which exist, at least partly, for the express purpose of enabling exploitation. This is the story of "surveillance capitalism." Digital advertising platforms have created sprawling, largely invisible ecosystems of data collectors and aggregators, analytics firms, and advertising exchanges, which data subjects—everyday people—know little about. They have brought into being a new set of relationships (e.g., the data aggregator/data subject relationship), designed from the ground up to facilitate one party extracting from the other.

We should expect more of this the more we integrate digital platforms into our lives. As platforms extend their reach, mediating new contexts, relationships, and activities, the data collection that comes in tow renders us—and our vulnerabilities—more visible, and as platforms become gatekeepers between us and more of the things we want and need—work, goods and services, information, communication—they create new opportunities to take advantage of what they learn.

IV. Conclusion

What are we to make of all of this? To conclude, I want to suggest that the language of exploitation is useful not only as a broad indictment against perceived abuses of power by big tech firms. Understanding platforms as vehicles of exploitation helps to illuminate normative issues central to the present conjuncture.

First, theories of exploitation highlight an important but under-appreciated truth, which challenges prevailing assumptions in debates about platform governance: exchange can be mutually beneficial, voluntary, and—still—wrong.[27] Which is to say, two parties can consent to an agreement, the agreement can serve both of their interests, and yet, nonetheless, it can be wrongfully exploitative. This idea, sometimes referred to as "wrongful beneficence," can be counter-intuitive, especially in the US and other liberal democratic contexts, where political cultures centered on individual rights often treat the presence of consent as settling all questions about ethical and political legitimacy. If two people come to an agreement, there is no deception or manipulation involved, and the agreement is good for both of them (all things considered), many assume the agreement is, therefore, beyond reproach.

Consider again paradigmatic cases of exploitation. When a price gouger sells marked-up goods to someone in need—scarce generators, say, to hurricane survivors—the buyer consents to the purchase and both parties leave significantly better off than they were. Likewise, when a sweatshop owner offers low-paying work in substandard conditions to local labourers and—given few alternatives—they accept, the arrangement is voluntary and serves both the owner's and the worker's interests.[28] Thus, if the price gouger and the sweatshop owner have done anything wrong in these cases, it is not that they have diminished the other parties' interests or forced them to act against their will. Rather, as we've seen, the former taking advantage of the latter is wrongfully exploitative because the treatment is unfair (i.e., the price of the generator is exorbitant, and the sweatshop pay is exceedingly low) and/or degrading (it fails to treat exploitees with dignity and respect).

This insight, that exploitation can be wrong even when mutually beneficial and voluntary, helps explain the normative logic of what Lewis Mumford (1964) called technology's "magnificent bribe"—the fact that technology's conveniences seduce us into tacitly accepting its

[27] As Joel Feinberg (1990, 176) put it, "a little-noticed feature of exploitation is that it can occur in morally unsavory forms without harming the exploitee's interests and, in some cases, despite the exploitee's fully voluntary consent to the exploitative behaviour." Wood (1995), Wertheimer (1996), Sample (2003), and others also emphasize this point.

[28] One might want to argue that the buyer in the first case and worker in the second are "coerced by circumstances," and therefore the exchanges are not truly voluntary. But as Chris Meyers (2004) points out, that's not the price gouger's or the sweatshop owner's fault—they didn't create the desperate conditions, and all they are doing is adding to the sets of options from which the other parties can choose. If in doing so they are wronging them (which, in cases of wrongful beneficence, they arguably are) it is not because they are forcing them to act against their will.

harms (Loeb 2021). Indictments against digital platforms are frequently met with the response that users not only accept the terms of these arrangements, they benefit from them. Mark Zuckerberg, for example, famously argued in the pages of the *Wall Street Journal* that Facebook's invasive data collection practices are justified because: "People consistently tell us that if they're going to see ads, they want them to be relevant. That means we need to understand their interests."[29] In other words, according to Zuckerberg, Facebook users find behaviourally targeted advertising (and the data collection it requires) beneficial, so they choose it voluntarily.[30] Similarly, as we have seen, gig economy companies deflect criticism by framing the labour arrangements they facilitate as serving the interests of gig workers, both economically and as a means of strengthening worker independence and autonomy.

The language of exploitation shows a way through this moral obfuscation. Implicit in tech industry apologia is the assumption that simply adding to people's options can't be wrong. But the price gouging and sweatshop labour cases reveal why it can be: if the only reason someone accepts an offer is they lack decent alternatives, and if the terms being offered are unfair or degrading, then the offer wrongfully takes advantage of them and their situation. So, while it is true that in many cases digital platforms expand people's options, giving them opportunities to benefit in ways they would otherwise lack, and which—given few alternatives—they sometimes voluntarily accept, that is not the end of the normative story. If platforms are in a position to provide the same benefits on better terms and simply refuse, they are engaging in wrongful exploitation and ought to be contested.

Second, having said that, the fact that people benefit from and willingly participate in these arrangements should not be ignored—it tells us something about the wider landscape of options they face. When people buy from price gougers or sell their labour to sweatshop factories they do so because they are desperate. From a diagnostic perspective, we can see that taking advantage of someone in such circumstances is morally wrong. But how, as a society, we should *respond* to that injustice is a more complicated matter. If there aren't better alternatives available to them, eliminating the option—by, for example, banning price gouging and sweatshop labour,

[29] Mark Zuckerberg (2019, January 25) in *The Facts About Facebook*.
[30] Of course, researchers have cast doubt on these claims about user preferences. See Joseph Turow and Chris Jay Hoofnagle (2019, January) in *Mark Zuckerberg's Delusion of Consumer Consent.*

or for that matter, gig work or behavioural advertising—could make the very people one is trying to protect even worse off, at least in the short run (Wood 1995, 156).

As Alan Wood (1995) argues, there are—at bottom—two ways to respond to exploitation: what he terms "interference" and "redistribution."[31] Interference focuses on the exploiter, stepping in to prevent them from exercising power to take advantage of others. Fair labour standards, for example, interfere with an employer's ability to exploit workers, and price controls interfere in the market to prevent gouging. Redistribution, by contrast, focuses on exploitees: rather than directly interfering to keep the powerful in check, redistributive strategies aim to empower the vulnerable. Universal basic income policies, for example, strengthen workers' ability to decline substandard pay and work conditions. Of course, economic support isn't the only way to help the vulnerable resist exploitation—one might think of certain education or job training programs as designed to acheive similar ends.

Differentiating between interference and redistribution strategies is useful for weighing the myriad proposals to rein in platform abuse. Some proposals adopt an interference approach, which focus on constraining the powerful—banning gig economy apps or behavioural advertising, for example, or imposing moratoria on face recognition technology.[32] Others aim to empower the vulnerable: digital literacy programs, for instance, equip people to make better decisions about how to engage with platforms, and forced interoperability policies would enable users to more easily switch platforms if they feel like they're being treated unfairly.[33] Some strategies combine interference and redistribution: if successful, efforts to revive antitrust enforcement in the technology industry would diminish the power of monopoly firms, weakening their ability to engage in exploitation, while also empowering users by increasing competition and thus strengthening their ability to refuse unfavorable terms.[34]

---

[31] Erik Malmqvist and András Szigeti (2021) argue that there is, in fact, a third option—what they term "remediation." To my mind, remediation is a form of redistribution.

[32] Bans and moratoria are frequently proposed, and sometimes implemented, as a strategy for bringing abuse by digital platforms under control. Uber, for example, has been directly banned or indirectly forced out of the market at various times and places (Rhodes 2017). Regulators, especially in Europe, have made compelling cases to eliminate behavioural advertising, especially when targeted at children. See, e.g., https://www.forbrukerradet.no/wp-content/uploads/2021/06/20210622-final-report-time-to-ban-surveillance-based-advertising.pdf And a number of cities in the US have imposed moratoria on the use of facial recognition technology by the police and other public actors, while at the same time it continues to find new applications. See, e.g., https://www.wired.com/story/face-recognition-banned-but-everywhere/

