*Original Article*

American Sociological Association

Sociological Theory
1–26

DOI: 10.1177/07352751231186797
st.sagepub.com
Sage

# The Moral Affordances of Construing People as Cases: How Algorithms and the Data They Depend on Obscure Narrative and Noncomparative Justice

**Barbara Kiviat**[1]

**Abstract**

Like many modes of rationalized governance, algorithms depend on rendering people as cases: discrete entities defined by regularized, atemporal attributes. This enables the computation behind the behavioral predictions organizations increasingly use to allocate benefits and burdens. Yet it elides another foundational way of understanding people: as actors in the unfolding narratives of their lives. This has epistemic implications because each cultural form entails a distinct information infrastructure. In this article, I argue that construing people as cases carries consequences for moral reasoning as well because different moral standards require different information. While rendering people as cases affords adjudications of comparative justice, parsing noncomparative justice often necessitates narrative. This explains why people frequently reach for stories that sit beyond the representations of individuals found in records and databases. With this argument, I contribute to the sociology of categorization/classification and draw broader conclusions about modern systems of bureaucratic, computational, and quantitative governance.

**Keywords**

categories, classification, morality, algorithms, allocation

Algorithms are everywhere, with organizations of all sorts now using personal data to predict how people will behave and then treating them differently as a result. Algorithms help decide who receives bail, gets a job interview, warrants a child-welfare investigation, qualifies for a loan, endures an audit, pays a higher price, and much more (Burrell and Fourcade 2021;

[1]Stanford University, Stanford, CA, USA

**Corresponding Author:**
Barbara Kiviat, Department of Sociology, Stanford University, 450 Jane Stanford Way, Stanford, CA 94305, USA.
Email: bkiviat@stanford.edu

Levy, Chasalow, and Riley 2021). The moral dimensions of this increasingly influential mode of allocating benefits and burdens is a key concern among scholars, not to mention real-world observers. Scholars study how particular ideas about right and wrong justify such sociotechnical arrangements (Fourcade and Healy 2017b; Kiviat 2019c), and they question whether algorithmic allocations are just, both directly and implicitly, by bringing to light inequalities by race, sex, and other traits (Eubanks 2017; O'Neil 2016; Zuboff 2019).

There is much here that seems new—novel moral terrain generated by the marriage of unprecedented amounts of data and computational savvy. Yet in this article, I argue that when it comes to how algorithms allocate benefits and burdens, we are also dealing with a much more long-standing, albeit underappreciated, tension. Algorithms extract insight by comparing people to one another, by detecting patterns across individuals. To do so, they necessarily rely on rendering people as cases—as discrete entities with particular attributes (e.g., as a line in a spreadsheet with $n$ columns of characteristics). This way of organizing cognition about people runs counter to a fundamentally different but no less prevalent way of understanding individuals, what they have done, and what they are likely to do in the future: as actors in the unfolding narratives of their lives (Bruner 1986; Heimer 2001). This, I argue, has implications for how people make moral sense of algorithmic allocation because construing people as cases versus actors in narratives makes legible different types of information, and, in turn, affords different forms of moral analysis. This tension has long been with us, but with the rise in algorithmic allocation, the stakes of understanding it have never been higher.

Of course, it is not just algorithms that depend on rendering people as cases. Nation-states, bureaucracies, and other modern forms of social organization also rely on recording people (and other entities to be governed) in standardized, detemporalized, depersonalized form (Lampland and Star 2009; Lipsky 1980:59; Prottas 1979; Scott 1998; Weber 1946). As scholars of categorization, classification, and quantification show, this involves systematically stripping away certain types of details—affect, intention, chronology, physical and social setting, and so forth (Beniger 1986; Bowker and Star 2000; Espeland and Stevens 1998, 2008; Starr 1987). This is not arbitrary, but a feature of large-scale systems built according to formal rationality. "Seeing" people at a distance, in abstract and formal ways, enables many great things, including statistical knowledge (Desrosières 1998; Porter 1994), but it also creates blind spots in understanding how the social world works (Scott 1998). Rendering people as cases has epistemic implications.

The argument of this article is that rendering people as cases—whether to serve bureaucracy, algorithm, or some other end—also funnels our *moral* faculties in particular directions. Taking algorithms as my example, I draw on the concept of affordances, or features of a built or technological environment that steer attention and behavior, to show that the formal aspects of algorithms and the data they rely on shape cognition in a way that makes applying certain moral standards more tractable than others. Specifically, I argue that rendering people as cases to be compared through logical manipulation affords *comparative* notions of justice while obscuring *noncomparative* ones. With people captured as cases, it is easier to figure out if they are being treated equally, say, than it is to evaluate whether each is getting what they deserve. Narrative, by contrast, provides the types of information people often reach for as they adjudicate noncomparative principles—which helps explain why people often insist on telling stories in debates about algorithm fairness (much to the chagrin of algorithm designers).[1]

More broadly, this article contributes to the literature on categorization and classification by demonstrating the utility of more explicitly theorizing the moral ramifications of how we represent people. Numerous scholars have written that categories and classification systems

have moral and ethical bearing on the social worlds they seek to describe, yet most of these observations stop with one-off examples—and nearly all imagine that the main source of moral tension is categorizing people one way versus another. Here, I explicitly connect to an extant philosophical framework, that of comparative and noncomparative justice, to more systematically analyze the moral dynamics at play. Moreover, I show there is moral power not just in the content of categories but also in the cultural form itself.

The rest of this article proceeds as follows. First, I further explain what it means to render people as cases versus as actors in narratives, and I review the scholarship that shows how flattening people into cases has bearing on social understanding (Bruner 1986; Espeland 1993; Heimer 2001; Porter 1994, 1995; Scott 1998). Next, I introduce my argument: that how we construe people also affords different moral evaluations. To do this, I draw on Feinberg's (1974) distinction between comparative and noncomparative justice and the concept of affordances (Davis 2020; Gibson 1979; Norman 2002). Then, I demonstrate how this approach can help us make better sense of when and how different moral sensibilities surface, drawing on two examples: computer scientists' take on algorithmic fairness and public policy debates over credit scoring. Finally, I consider the broader implications of my core theoretical insight: that the rationalization of human life in categorical and computational ways has bearing not only on how we understand the social world but also on how we make moral sense of it.

## CASE VERSUS NARRATIVE: COMPETING WAYS OF SEEING AND MAKING SENSE OF PEOPLE

What does it mean to render a person as a "case"? In short, it means slotting people—amorphous, dynamic, and contradictory beings—into discrete, regularized, and atemporal categories. When a person appears in a census record, or on an intake form, or in an administrative database, that person is represented as an entity with a particular constellation of attributes (Abbott 2001; Dourish 2017; Lipsky 1980:59). A person is construed as a combination of features, as a set of boxes that go either checked or unchecked. Rendering people as cases is thus one way of *organizing cognition* about them (Heimer 2001). It is a way that individuals and organizations culturally represent the irreducibly complex reality of a human life through the processes of categorization and classification with the key task being to document what *kind* of person this is on multiple dimensions (Bowker and Star 2000; Zerubavel 1996).

As scholars of categorization, classification, and quantification show, the construal of people as cases erases idiosyncratic and ambiguous aspects of each person's life and focuses attention on a set of standardized, context-free characteristics, which makes it easier to compare people to one another (Bowker and Star 2000; