Decision Time: Normative Dimensions of Algorithmic Speed

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

ABSTRACT

Existing discussions about automated decision-making focus primarily on its inputs and outputs, raising questions about data collection and privacy on one hand and accuracy and fairness on the other. Less attention has been devoted to critically examining the *temporality* of decision-making processes—the speed at which automated decisions are reached. In this paper, I identify four dimensions of algorithmic speed that merit closer analysis. Duration (how much time it takes to reach a judgment), timing (when automated systems intervene in the activity being evaluated), frequency (how often evaluations are performed), and lived time (the human experience of algorithmic speed) are interrelated, but distinct, features of automated decision-making. Choices about the temporal structure of automated decision-making systems have normative implications, which I describe in terms of "disruption," "displacement," "re-calibration," and "temporal fairness," with values such as accuracy, fairness, accountability, and legitimacy hanging in the balance. As computational tools are increasingly tasked with making judgments about human activities and practices, the designers of decision-making systems will have to reckon, I argue, with when—and how *fast*—judgments ought to be rendered. Though computers are capable of reaching decisions at incredible speeds, failing to account for the temporality of automated decision-making risks misapprehending the costs and benefits automation promises.

CCS CONCEPTS

• **Applied computing** → **Law, social and behavioral sciences**; • **Social and professional topics** → **Computing / technology policy**.

KEYWORDS

automated decision-making, automation, speed, time, temporality, AI ethics, data ethics

ACM Reference Format:
Daniel Susser. 2022. Decision Time: Normative Dimensions of Algorithmic Speed. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), June 21–24, 2022, Seoul, Republic of Korea.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3531146.3533198

FAccT '22, June 21–24, 2022, Seoul, Republic of Korea

ACM ISBN 978-1-4503-9352-2/22/06…$15.00
https://doi.org/10.1145/3531146.3533198

1 INTRODUCTION

We are automating more and more high-stakes decisions. Judgments about, for example, who gets access to financial credit, who is offered a job, and the risks criminal defendants pose to society are increasingly delegated to machines [84]. The computational tools driving this automation are attracting attention in equal measure from researchers, advocates, and policymakers who question whether such systems conform to or defy shared ethical and social commitments. Existing discussions about automated decision-making systems focus primarily on their inputs and outputs, raising concerns about data collection and privacy on one hand and accuracy and fairness on the other. Less attention has been devoted to critically examining the *speed* at which automated decisions are reached, and its impact on the temporality of decision-making—how decision-making processes unfold through time. In what follows, I argue this is a significant oversight.

The temporality of automated decision-making demands our attention for at least two reasons. First, human behaviors, activities, and practices are inextricably temporal phenomena: they take place, often indeterminately, through time. Conversations on social media, for example, which are increasingly subject to both human and automated judgment, do not happen in an instant. A teacher's impact on students, which is frequently subject to automated performance evaluation, builds and develops over months or years. While certain decisions can be reached by examining a time-slice (e.g., if a driver exceeds the speed limit even for a moment, they are liable to get a ticket), in most cases, evaluating human activity—determining whether it conforms to or defies shared values—necessarily involves accounting for how that activity unfolds over some period. Judgment takes time.

Second, judgment is not simply an epistemic practice, it is a social practice too, and the speed at which judgments are rendered can have normative consequences. Which is to say, the purpose of judgment is not just to find the truth but to *use* that finding to further certain goals. Evaluating someone's creditworthiness or their risk of developing an expensive health condition is not an exercise undertaken purely for the sake of knowing. It is a means of deciding how much money to lend someone and under what terms, or whether to insure them and at what cost. Thinking about decision-making in exclusively epistemic terms (i.e., only attending to the truth and accuracy of decision-making processes) thus misses something essential. To fully appreciate the significance of automated decision-making we have to take stock of its impacts on the social functions of judgment, which requires accounting for changes in decision-making speed. This is especially true when decisions are made in certain social and institutional contexts—for example, as the above cases suggest, in the context of financial lending or the provision of health insurance, but also, as we will see, in contexts such as law enforcement and the regulation of

public speech—where questions about *how judgments are rendered* and *how interventions are made on the basis of those judgments* have special normative significance. We expect decision-makers to be deliberate, impartial, and accountable, and their judgments to be transparent, accurate, and fair. Speeding up decision-making processes can change whether and how they live up to these values.

The literature on fairness in artificial intelligence and machine learning reflects a growing awareness that merely focusing on the statistical distributions of decision-making outcomes is insufficient, either for ensuring that digital technologies don't exacerbate existing systems of marginalization and oppression, or (especially) for developing technologies that actively advance social justice. In this vein, recent work has called for closer attention to what is bracketed by certain fundamental computational abstractions [72], pointed to distinctions between procedural and substantive conceptions of fairness [6, 31], and emphasized, more generally, the limits of bias-oriented antidiscrimination frameworks [40]. Likewise, the discussion that follows aims to widen the lens of technology ethics and policy research, drawing attention to a feature of automation that has been relatively neglected—its temporality—but which, I argue, has important implications for the field.

We are more used to thinking spatially than temporally—e.g., via metaphors of space—so identifying and wrestling with questions of time and temporality can be difficult. For the sake of clarity, I distinguish effects of algorithmic speed along four dimensions—duration, timing, frequency, and lived time—and I discuss the normative implications of these changes. It is important to emphasize, however, that these are analytical distinctions, drawn in order to make the problems I want to raise easier to grasp.[1] In reality, duration, timing, frequency, and the lived experience of time are deeply intertwined. Timing (when something happens) is often a function of its duration (how long it takes), and the shorter the duration of the process, the more frequently it can occur in the same amount of time. Of course, we experience all of this, so adjusting any of these variables changes those experiences.

Furthermore, the schema I put forward is not intended to be comprehensive. There are other temporal dimensions of automation that deserve careful attention and analysis. My aim in what follows is to demonstrate the importance of time and temporality for broader conversations about the ethics and politics of automated decision-making, and to encourage further exploration in this space.

Having described various impacts of algorithmic speed on the temporalities of socially significant decision-making processes, I explore some normative consequences of these changes. Here again, I draw analytical distinctions, highlighting normative dynamics that are, in actuality, closely interrelated. And here too my argument is not exhaustive. There are, I am sure, other social, ethical, and political issues related to algorithmic speed that deserve our attention. In this paper, I describe four: (1) normative *disruption*—replacing pre-existing temporal norms with new norms; (2) *re-calibration*—changing the balance of normative trade-offs or the relative weights assigned to competing values; (3) *displacement*—shifting costs from one part of the decision-making process to another; and (4) *temporal fairness*—how the risks and benefits of automated decision-making are distributed across decision-making subjects.

Finally, to be clear, I am not arguing that speed is bad—at least not always. Rather, my aims are (1) to dislodge the assumption that speed and efficiency are value-neutral, or that they are axiomatically good (an unstated premise in many discussions about technology and an official mantra in Silicon Valley, where there is a cultural imperative to "move fast and break things"), and (2) to offer conceptual tools for understanding the normative implications of speeding up important decision-making processes. This project thus joins with related calls for more "friction" [20] and "seamfulness" [12] in user interface design, consideration of "desirable inefficiencies" [63], and other critical technology scholarship challenging the notion that faster, smoother, more efficient technologies necessarily serve individual and collective ends.[2]

2 NORMATIVE TIME

Algorithms save time; this is a chief source of their appeal. But using computational tools to speed up (or slow down) certain decisions is not a "neutral" adjustment without further implications. Theorists across many disciplines, including historians, philosophers, social and political theorists, have explored the idea that time and temporality can be normative. Which is to say, how activities unfold through time—how slowly or quickly, at what intervals, and so on—is value-laden. Space permits only a brief and selective overview of some of those accounts here, but such a gloss will set the stage for the discussion that follows.

Philosophers have long raised questions about the nature of time and its significance for human experience, from the distinction in ancient Western philosophy between *chronos* (roughly, objective time) and *kairos* (temporal context) [74] to Kant's theory of space and time as "pure forms of intuition" [33] to Marx's theory of value as a function of "socially necessary labor time" [16]. In the 19th and early 20th centuries, especially among philosophers in the European Continental tradition, time and temporality became more central—for example, in Henri Bergson's response to Kant in *Time and Free Will* [4], Martin Heidegger's metaphysics and theory of meaning in *Being and Time* [36], and Edmund Husserl's phenomenological investigation of time consciousness [44].