[33] See, e.g., https://www.eff.org/deeplinks/2019/07/interoperability-fix-internet-not-tech-companies

[34] See, e.g., https://www.newyorker.com/magazine/2021/12/06/lina-khans-battle-to-rein-in-big-tech

There are trade-offs involved in the decision to utilize one or the other type of approach. People voluntarily accept unfair terms of exchange when they lack decent alternatives, so interference strategies could do more harm than good if they aren't accompanied by redistributive efforts designed to expand people's options. If people are reliant on crowdwork, for example, because they can't find better paying or more secure jobs, then limiting opportunities for such work might—on balance—make them worse off rather than better, putting them in an even more precarious financial position than where they started.[35] Similar concerns have been raised about behavioural advertising. Despite its harms, observers point out that digital ad markets are "the critical economic engine underwriting many of the core [internet] services that we depend on every day" (Hwang 2020, 1). Interfering in these markets haphazardly could threaten the whole system.[36]

If we step back, however, these insights together paint a clearer and more damning picture than is perhaps first suggested by the careful way I have parsed them above. They suggest that the platform age emerged against a backdrop of deep social and economic vulnerability—a world in which many lacked adequate options to begin with—and platform companies responded by developing technologies and business models designed to perpetuate and exploit them. It is a picture, in other words, of many platforms as fundamentally predatory enterprises: high-tech tools for capturing and hoarding value, and not—as their proponents would have us believe—marvels of value creation. This is, I think, the basic normative intuition behind claims that digital platforms are exploitative, and we shouldn't let our efforts to unspool its implications distract us from the moral clarity driving it.

Moreover, as the Marxist critique emphasizes, what makes exploitation particularly insidious is the thin cover of legitimacy it creates to conceal itself, the veneer of willingness by all parties to participate in the system—their consent and mutual benefit—that obscures the unfairness and degradation hiding just below the surface. As more and more people see through this normative fog, long-held assumptions that digital platforms (as they currently exist) are, at

[35] Once again, questions about these trade-offs mirror debates about how to respond to exploitative sweatshop labor. For a helpful overview of these debates, see Snyder (2010).
[36] Hwang (2020) suggests "controlled demolition" instead. For a more nuanced history and political economy of digital advertising markets see Lee McGuigan (2023) in *Selling the American People: Advertising, Optimization, and the Origins of Adtech*.

bottom, forces for good are losing strength, space is opening up to imagine new, different sociotechnical arrangements, and conditions are improving to advance them.

Bibliography

Alessandro Acquisti, Curtis Taylor, and Liad Wagman (2016). "The Economics of Privacy." *Journal of Economic Literature* 54(2), 442–92.

Mark Andrejvic (2012). "Exploitation in the Data Mine," in *Internet and Surveillance: The Challenges of Web 2.0 and Social Media*, edited by Christian Fuchs, Kees Boersma, Anders Albrechtslund, and Marisol Sandoval. Routledge Press.

Bethany Biron (August 3, 2022). "Number of Uber drivers hits record high of 5 million globally as cost of living soars — with 70% citing inflation as their primary reason for joining the company." *Business Insider*. https://www.businessinsider.com/uber-drivers-record-high-5-million-cost-living-inflation-2022-8

Oren Bar-Gill (2019). "Algorithmic Price Discrimination When Demand Is a Function of Both Preferences and (Mis)Perceptions." *The University of Chicago Law Review* 86(2), 217-254.

Ryan Calo (2014). "Digital Market Manipulation." *The George Washington Law Review* 82(4), 773-802.

Ryan Calo and Alex Rosenblat (2017). "The Taking Economy: Uber, Information, and Power," *Columbia Law Review.*

G. A. Cohen (1979). "The Labor Theory of Value and the Concept of Exploitation," *Philosophy & Public Affairs* 8(4).

Julie Cohen (2019). *Between Truth and Power: The Legal Constructions of Informational Capitalism*. Oxford University Press.