That we live in an "accelerating" society is an idea familiar from late twentieth century history and social theory. Harmut Rosa gives this idea its clearest expression, arguing that "social acceleration" has three components: "Technological acceleration" is the speeding up of communication, transportation, and production (including the production of new technologies themselves) [68]. The "acceleration of social change" points to transformations in the way people live their lives. As Rosa writes, "attitudes and values as well as fashions and lifestyles, social relations and obligations as well as groups, classes, or milieus, social languages as well as forms of practice and habits are said to change at ever increasing rates" [68, p. 7]. And finally, "the pace of life" is accelerating. Not only is the way we live

[1] Already, I've reverted to spatial metaphors.

[2] While "speed" and "efficiency" are conceptually distinct, I generally use the two terms in tandem. In computing, efficiency is often defined as faster processing with fewer resources: famously "Moore's Law" predicted that processor speed would steadily increase over time as transistors got physically smaller. And, as Ohm and Frankle argue, imposing *in*-efficiency generally means designing "speed bumps" into software or hardware systems—i.e., making them slower [63].

changing at a fast clip, but it is changing *into* a faster way of living. Our days are more harried. We literally do more and more—engage in a greater number of activities—in the same periods of time [68, pp. 9-10].

These changes matter. According to Rosa, social acceleration has led to the "disintegration" of both individual personality and political organization. Given the frenzy of contemporary life, it is difficult for individuals to orient themselves around "time-resistant priorities," with many people reporting feeling "directionless" as a result [68, p. 20]. Analogously, governments are confronted with more decisions to make and less time in which to make them, challenging their capacity to enact the "deliberate and democratic political shaping of our society and form of life" [68, p. 24].

More specifically, William Scheuerman argues that social acceleration destabilizes specific liberal democratic political norms. Most importantly, it throws off-balance the separation of powers between the legislative, executive, and judicial functions of government. Traditionally, the legislative branch was meant to be slow, deliberative, and future-oriented, creating stable, long-lasting laws that enable individuals to organize their lives and groups to coordinate [71, pp. 29-30]. By contrast, there is an "intellectual association of the executive with dispatch" [71, p. xvii]. As the world becomes faster, Scheuerman argues, we become more presentist—concerned less with planning for the future than with managing the moment. Since that is understood to be the executive's strength, power gravitates toward it from the other branches of government.

Recently, Elizabeth Cohen has rendered "the political value of time" more concrete. Time is an "inextricable part of political foundation" [15, p. 6], she argues, because it is "one of the most precious and finite resources required for the accomplishment of human purposes" [15, p. 1]—rights and liberties are not of much value if we don't have time to use them. Moreover, states use temporal boundaries to structure democratic processes and to condition access to rights, goods, and opportunities: Holding elections at regular intervals gives citizens influence over the direction of government and serves as an imprimatur of democratic consent. Statutes of limitation condition punishment of criminals and redress for victims of crime. Mandatory waiting periods structure access to legally protected rights, such as abortion. For these reasons and others, Cohen argues, how the state treats people's time and how political processes are temporally structured are central to social justice [15].

The relationship between time and *technology* (beyond the fact that technology is an engine of acceleration) is explored in sociology and science and technology studies.[3] Paul Virilio's "dromology," or theory of speed, is a common point of departure. For Virilio, speed is an instrument of power [46, 81]. Robert Hassan focuses specifically on information and communication technologies (ICTs), and on the "network time" we are forced, by our reliance on ICTs, to inhabit [34]. Judy Wajcman resists that kind of determinism in her investigation of the "time-pressure paradox"— the tension between people's reported experiences of time poverty and the fact that, on average, we have more leisure time than ever—insisting that speed is just as much a social imperative we design into our technologies as it is a value they impose on us [82]. Donald Mackenzie explores the financial sector's construction of ever faster communications technologies to facilitate high-frequency trading, illustrating the sociomaterial conditions of digital speed [59]. And Steven Jackson urges a temporal re-orientation in critical technology scholarship, away from the high speed of "cutting-edge" innovation, toward the slow, enduring work of building and maintaining technological infrastructure [45].

Several themes from this expansive literature are instructive for present purposes and will shape the discussion that follows. First, temporality can be normative: many human behaviors and practices are guided by norms that determine how fast or slow they should happen, at what intervals, and so on, and changing those norms has social, ethical, and political consequences. Second, digital technologies tend to speed things up—indeed, the speed and efficiency of digital technologies are often their main selling points—and accelerating social practices can mean changing the temporal norms that govern them. Third, acceleration is not inevitable. We control, to a large extent, the temporal effects of our technologies, and we can decide when speed is and isn't desirable. Fourth, these issues are essential to the pursuit of just technology.

Given the individual, social, and political significance of time, and the manifest impacts of digital technologies on how we live and experience time, one might expect questions about time and digital automation to feature prominently in discussions about technology ethics and policy. Yet such questions have attracted surprisingly little attention.[4] One important strand of scholarship approaches temporal questions obliquely, exploring how predictive systems trained on data that reflects historical injustices can, over time, "compound" them [27, 37, 38]. More relevant for the present discussion, critics and scholars have described the temporal effects of automation in certain specific domains, such as law enforcement. Some have argued, for example, that the use of predictive analytics by police embeds in police practice a pre-emptive logic that is incompatible with human agency and politics [2], that this logic exemplifies a dangerous, "paranoid style" of governmentality [73], and that it fails to treat the people policed as full moral subjects [77].

These discussions are fruitful and I build on their insights, but they are relatively abstract. I want to shed light on the specific effects of algorithmic speed—its concrete forms and characteristics—and their normative implications. What follows thus complements these other projects, framing questions about time and automation, and providing conceptual tools for analyzing them, that are, I hope, especially relevant to practical debates about technology governance.

3 TEMPORALITIES OF AUTOMATED DECISION-MAKING

3.1 Duration

Whether a human or machine is at the helm, reaching a decision takes time. How much time—the duration of decision-making processes—is, in part, a function of the quantity of inputs and the capacity of decision-makers to process them. Capable of churning through huge amounts of information in a small amount of

[3] Some interesting examples can be found in [35] and [83].

[4] An important, if brief, exception was the "Speed Conference" held at Cornell Tech's Digital Life Initiative in 2018: https://www.dli.tech.cornell.edu/speed

time, computers can dramatically reduce the duration of decision-making, relative to the manual processes they replace. For example, the infamous COMPAS risk assessment algorithm, which takes as inputs responses to a long, written survey, is able to output a risk score instantly [3]. The decision-making process is shorter in duration than it would be if a human had to review all of the same information. It takes less time.

What is often ignored, however, is that the *objects of judgment*—the things being evaluated—are often behaviors or practices that have their own durations. Human activity unfolds over time, and judgments rendered too quickly can misapprehend it.

Consider the case of automated content moderation, much discussed in the literature on automated decision-making [See, e.g., 7, 10, 17, 30, 51, 67]. Algorithms that detect individual words or phrases are easy to create, and advanced techniques in natural language processing (NLP) can classify the general sentiments expressed in a text with some accuracy [85]. But a variety of technical challenges stand in the way of developing computational tools that can perform the kind of sophisticated interpretive tasks content moderation demands, such as distinguishing between extreme and extremist speech or hateful and hate speech. Machine learning algorithms, which power content moderation systems, require huge amounts of human-labeled training data, and training data is needed in each language the systems will encounter; they need clearly defined definitions of problematic content, often elusive in areas of social policy where considerable disagreement exists about norms; and they are domain-specific, meaning models trained to detect hate speech can't also detect extremist content [25].