Tressie McMillan Cottom (2020). "Where Platform Capitalism and Racial Capitalism Meet: The Sociology of Race and Racism in the Digital Society." *Sociology of Race and Ethnicity* 6(4), 441-449.

Valerio De Stefano (2015). "The Rise of the 'Just-in-Time Workforce': On-Demand Work, Crowd Work and Labour Protection in the 'Gig-Economy.'" *SSRN Electronic Journal.* https://doi.org/10.2139/ssrn.2682602.

Veena Dubal. (2020). "The Time Politics of Home-Based Digital Piecework." *Center for Ethics Journal: Perspectives on Ethics*. *https://c4ejournal.net/2020/07/04/v-b-dubal-the-time-politics-of-home-based-digital-piecework-2020-c4ej-xxx/*

Veena Dubal (2021). "The New Racial Wage Code." *Harvard Law and Policy Review* 15, 511-549.

Veena Dubal (2023). "On Algorithmic Wage Discrimination." *Columbia Law Review* 123, 1929-1992.

Ariel Ezrachi and Maurice E. Stucke. (2016). "The Rise of Behavioural Discrimination." *European Competition Law Review* 37(2), 485–92.

Joel Feinberg. (1990). *The Moral Limits of the Criminal Law. Vol 4: Harmless Wrongdoing*. Oxford University Press.

Nancy Fraser (2016). "Expropriation and Exploitation in Racialized Capitalism: A Reply to Michael Dawson." *Critical Historical Studies* 3(1), 163-178.

Christian Fuchs (2010). "Labor in Informational Capitalism and on the Internet." *The Information Society* 26(3) 179–96.

Christian Fuchs (2017). *Social Media: A Critical Introduction*. SAGE Publications.

Tarleton Gillespie (2010). "The Politics of 'Platforms.'" *New Media & Society* 12(3), 347–64.

Robert Goodin (1987). "Exploiting a Situation and Exploiting a Person." In *Modern Theories of Exploitation*, edited by Andrew Reeve. SAGE Publications.

Mary L. Gray and Siddharth Suri (2019). *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass.* Harper.

Interactive Advertising Bureau (2022). "Advance Notice of Proposed Rulemaking for a Trade Regulation Rule on Commercial Surveillance and Data Security." https://www.iab.com/wp-content/uploads/2022/11/IAB-ANPRM-Comments.pdf

Tim Hwang (2020). *Subprime Attention Crisis: Advertising and the Bomb at the Heart of the Internet*. Farrar, Straus and Giroux.

Tim Jordan (2015). *Information Politics: Liberation and Exploitation in the Digital Society*. Pluto Press.

Orly Lobel. (July 24, 2019). "The Debate Over How to Classify Gig Workers Is Missing the Bigger Picture." *Harvard Business Review*. https://hbr.org/2019/07/the-debate-over-how-to-classify-gig-workers-is-missing-the-bigger-picture

Zachary Loeb (October 25, 2021). "The Magnificent Bribe." *Real Life Magazine*. https://reallifemag.com/the-magnificent-bribe/.

Alexander MacKay and Samuel Weinstein (2020). "Dynamic Pricing Algorithms, Consumer Harm, and Regulatory Response." https://www.ssrn.com/abstract=3979147

Erik Malmqvist and András Szigeti. (2021). "Exploitation and Remedial Duties." *Journal of Applied Philosophy* 38(1), 55–72.

Lee McGuigan (2014). "After Broadcast, What? An Introduction to the Legacy of Dallas Smythe," In *The Audience Commodity in the Digital Age,* edited by Lee McGuigan and Vincent Manzerolle. Peter Lang.

Chris Meyers (2024). "Wrongful Beneficence: Exploitation and Third World Sweatshops." *Journal of Social Philosophy* 35(3), 319–33.

Pegah Moradi and Karen Levy. (2020). "The Future of Work in the Age of AI: Displacement or Risk-Shifting?" In *The Oxford Handbook of Ethics of AI*, edited by Markus D. Dubber et al., Oxford University Press, 269–88.