Future advances in NLP may solve some of these problems.[5] However, automated content moderation faces non-technical challenges too, including problems of temporal duration. Not only are the scale and pace of content creation difficult to manage [30], speech is deeply *contextual*, and the context needed to interpret and understand particular utterances often takes time to emerge; thus, the information needed to normatively evaluate speech might not exist at the time of decision. What might appear abusive at time t_1 could turn out to be slang, satire, or jest between friends. Consider slurs that have been reclaimed by their targets, such as affirming uses of the term "queer" by many in LGBTQ+ communities [58]. A message reading "What a queer!" carries a different connotation when exchanged between two queer-identified people than it does when exchanged between straight people. And having access to information about the interlocutors may be insufficient to evaluate the exchange (perhaps one is gay but not out publicly). In many cases, the missing ingredient will be time—further discussion that situates the utterance in a larger discourse.

As another example, imagine a tweet that reads, "You bastard! I'm gonna kill you!!" On the face of it, that looks bad. But later posts might reveal that the recipient of the message is a friend of the sender who lured him to a surprise party. Of course, more time will not always solve the problem. When interlocutors know each other they can implicitly interpret speech through background assumptions (e.g., "My friend would not speak to me in an abusive way"). In such cases the necessary context might never be made explicit. But in many cases the context will arrive, once the conversation has had time to unfold. The fairness and accuracy of automated content moderation systems can depend, then, on how much time they allow before rendering judgment. When designing such systems, one may need to consider trade-offs between speed and efficiency, on one hand, and accuracy and fairness, on the other.

Indeed, the problem is even more complicated than that. Conversations are not just contextual, they are *indeterminate*—the direction a conversation will take is unknowable in advance.[6] What begins as casual banter sometimes turns serious. Disagreement may lead to consensus. Acrimony can resolve into appreciation. Content moderation systems miss this when they fail to take the duration of conversations into account. A rush to judgment might thwart it entirely, ending a conversation before it has time to reach resolution, foreclosing the good that dialogue can bring. Beyond unfairness and inaccuracy, inattention to duration can thus lead to a kind of irreparable interpersonal and social harm, depriving people of the opportunity for transformation and reconciliation. To be sure, protecting people from extremism, hatred, and abuse is a worthy goal, and efforts to automate content moderation in order to facilitate such protections at scale are laudable. But such systems are likely to remain imperfect, more so if they neglect issues of temporal duration. The costs described above must therefore be weighed against the benefits such imperfect systems promise.

3.2 Timing

Questions about the duration of automated decision-making are interrelated with, but distinct from, questions about its timing. *When* should judgment be rendered? As the content moderation example in the previous section illustrated, decisions about when to judge—and when to intervene—can (implicitly or explicitly) involve important trade-offs. This question admits of a great deal of nuance and specificity, as modern technology allows for decision-making to be very precisely timed. When the question has figured in previous discussions about digital technologies and automated decision-making, however, it has generally been framed in binary terms: should judgment and intervention be rendered ex ante or ex post, before or after the fact? Which is to say, should decision-making systems try to prevent undesirable behavior or should they wait and sanction it afterward?

Examples of ex ante enforcement strategies abound. Rather than punishing people for entering rooms they are not permitted in, we lock the door. Instead of waiting for firearms to be illegally discharged, we find and confiscate them using metal detectors and body scanners. Urban planners and security experts refer to this kind of preemptive strategy as "target hardening"—to minimize crime, advocates argue, construct the built environment in ways that make committing crime as difficult as possible.[7]

[5] New software used by intelligence agencies to detect online disinformation may signal some progress [52].

[6] See, e.g., Hans-Georg Gadamer's *Truth and Method*: "We say that we 'conduct' a conversation, but the more genuine a conversation is, the less its conduct lies within the will of either partner. Thus a genuine conversation is never the one that we wanted to conduct. Rather, it is generally more correct to say that we fall into conversation, or even that we become involved in it. The way one word follows another, with the conversation taking its own twists and reaching its own conclusions, may well be conducted in some way, but the partners conversing are far less the leaders of it than the led" [29, p. 385].

[7] Cozens and Love describe target hardening as one of seven principles of "Crime Prevention through Environmental Design (CPTED)" [21].

Preemptive enforcement is common in digital contexts too. Access controls, such as passwords and biometric authentication, are the equivalent of digital locks [50]. Internet filters help parents keep children from visiting websites for adults, and they allow governments to block citizens from accessing or distributing dissident news.[8] There are even websites and apps that facilitate self-directed preemption—tying oneself to the mast, so to speak, digitally. For example, some Orthodox Jews use internet filters to keep from accidentally encountering "unkosher" content, and apps like Freedom disable internet access on computers or smartphones for a certain amount of time, to "free" users from unwanted (but tempting) distraction [61].

Some scholars have argued that ex ante intervention might be especially warranted in digital contexts, because the efficiency with which digital crimes can be committed needs to be counteracted. Neal Katyal refers to this as "cost deterrence." "Cyberspace," he argues, "presents unique opportunities for criminals to reduce their perpetration costs; the probability of success in inflicting a certain level of harm while holding expenditures constant is greater. Accordingly, the law should develop mechanisms to neutralize these efficiency advantages" [49, p. 1006]. Conveniently, ex ante intervention is especially easy to implement in digital environments, because their design is entirely subject to our control. As Katyal puts it, following internet law scholars like Lawrence Lessig and Joel Reidenberg: "the Internet, an artificial environment, is all architecture (or code) and thus infinitely malleable, at least in theory" [48, pp. 1041-2][56, 65].

Yet this kind of strategy carries considerable risks. An early lesson in the implications of using digital tools as a mechanism for ex ante enforcement centered on the development of digital rights management (DRM)—encryption techniques that restrict access to digital media—which enjoyed widespread adoption in the late 1990s and early 2000s, when the music and film industries began to contend seriously with the proliferation of their content online. Much discussed in the academic literature, DRM exemplified the desire for what legal scholar Jonathan Zittrain termed "perfect enforcement" of laws, norms, and market protections [86].

As Zittrain argues, when the speed and efficiency of ex ante strategies begin to approach perfect (i.e., inescapable) enforcement the risk of over-prevention become especially worrying. First, because rules are often poorly articulated and imperfectly operationalized, enforcing them too effectively can "amplify" and "lock-in" errors [86, pp. 114-17]. Second, the *inability* to perfectly enforce laws is a check against government overreach; ipso facto the nearer we get to perfect enforcement, the greater government's power grows [86, pp. 117-19]. Third, following Tim Wu, Zittrain argues that there may be cases where even the alleged victims of rule-breaking would prefer the rule not to be enforced—cases of so-called "tolerated uses." For example, owners of copyrighted materials shared online may have standing to object, but they might choose not to, preferring instead to enjoy the free promotion [86, pp. 119-22]. To the extent that preemptive enforcement efforts preclude such tolerated uses, they could make worse off even those the rules are designed to protect.

[8]China's "great firewall" is the most notorious example [14].

The near-instantaneous speed of algorithms opens the door to many more practices of this kind. The examples I began with illustrate enforcement that is automated (the work of preempting bad behavior is delegated to machines), but the *decision* to enforce, to determine when enforcement is warranted, for the most part, is not. Locks are installed at a particular time and the keys are thoughtfully distributed. DRM techniques are applied at the point of production. Internet filters are manually configured in advance.

The speed of automated decision-making enables preemptive decisions even in cases where judgment is required *in the moment.* Banks deploy fraud detection algorithms that instantly evaluate and halt suspicious payments [32]. Websites featuring user-generated content use automated identification tools—such as YouTube's "ContentID" system—to prohibit the sharing of copyrighted material (or to facilitate payment for it) [See, e.g., 57]. Lawmakers have proposed legislation requiring new vehicles to have built-in breathalyzer devices that prevent cars from starting if a driver's blood alcohol level is above a certain threshold [55]. Or, consider again the example of online content moderation. Pressure is mounting for social media companies not only to manage hate speech and extremist content on their platforms, but to prevent it entirely [7]. Algorithms make this possible: in the split second between clicking "post" and messages appearing on others' screens, content filters can detect key words or images and block or degrade their transmission [51].

In this way, algorithmic speed makes possible new, dynamic forms of target hardening, enabling preemptive enforcement on-the-fly [2]. The question, of course, is whether—in any particular case—this is desirable. Unless the designers of decision-making systems treat the issue of timing with care and attention, preemption may become the default. The value of speed and efficiency is, to many, self-evident—absent explicit reasons to resist it, technological development tends in its direction. As the speed of automated decision-making increases—as the time it takes to render decisions approaches zero—we may find more and more judgment occurring ex ante, whether or not that serves agreed-upon ends.

3.3 Frequency

To understand the significance of duration, above, we had to consider the temporal nature of the behaviors or activities being judged, asking how much time it takes for them to unfold before meaningful judgment can be rendered. To understand the significance of timing, our attention turned to the impacts of speed on decision-making interventions—specifically, the possibility that automation will lead to more preemptive enforcement. The issue of frequency (how often decisions are reached) raises related questions, and also prompts additional concerns about the *purpose* of judgment. By speeding-up decision-making processes, algorithmic decision-making enables more frequent judgment. Indeed, in some cases, where once we were subject to occasional evaluation we are now evaluated constantly.

One purpose of judgment is to detect rule breaking. In such cases, the decision is binary: one has or hasn't broken the rules. Let's call this *detecting infringement.* Examples of detecting infringement—automatically or not—are commonplace. Traffic enforcement technologies, such as red light and speed cameras, attempt to determine whether or not someone has broken rules of the road. Ankle

and wrist monitors alert authorities if someone under physical confinement has violated the terms of their detention by leaving their home or workplace [1]. Face recognition systems identify people illegally crossing borders [84].

Using automated decision-making tools to detect infringement is like setting a tripwire—the system checks constantly to see if a rule has been broken and raises a flag if it has. Here, the fact that algorithmic speed enables more frequent judgment simply means the tripwire is more precise. The more times per minute, say, an ankle monitor can detect its location via GPS, the faster it will alert authorities when it moves beyond the prescribed perimeter. Many of the same normative concerns discussed above, such as worries about perfect enforcement, attach to these technologies. Moreover, how they are distributed—who is made subject to their judgment, and who is not—raises urgent moral and political questions. For example, black Americans are disproportionately subjected to electronic monitoring, and therefore disproportionately bear the risks they create.[9] I discuss these fairness issues below.

But detecting infringement is not the only purpose automated decision-making systems serve. In many cases, rather than checking to see if a rule has been broken the goal of automated decision-making is to evaluate and categorize people according to some more flexible standard. For example, credit scoring companies use digital tools to determine how creditworthy consumers are [64]. Courts use predictive systems, such as COMPAS, discussed above, to evaluate how likely criminal defendants are to recidivate [3]. Universities use data-driven technologies to calculate the risk of students dropping out [8]. Social services agencies use algorithms to determine who is in greatest need of public assistance [28].

I refer to these kinds of judgments as *assessments,* which differ from detecting infringement in at least two ways. First, unlike detecting infringement, which is usually a binary decision—someone either has or hasn't broken the rules—assessments involve categorizing people or situating them along some kind of spectrum.[10] Today, what is often assessed is *risk*, as the examples above illustrate. But whether we realize it or not we are each subject to endless assessments. Our health and wealth are "scored," as is everything from our energy use to social influence to the likelihood we will be the victims (or perpetrators) of fraud [13, 24]. Second, infringement is usually synchronic: there is a moment when rule-breaking occurs. By contrast, assessment tends to be diachronic: when evaluating someone's risk of recidivating or the probability that they will drop out of school, the inputs to that assessment stretch across time. A person's creditworthiness is not judged from information gathered about them in a single moment; it requires a credit *history*, a record of debits and credits, payments and non-payments.

Whereas increasing the frequency of judgment in order to detect infringement simply produces more precise, finely-calibrated results, doing the same in cases of assessment transforms the very nature of the decision-making process. Assessment becomes more like detecting infringement: rather than slowly developing a composite picture of someone and rendering judgment about their behavior as a whole, automated decision-making systems impose endless tripwires and wait for them to be set off. China's "social credit" system exemplifies this phenomenon in the extreme, where a wide range of everyday behaviors are monitored and one perceived misstep can trigger immediate sanction. "Caught jaywalking, don't pay a court bill, play your music too loud on the train—you could lose certain rights, such as booking a flight or train ticket" [53]. But, in principle, this is not any different from the kinds of scoring people elsewhere face. The Chinese system only appears more totalizing because the digital infrastructure that facilitates it is more centralized and its assessments are used to accomplish traditionally public, institutional functions [22].

Indeed, as digital surveillance has become ubiquitous, data about our behavior—practically at every moment—now feeds automated decision-making systems. The result, as everyone has experienced, is *relentless assessment*. Instead of being evaluated occasionally, at relatively well-understood intervals, we are judged constantly, every move we make serving as indication that we are more or less trustworthy, dangerous, healthy—whatever those judging us care to assess [64]. And once again, these transformations are not "neutral." In cases of assessment, increasing the frequency of judgment—making judgment constant, *relentless*, rather than occasional, episodic—means taking some "give" out of the system, making judgment more reactionary, more rigid, and more myopic. Decision-making subjects have fewer opportunities to make up for missteps, less time to notice they are moving in the wrong direction and to correct course. When each data point immediately impacts an assessment less attention is given to the broader context, and there are fewer opportunities for grace.

3.4 Lived Time

The duration, timing, and frequency of automated decision-making are objective dimensions of its temporality, design choices that can be measured with clocks and calendars. But judgment is also *experienced* by the judged, and the temporal transformations discussed above change those lived experiences. I can only speculate here about how such changes are felt and about their significance for decision-making subjects. Nevertheless, the effects of algorithmic speed on lived experience are too important to ignore, so I raise the following as prompts for further reflection.

First, there is evidence that people find relentless assessment, and the surveillance that fuels it, *creepy* [64, 78]. Think of the widespread suspicion that smartphones persistently listen to our conversations and use what they hear to target advertisements. What people find disconcerting is the *timing*: "I was having a conversation with my sister about buying clothes and she recommended Etsy, then *moments after* that conversation ended I was browsing Reddit and saw an Etsy ad out of the blue," one person reported [43]. "A friend and I were sitting at a bar, iPhones in pockets, discussing our recent trips in Japan and how we'd like to go back. *The very next day*, we both received pop-up ads on Facebook about cheap return flights to Tokyo," recounted another [62]. Despite little evidence this is really going on and more plausible explanations for what people are experiencing, the sense that our devices are constantly

[9] "The reality of electronic monitoring today is that Big Brother is watching some groups more than others. No national statistics are available on the racial breakdown of Americans wearing ankle monitors, but all indications suggest that mass supervision, like mass incarceration, disproportionately affects black people. In Cook County, Illinois, for instance, black people make up 24% of the population, and 67% of those on monitors" [54]. See, generally, [11].

[10] That is not to say whether or not someone has broken a rule isn't sometimes ambiguous and contestable; oftentimes it is. See, e.g., [9]

listening to us appears difficult for most people to shake [66]. And while some might dismiss these feelings as inconsequential—mere paranoia—there is reason to take them seriously. Frank Pasquale suggests that they are "world disclosive"—"an emotional reaction that alerts us to the real possibility of harm" [64, pp. 37-38]. A leading indicator, perhaps, that businesses are using digital technologies in ways that contravene important social norms, like expectations of privacy [75, 78].

Second, at an individual level, without time to reflect on, contextualize, and understand decisions, it is difficult for people to meaningfully internalize—or resist—them. While in some cases judgment merely serves to distribute goods, mitigate risks, or accomplish some other momentary goal, in many contexts of interest to technology ethics and policy scholars there are presumably deeper, longer-term purposes at work. Credit scores are not just a tool for banks to manage risk, they give feedback to borrowers about what constitutes good financial behavior [41]. More than simply ridding platforms of harmful media, content moderation should signal community speech standards [60]. In other words, we judge people not only to decide how to treat them, but also to *influence* their future behavior—we aim to discipline, if also to punish. But impacting people in this way, shaping their behaviors and habits, requires giving them time to reflect on and understand decisions reached about them [76]. And systems of accountability that provide for feedback in the other direction—allowing people to question and challenge decisions—requires time to deliberate about and evaluate those judgments.