Jeffrey Moriarty (2021). "Why Online Personalized Pricing Is Unfair." *Ethics and Information Technology* 23(3), 495–503.

James Muldoon (2022). *Platform Socialism: How to Reclaim Our Digital Future from Big Tech*. Pluto Press.

Lewis Mumford (1964). "Authoritarian and Democratic Technics." *Technology and Culture* 5(1), 1-8.

"Open Letter to President Biden from Tech Workers in Kenya" (May 22, 2024), https://www.foxglove.org.uk/open-letter-to-president-biden-from-tech-workers-in-kenya/

Sebastian Pfotenhauer, Brice Laurent, Kyriaki Papageorgiou, and Jack Stilgoe. (2022). "The Politics of Scaling." *Social Studies of Science* 52(1), 3–34.

Jeffrey Reiman (1987). "Exploitation, Force, and the Moral Assessment of Capitalism: Thoughts on Roemer and Cohen." *Philosophy & Public Affairs* 16(1), 3-41.

John Roemer (1985). "Should Marxists Be Interested in Exploitation?" *Philosophy & Public Affairs* 14(1), 30-65.

Ruth Sample (2003). *Exploitation: What It Is and Why It's Wrong*. Rowman and Littlefield.

Peter Seele, Claus Dierksmeier, and Reto Hofstetter (2021). "Mapping the Ethicality of Algorithmic Pricing: A Review of Dynamic and Personalized Pricing." *Journal of Business Ethics* 170 (4), 697-719.

Jeremy Snyder (2010). "Exploitation and Sweatshop Labor: Perspectives and Issues." *Business Ethics Quarterly* 20(2), 187–213.

Hillel Steiner (1984). "A Liberal Theory of Exploitation." *Ethics* 94(2), 225-241.

Daniel Susser (2019). "Notice After Notice-and-Consent: Why Privacy Disclosures are Valuable Even if Consent Frameworks Aren't." *Journal of Information Policy* 9, 37-62.

Daniel Susser, Beate Roessler, and Helen Nissenbaum (2019). "Online Manipulation: Hidden Influences in a Digital World." *Georgetown Law Technology Review* 4.

Daniel Susser and Vincent Grimaldi (2021). "Measuring Automated Influence: Between Empirical Evidence and Ethical Values." *Proceedings of the ACM Conference on AI, Ethics, and Society*, 1-12.

Zephyr Teachout (2023). "Algorithmic Personalized Wages." *Politics and Society* 51(3), 436-458.

Tiziana Terranova (2000). "Free Labor: Producing Culture for the Digital Economy," *Social Text* 18(2), 33-58.

Zeynep Tufekci (March 19, 2018). "Facebook's Surveillance Machine," *The New York Times*. https://www.nytimes.com/2018/03/19/opinion/facebook-cambridge-analytica.html

Joseph Turow and Chris Hoofnagle. (January 29, 2019). "Mark Zuckerberg's Delusion of Consumer Consent." *The New York Times.* https://www.nytimes.com/2019/01/29/opinion/zuckerberg-facebook-ads.html

US Federal Trade Commission (September 15, 2022), "Policy Statement on Enforcement Related to Gig Work," https://www.ftc.gov/legal-library/browse/policy-statement-enforcement-related-gig-work

Alan Wertheimer (1996). *Exploitation*. Princeton University Press.

Allen Wood (1995). "Exploitation." *Social Philosophy and Policy* 12(2), 136–58.

Jonathan Zittrain. (December 7, 2009). "The Internet Creates a New Kind of Sweatshop." *Newsweek.* https://www.newsweek.com/internet-creates-new-kind-sweatshop-75751

Shoshana Zuboff (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs.

Matt Zwolinski (2012). "Structural Exploitation." *Social Philosophy and Policy* 29(1), 154–79.

Matt Zwolinski, Benjamin Ferguson, and Alan Wertheimer (2022). "Exploitation." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta & Uri Nodelman. https://plato.stanford.edu/archives/win2022/entries/exploitation/