Third, at a social level, how people perceive automated decisions—if they experience them as creepy, invasive, or unfair, or if they happen so fast they aren't perceived at all—could affect their legitimacy. Which is to say, unless decisions are experienced in certain ways, people may refuse to voluntarily accept them [79, 80]. Numerous critics have warned that incorporating digital technologies into socially significant decision-making processes could threaten their legitimacy. For example, Danaher argues that the opacity of some algorithmic decisions—i.e., difficulty understanding their motivating logic—is a fundamental problem for their legitimacy [23].[11] For Binns, legitimacy hinges on whether decisions meet standards of public reason—sometimes automated decisions meet those standards, sometimes they do not [5]. Henin and Le Métayer suggest what is really essential is that decisions are *contestable*—that people have the opportunity to challenge them—and whether or not automated decision-making systems provide such opportunities is a function of their design [39]. Whatever one's view, algorithmic speed calls all of these conditions into question. If decisions are rendered too quickly and too often, people are likely to experience them as beyond understanding, contestation, and accountability.

4 NORMATIVE EFFECTS

In the previous section, I outlined four dimensions of algorithmic speed—four ways the speed of digital automation is transforming the nature, function, and social effects of important decision-making processes. That such changes have normative significance is obvious, I hope, from those descriptions. But to give them clearer form, in this section I describe the normative implications of algorithmic speed in more general terms. As in the previous section, the discussion that follows is not exhaustive. There are, to be sure, important implications of speeding-up (or slowing down) decision-making beyond those I identify here. My aim is merely to begin exploring these issues systematically.

4.1 Disruption

As we've seen, decision-making speed is determined, in part, by how much information the decision-maker has to consider, and how quickly they are able to process it. The reason algorithmic decision-making systems generally outpace human deciders is that computers can churn through huge amounts of information at incredible speeds. As I discussed in section 2, however, the temporalities of decision-making processes can also be shaped or constrained by social, ethical, legal and other norms. The US Constitution, for example, guarantees criminal defendants the right to a "speedy trial." Some cases are so complex, and the costs of error so high, one might think that there should be no limit on the time and resources spent adjudicating them. But we have decided otherwise: both for the sake of the defendant (who might be detained while awaiting a verdict) and the public (whose resources are required to carry out the trial), a decision must be reached within a given period—a limited duration—even if that means excluding relevant information from the process.[12]

By contrast, legislative processes are supposed to be "slow going, deliberate" [71, p. xvi]. Here, the mandate is precisely the opposite of the speedy trial—suspicion is cast not upon too-long decision-making processes but too-short ones. The US Senate, famously described as "the world's greatest deliberative body," is a procedural Rube Goldberg machine designed to slow overeager lawmakers. Of course, in the right hands and in the right circumstances it can be made to move quickly. Days after the September 11th terrorist attacks, the US Congress quickly passed the Authorization for Use of Military Force of 2001, a law that still—more than two decades later—empowers US presidents to engage in virtually unconstrained military action abroad. But such cases, where legislative bodies act too expeditiously, often provoke outcry, becoming exceptions that prove the rule.

Having surfaced these temporal norms—trials should be speedy, legislation slow—we can see why the incorporation of digital technologies into such decision-making processes might be more or less desirable. Technologies that speed up lawmaking may be a problem, whereas we might welcome those that help speed up trials. For example, "e-discovery" tools promise to help lawyers more quickly sift through the mountains of documents involved in many court battles [70]. All else equal, the incorporation of such tools into legal work may therefore signal a net good.

The same is true of automating other important decision-making processes: some kinds of decision-making are governed (often implicitly) by pre-existing temporal norms, which automation can disrupt. Worries about the increasing frequency of judgment, discussed above, are illustrative. It was once reasonable to expect, for example, that one's credit score would be updated more-or-less monthly—increased if bills were paid on time and decreased if

[11] Others argue this problem is overblown [69].

[12] https://www.americanbar.org/groups/criminal_justice/publications/criminal_justice_section_archive/crimjust_standards_speedytrial_blk/

they weren't. Now, the speed of algorithmic judgment means that credit bureaus adjust scores constantly, in response to information about anything a borrower does, received from any of the countless sources of digital behavioral tracking [64]. As technology ethics and policy scholars have argued, this makes understanding and contesting scores increasingly difficult, raising urgent questions about transparency, due process, accountability, and legitimacy [13, 64].

Of course, disruption is not always bad. There may be cases where the status quo norms regulating decision-making speed are unjust, and where new temporal norms are called for. In such cases, automation and the speed that comes with it could be an improvement, all things considered. What I want to draw attention to here is that by changing the temporality of decision-making automation can alter the normative landscape. Determining whether such changes are ultimately good or bad, just or unjust, requires further analysis. In many cases it likely calls for public discussion and debate.

4.2 Re-Calibration

A second effect of the temporal changes described above is they can—sometimes inadvertently—upend a previously settled balance of normative trade-offs. The kinds of decisions discussed throughout this paper almost always require carefully calibrating competing values. Because the speed at which decisions are made can privilege certain values over others, speeding up or slowing down decision-making can mean changing how different values are weighed.

Consider trade-offs between speed and accuracy, discussed above. On one hand, we think decision-makers should strive for maximal accuracy: their assessments should be factually correct and logically sound. On the other hand, real-world decisions usually have deadlines: to be useful or fair, the decision must arrive in a reasonable amount of time. Yet these two goals are often in tension. In general, more time and effort deliberating—less efficient decision-making—leads to more accurate results. This trade-off is at the heart of behavioral economics, typified by Daniel Kahneman and Amos Tversky's "thinking fast and slow" model of cognition [47]. So-called "System 1 thinking" is intuitive and fast. But the mechanisms responsible for that speed—biased cognitive heuristics—are prone to error. By contrast, "System 2 thinking" is slow and deliberative, more accurate but also less efficient. Good decision-making requires balancing the efficiency of System 1 and the accuracy of System 2. As Kahneman writes, "it minimizes effort and optimizes performance" [47, p. 25].

The same trade-offs have to be managed in social and political decision-making processes. A standard example, previously mentioned, is the US right to speedy trials. There are good reasons to value both accuracy in judicial decision-making and efficiency in judicial decision-making processes, and striking the right balance between these two values—agreeing about how much error we can tolerate and how much time we can devote to minimizing it—is a complex normative challenge. Likewise in myriad other contexts: How much time should food safety regulators spend inspecting each production facility? How often should social workers check in on at-risk children? How long should a physician spend with each patient in order to understand what they are experiencing and how to respond?

Examples of this problem are everywhere, as Cooper et al. argue: "it can be observed in a diverse range of disciplines, including economics, law, and medicine," and—as they explore—in every area of computing [18, p. 2]. "There is an inherent tension between *how correct computations are* and *how long it takes* to compute them," [18, p. 1] and when computational tools are incorporated into high-stakes decision-making these trade-offs come with them. To illustrate, Cooper et al. discuss autonomous vehicles. Driving involves endless split-second decisions: Is the road ahead wet or icy? Speed up or stop short at the intersection as the light turns red? Drive over the pothole or swerve around it? In each case, accuracy is critical and time is of the essence—what is the correct balance?

Incorporating computational tools into decision-making processes carries with it value-laden choices about how to weigh these trade-offs, choices that may or may not conform with previously settled arrangements. As we saw in the section on duration, misalignments between the speed of decision-making and the temporality of the behavior or activity being judged—e.g., content moderation systems that rush to judgment before relevant context can emerge—risk decision-making processes that systematically misapprehend their subjects. Choosing to use such systems without attending to questions about their temporalities therefore implicitly prioritizes speed and efficiency over decision-making accuracy.

In other cases, recalibrating these trade-offs is precisely the point of automation. Virginia Eubanks has chronicled Indiana's long experiment with automating the state's welfare enrollment system. In delegating decisions about who is and is not eligible for welfare to machines, trade-offs between accuracy and efficiency played a central (if largely unspoken) role. Prior to automation, human caseworkers helped applicants navigate the complex bureaucracy—often a long, time-consuming process. But with the introduction of IBM's digital enrollment system, such "inefficiencies" were rooted out as "performance metrics designed to speed eligibility determinations created perverse incentives for call center workers to close cases prematurely" [28, p. 50]. As Eubanks tells it, "The goals of the project were consistent throughout the automation experiment: maximize efficiency and eliminate fraud by shifting to a taskbased system and severing caseworker-to-client bonds. They were clearly reflected in contract metrics: response time in the call centers was a key performance indicator; determination accuracy was not" [28, p. 74].

4.3 Displacement

In addition to re-calibrating normative trade-offs, algorithmic speed can displace the costs of automation, making it even more difficult to weigh the risks and benefits automation promises.[13] Consider, again, speed vs. accuracy. In contexts where decisions are subject to review, speeding up decision-making on the front end can simply *move* the costs of error to a back-end appeals process. Viewed from the front, the process thus seems more efficient; but from above one can see that the cost of slow, careful deliberation has simply shifted from one part of the decision-making process to another.

[13] I am grateful to Kiel Brennan-Marquez for pointing this out.

That seems to be what frequently happens with efforts to automate content moderation. Despite claims to the contrary from social media companies, we've seen that NLP and related technologies necessary for machine-driven content analysis have not yet sufficiently matured, and as a result, automated decision-making about online speech is rife with error. For some time, these inaccuracies were simply accepted as the cost of doing content moderation at scale. Eventually, however, users demanded the right to contest decisions, and social media companies were forced to implement complex—and time-consuming—paths for appeal [26]. What looked initially like gains in efficiency were actually, in the end, displacements of time, money, and human attention from one part of the content review system to another.

The same thing can happen in relation to values like privacy, security, and fairness. The field of fair AI/ML has produced countless examples of automation promising to replace expensive and biased human judges with the alleged efficiency, impartiality, and value-neutrality of mathematics, only to deliver flawed, discriminatory algorithms that necessitate costly audits, impact assessments, and technical de-biasing solutions in response. Again, what appears at first to be cost reduction—cheaper, higher quality, automated outputs—is, upon closer inspection, merely cost displacement. Freed for a moment from expensive and unreliable human deciders, organizations deploying automated decision-making technologies are eventually forced to spend significant resources managing unfair and unreliable algorithms, and the costly public relations fiascoes they create.

By displacing costs, algorithmic speed can make it more difficult to understand and evaluate automation's utility. Are the speed and efficiency digital technologies promise worth anything, on balance, if slow, costly, human review will eventually have to be built back into the process anyway? Are they worth it if the technologies end up requiring new, costly quality control measures? Answering these questions is important not only for organizations evaluating automation's value proposition, but also for anyone trying to understand the kinds of normative trade-offs discussed above. When costs are displaced—and therefore made more difficult to detect and apprehend—weighing them becomes more challenging.

4.4 Temporal Fairness

Finally, the above considerations raise important questions about how the costs and benefits of algorithmic speed are distributed. If more speed can mean less accuracy, accountability, legitimacy, and so on, then we ought to ask: are these costs imposed on everyone equally? We might worry, for example, that if the performance of predictive systems generally improves over time, then whoever is subjected to their decisions early on will be more likely to bear the costs of error than those interacting with such systems later. Or we might ask whether some are being subjected to automated decision-making while others continue to benefit from the advantages of human deciders.

At least in some contexts, there is reason to suspect that decision-making speed has already become a means of advantaging some and disadvantaging others. In 2021, the Wall Street Journal revealed that Facebook gave certain users—politicians, celebrities, and other influential public figures—access to a slower, more time-consuming, and deliberative content moderation process than everyone else. The default system, a hybrid of machine and human moderators, produces a significant number of false positives, removing content that (by Facebook's own rules) should be allowed on the platform. Worried about the public relations fall-out from high-profile misfires, the company created an additional layer of review—known internally as "XCheck"—that devotes additional time and human attention to judging the posts of high-status users, before deciding whether to delete them. "For ordinary users, Facebook dispenses a kind of rough justice in assessing whether posts meet the company's rules against bullying, sexual content, hate speech and incitement to violence," the Wall Street Journal reported, while those given access to XCheck are treated more "deferentially": "If Facebook's systems conclude that one of those accounts might have broken its rules, they don't remove the content—at least not right away, the documents indicate. They route the complaint into a separate system, staffed by better-trained, full-time employees, for additional layers of review" [42].

Who is subjected to ex ante vs. ex post enforcement can follow a similar pattern. As we've seen, the instantaneous judgments automation enables can blur the line between prevention and punishment, and the decision to pursue one or the other strategy carries different risks and benefits—among other things, the risk of over- or under-enforcement. As the example above illustrates, choosing instantaneous, functionally ex ante judgment for some people and slower, ex post judgment for others means giving the latter group the benefit of the doubt. Or consider the case of auto lending, which demonstrates even more forcefully how these differences are felt, concretely, by the people involved. Banks routinely install technology in cars (often without the driver's knowledge) that allows for remotely disabling a vehicle the moment a borrower misses a payment. However, all are not subjected to this enforcement mechanism—it is primarily used to police "subprime" borrowers [19]. Those with higher credit scores get the benefit of the doubt if they miss a payment, and—consequently—more time to make good on their loans, before being deprived of access to their cars.

In this way, choosing fast, automated decision-making for some people and slow, human decision-making for others can be a means of marginalization. As Cohen argues, how state institutions treat people's time reflects judgments about their character, political status, and worth. For example, making one person wait longer for access to some right or opportunity—say, voting—than another, similarly situated person, implies that one person's time is more valuable than the other's. Likewise, sentencing two similarly situated people to different prison terms for the same crime implies that they require different degrees of rehabilitation. "[Treating] the time of some citizens as having more value than that of other citizens," Cohen writes, means "disadvantaging those whose time is accorded less or no political value" [15, p. 141]. The same is true in other institutional contexts besides the exercise of state power. By devoting more time, resources, and attention to judging some users' content than others, Facebook implicitly accords them more value. Giving some borrowers time to account for missed payments, while denying other borrowers that same opportunity, signals that some and not others are worth the time.

Moreover, in addition to worrying about the costs of algorithmic speed finding unequal distribution, we should be attentive to who

enjoys its benefits. Automated decision-making promises quicker, more efficient access to rights, goods, and opportunities—e.g., fast-lanes through airport security, shorter visa processing times, instant loan approvals, and so on. Yet even when digital technologies deliver on their promises they tend not to deliver on them equally for everyone.

5 CONCLUSION

The field of technology ethics and policy has itself picked up pace. A new sense of urgency in academia, industry, and the public sphere is driving research into the normative implications of new technologies (especially automated decision-making) to grow and develop at incredible speed. One can understand why speed and its effects are rarely commented on in these discussions: as many who study time and temporality point out, speed is too much a part of our contemporary experience to garner much notice.

We ought to give it more attention. The temporality of decision-making—"decision time"—reflects important norms and values, and as in other areas, incorporating digital technologies into socially significant decision-making processes impacts the temporal norms and values that govern them. The speed that attends automation can modulate the duration, timing, frequency, and lived experience of decision-making, and in doing so it can disrupt prior norms, re-calibrate carefully balanced trade-offs, and displace (and thus disguise) automation's costs. As ever, the risks and benefits of these transformations are often unfairly distributed.

Speed and efficiency can bring enormous advantage. My aim in this paper has been to cast light, not a shadow, on these new temporalities of automated decision-making. But for speed to serve us, individually and collectively, we must understand and thoughtfully guide it.

ACKNOWLEDGMENTS

Special thanks to Kiel Brennan-Marquez and Karen Levy for early discussions that motivated and shaped the ideas in this paper, and to my colleagues at Penn State's Rock Ethics Institute for invaluable feedback. Thanks also to the engaged, thoughtful audiences at 4S, the Society for Philosophy and Technology, and the Monash University Philosophy Department—especially Rob Sparrow—for their generous insights.

REFERENCES

[1] Michelle Alexander. 2018. The Newest Jim Crow. *The New York Times* (Nov 2018). https://www.nytimes.com/2018/11/08/opinion/sunday/criminal-justice-reforms-race-technology.html
[2] Mark Andrejevic, Lina Dencik, and Emiliano Treré. 2020. From pre-emption to slowness: Assessing the contrasting temporalities of data-driven predictive policing. *New Media & Society* 22, 9 (Sept. 2020), 1528–1544. https://doi.org/10.1177/1461444820913565
[3] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine Bias. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing
[4] Henri Bergson. 2001. *Time and Free Will: An Essay on the Immediate Data of Consciousness.* Dover Publications.
[5] Reuben Binns. 2018. Algorithmic Accountability and Public Reason. *Philosophy & Technology* 31, 4 (Dec 2018), 543–556. https://doi.org/10.1007/s13347-017-0263-5
[6] Reuben Binns. 2018. Fairness in Machine Learning: Lessons from Political Philosophy. In *Proceedings of Machine Learning Research.* 11.
[7] Hannah Bloch-Wehba. 2020. Automation in Moderation. *Cornell International Law Journal* 53 (2020). https://papers.ssrn.com/abstract=3521619
[8] Alexis Blue. 2018. Researcher Looks at "Digital Traces" to Help Students. https://news.arizona.edu/story/researcher-looks-digital-traces-help-students
[9] Kiel Brennan-Marquez. 2019. Extremely Broad Laws. *Arizona Law Review* 61 (2019), 26.
[10] Kiel Brennan-Marquez, Karen Levy, and Daniel Susser. 2019. Strange Loops: Apparent Versus Actual Human Involvement in Automated Decision Making. *Berkeley Technology Law Journal* 34 (2019). https://doi.org/10.15779/Z385X25D2W
[11] Simone Browne. 2015. *Dark Matters: On the Surveillance of Blackness.* Duke University Press, Durham.
[12] 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
[13] Danielle Keats Citron and Frank Pasquale. 2014. The Scored Society: Due Process for Automated Predictions. *Washington Law Review* 89 (2014), 34.
[14] Grant Clark. 2018. The Great Firewall of China. *Bloomberg.com* (2018). https://www.bloomberg.com/quicktake/great-firewall-of-china
[15] Elizabeth F. Cohen. 2018. *The Political Value of Time: Citizenship, Duration, and Democratic Justice* (1 ed.). Cambridge University Press. https://doi.org/10.1017/9781108304283
[16] G. A. Cohen. 1980. *The Labor Theory of Value and the Concept of Exploitation.* Princeton University Press, 135–157. https://doi.org/10.1515/9781400853557.135
[17] Julie E. Cohen. 2019. *Between Truth and Power: The Legal Constructions of Informational Capitalism.* Oxford University Press, New York, NY.
[18] A. Feder Cooper, Karen Levy, and Christopher De Sa. 2020. Regulating Accuracy-Efficiency Trade-Offs in Distributed Machine Learning Systems. *SSRN Electronic Journal* (2020). https://doi.org/10.2139/ssrn.3650497
[19] Michael Corkery and Jessica Silver-Greenberg. 2014. Miss a Payment? Good Luck Moving That Car. *New York Times Dealbook* (Sep 2014). https://dealbook.nytimes.com/2014/09/24/miss-a-payment-good-luck-moving-that-car/
[20] Anna L. Cox, Sandy J.J. Gould, Marta E. Cecchinato, Ioanna Iacovides, and Ian Renfree. 2016. Design Frictions for Mindful Interactions: The Case for Microboundaries. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems.* ACM, 1389–1397. https://doi.org/10.1145/2851581.2892410
[21] Paul Cozens and Terence Love. 2015. A Review and Current Status of Crime Prevention through Environmental Design (CPTED). *Journal of Planning Literature* 30, 4 (Nov. 2015), 393–412. https://doi.org/10.1177/0885412215595440
[22] Xin Dai. 2018. Toward a Reputation State: The Social Credit System Project of China. *SSRN Electronic Journal* (2018). https://doi.org/10.2139/ssrn.3193577
[23] John Danaher. 2016. The Threat of Algocracy: Reality, Resistance and Accommodation. *Philosophy & Technology* 29, 3 (Sept. 2016), 245–268. https://doi.org/10.1007/s13347-015-0211-1
[24] Pam Dixon and Robert Gellman. 2014. *The Scoring of America: How Secret Consumer Scores Threaten Your Privacy and Your Future.* http://www.worldprivacyforum.org/wp-content/uploads/2014/04/WPF_Scoring_of_America_April2014_fs.pdf
[25] Natasha Duarte, Emma Llanso, and Anna Loup. 2017. *Mixed Messages: The Limits of Automated Social Media Content Analysis.* https://cdt.org/wp-content/uploads/2017/11/Mixed-Messages-Paper.pdf
[26] Ángel Díaz and Laura Hecht-Felella. 2021. Double Standards in Social Media Content Moderation. *Brennan Center for Justice* (Aug 2021), 40.
[27] Benjamin Eidelson. 2021. Patterned Inequality, Compounding Injustice, and Algorithmic Prediction. *The American Journal of Law and Equality (Forthcoming)* (2021).
[28] Virginia Eubanks. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor.* St. Martin's Press, New York.
[29] Hans-Georg Gadamer. 2004. *Truth and Method* (2nd, rev. ed ed.). Continuum.
[30] Tarleton Gillespie. 2018. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media.* Yale University Press.
[31] Ben Green. 2021. Escaping the Impossibility of Fairness: From Formal to Substantive Algorithmic Fairness. *SSRN* (Dec 2021). https://doi.org/10.2139/ssrn.3883649
[32] Rachel Green. 2019. Finastra and Fico Introduced Two AI and Machine Learning Fraud Prevention Solutions. https://www.businessinsider.com/finastra-fico-launch-artificial-intelligence-solutions-to-stop-fraud-2019-9
[33] Paul Guyer. 2016. Kant, Immanuel (1724–1804). In *Routledge Encyclopedia of Philosophy* (1 ed.). Routledge, London. https://doi.org/10.4324/9780415249126-DB047-1
[34] Robert Hassan. 2003. Network Time and the New Knowledge Epoch. *Time & Society* 12, 2–3 (Mar 2003), 226–241. https://doi.org/10.1177/0961463X030122004
[35] Robert Hassan and Ronald E. Purser (Eds.). 2007. *24/7: Time and Temporality in the Network Society.* Stanford Business Books, Stanford, Calif. OCLC: ocm85783257.
[36] Martin Heidegger. 1962. *Being and Time.* Blackwell Publishing Ltd., Oxford.
[37] Deborah Hellman. 2020. Sex, Causation, and Algorithms: How Equal Protection Prohibits Compounding Prior Injustice. *Wash. UL Rev.* 98 (2020), 481.

[38] Deborah Hellman. 2021. Big Data and Compounding Injustice. *Journal of Moral Philosophy, forthcoming, Virginia Public Law and Legal Theory Research Paper* 2021-27 (2021).

[39] Clément Henin and Daniel Le Métayer. 2021. Beyond Explainability: Justifiability and Contestability of Algorithmic Decision Systems. *AI & Society* (Jul 2021). https://doi.org/10.1007/s00146-021-01251-8

[40] Anna Lauren Hoffmann. 2019. Where Fairness Fails: Data, Algorithms, and the Limits of Antidiscrimination Discourse. *Information, Communication & Society* 22, 7 (June 2019), 900–915. https://doi.org/10.1080/1369118X.2019.1573912

[41] Tatiana Homonoff, Rourke O'Brien, and Abigail B. Sussman. 2021. Does Knowing Your Fico Score Change Financial Behavior? Evidence from a Field Experiment with Student Loan Borrowers. *The Review of Economics and Statistics* 103, 2 (May 2021), 236–250. https://doi.org/10.1162/rest_a_00888

[42] Jeff Horwitz. 2021. Facebook Says Its Rules Apply to All. Company Documents Reveal a Secret Elite That's Exempt. *Wall Street Journal* (Sep 2021). https://www.wsj.com/articles/facebook-files-xcheck-zuckerberg-elite-rules-11631541353

[43] Tatum Hunter. 2021. Ask Help Desk: No, your phone isn't listening to your conversations. Seriously. *Washington Post* (Nov 2021). https://www.washingtonpost.com/technology/2021/11/12/phone-audio-targeting-privacy/

[44] Edmund Husserl. 1992. *The Phenomenology of Internal Time-Consciousness.* University Microfilms International.

[45] Steven J Jackson. 2016. Speed, time, infrastructure: Temporalities of breakdown, maintenance, and repair. *The sociology of speed: Digital, organizational, and social temporalities* (2016), 169–85.

[46] Ian James. 2007. *Paul Virilio.* Routledge, New York, NY.

[47] Daniel Kahneman. 2013. *Thinking, Fast and Slow* (1st pbk. ed ed.). Farrar, Straus and Giroux, New York. OCLC: ocn834531418.

[48] Neal Katyal. 2002. Architecture as Crime Control. *The Yale Law Journal* (2002), 103.

[49] Neal K Katyal. 2001. Criminal Law in Cyberspace. *University of Pennsylvania Law Review* (2001), 113.

[50] Ian R Kerr. 2010. Digital Locks and the Automation of Virtue. *From" Radical Extremism" to" Balanced Copyright": Canadian Copyright and the Digital Agenda* (2010), 247.

[51] Kate Klonick. 2018. The New Governors: The People, Rules, and Processes Governing Online Speech. *Harvard Law Review* 131 (2018), 73.

[52] Will Knight. 2020. The AI Company Helping the Pentagon Assess Disinfo Campaigns. *Wired* (2020). https://www.wired.com/story/ai-helping-pentagon-assess-disinfo-campaigns/

[53] Nicole Kobie. 2019. The Complicated Truth About China's Social Credit System. *Wired UK* (Jul 2019). https://www.wired.co.uk/article/china-social-credit-system-explained

[54] Ava Kofman. 2019. Digital Jail: How Electronic Monitoring Drives Defendants Into Debt — ProPublica. *ProPublica* (Jul 2019). https://www.propublica.org/article/digital-jail-how-electronic-monitoring-drives-defendants-into-debt

[55] Frederick Kunkle. 2019. Bill would require new vehicles to have alcohol detectors to fight drunken driving. *Washington Post* (Oct 2019). https://www.washingtonpost.com/transportation/2019/10/16/bill-would-require-new-vehicles-have-alcohol-detectors-fight-drunken-driving/

[56] Lawrence Lessig. 2006. *Code* (version 2.0 ed.). Basic Books, New York. OCLC: ocm77638613.

[57] Yafit Lev-Aretz. 2011. Second Level Agreements. *Akron L. Rev.* 45 (2011), 137.

[58] Dottie Lux. 2017. Facebook's Hate Speech Policies Censor Marginalized Users. *Wired* (2017). https://www.wired.com/story/facebooks-hate-speech-policies-censor-marginalized-users/

[59] Donald MacKenzie. 2016. Capital's Geodesic: Chicago, New Jersey, and the Material Sociology of Speed. In *The Sociology of Speed: Digital, Organizational, and Social Temporalities.* Oxford University Press, 55.

[60] Sarah Myers West. 2018. Censored, Suspended, Shadowbanned: User Interpretations of Content Moderation on Social Media Platforms. *New Media & Society* 20, 11 (Nov 2018), 4366–4383. https://doi.org/10.1177/1461444818773059

[61] Josh Nathan-Kazis. 2012. Orthodox Rally for a More Kosher Internet. https://forward.com/news/156102/orthodox-rally-for-a-more-kosher-internet/

[62] Sam Nichols. 2018. Your Phone Is Listening and it's Not Paranoia. https://www.vice.com/en/article/wjbzzy/your-phone-is-listening-and-its-not-paranoia

[63] Paul Ohm and Jonathan Frankle. 2018. Desirable Inefficiency. *Florida Law Review* 70 (2018), 63.

[64] Frank Pasquale. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information.* Harvard University Press, Cambridge.

[65] Joel R Reidenberg. 1997. Lex informatica: The formulation of information policy rules through technology. *Tex. L. Rev.* 76 (1997), 553.

[66] Dana Rezazadegan. 2021. Is your phone really listening to your conversations? Well, turns out it doesn't have to. http://theconversation.com/is-your-phone-really-listening-to-your-conversations-well-turns-out-it-doesnt-have-to-162172

[67] Sarah T. Roberts. 2019. *Behind the Screen: Content Moderation in the Shadows of Social Media.* Yale University Press.

[68] Hartmut Rosa. 2003. Social Acceleration: Ethical and Political Consequences of a Desynchronized High-Speed Society. *Constellations* 10, 1 (March 2003), 3–33. https://doi.org/10.1111/1467-8675.00309

[69] Henrik Skaug Saetra. 2020. A Shallow Defence of a Technocracy of Artificial Intelligence: Examining the Political Harms of Algorithmic Governance in the Domain of Government. *Technology in Society* 62 (2020). https://doi.org/10.1016/j.techsoc.2020.101283

[70] Ajith Samuel. 2019. Artificial Intelligence Will Change E-Discovery in the Next Three Years. https://www.lawtechnologytoday.org/2019/04/artificial-intelligence-will-change-e-discovery-in-the-next-three-years/

[71] William E. Scheuerman. 2004. *Liberal Democracy and the Social Acceleration of Time.* Johns Hopkins University Press.

[72] Andrew D Selbst, danah boyd, Sorelle Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2018. Fairness and Abstraction in Sociotechnical Systems. In *FAT* ML '18.* 17.

[73] Bonnie Sheehey. 2018. Algorithmic Paranoia: The Temporal Governmentality of Predictive Policing. *Ethics and Information Technology* (Nov. 2018). https://doi.org/10.1007/s10676-018-9489-x

[74] John E. Smith and The Hegeler Institute. 1969. Time, Times, and the 'Right Time'; Chronos and Kairos:. *Monist* 53, 1 (1969), 1–13. https://doi.org/10.5840/monist196953115

[75] Daniel Susser. 2016. Information Privacy and Social Self-Authorship. *Techné: Research in Philosophy and Technology* 20, 3 (2016), 216–239. https://doi.org/10.5840/techne201671548

[76] 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, 27–44.

[77] Daniel Susser. 2021. Predictive Policing and the Ethics of Preemption. In *The Ethics of Policing: New Perspectives on Law Enforcement*, Ben Jones and Eduardo Mendieta (Eds.). NYU Press.

[78] Omer Tene and Jules Polonetsky. 2013. A Theory of Creepy: Technology, Privacy and Shifting Social Norms. *Yale Journal of Law and Technology* 16 (2013), 59–102.

[79] Tom R. Tyler. 1997. The Psychology of Legitimacy: A Relational Perspective on Voluntary Deference to Authorities. *Personality and Social Psychology Review* 1, 4 (Nov 1997), 323–345. https://doi.org/10.1207/s15327957pspr0104_4

[80] Tom R. Tyler. 2006. *Why People Obey the Law.* Princeton University Press.

[81] Paul Virilio. 2006. *Speed and Politics.* Semiotext(e), Los Angeles, CA.

[82] Judy Wajcman. 2014. *Pressed for Time: The Acceleration of Life in Digital Capitalism.* University of Chicago Press. https://doi.org/10.7208/chicago/9780226196503.001.0001

[83] Judy Wajcman and Nigel Dodd (Eds.). 2017. *The Sociology of Speed: Digital, Organizational, and Social Temporalities* (first edition ed.). Oxford University Press, Oxford, United Kingdom. OCLC: ocn952384327.

[84] Algorithm Watch. 2020. Automating Society Report 2020. (Oct. 2020). https://automatingsociety.algorithmwatch.org/wp-content/uploads/2020/10/Automating-Society-Report-2020.pdf

[85] Lin Yue, Weitong Chen, Xue Li, Wanli Zuo, and Minghao Yin. 2019. A Survey of Sentiment Analysis in Social Media. *Knowledge and Information Systems* 60, 2 (Aug 2019), 617–663. https://doi.org/10.1007/s10115-018-1236-4

[86] Jonathan Zittrain. 2008. *The Future of the Internet and How to Stop It.* Yale University Press, New Haven [Conn.]. OCLC: ocn179832240.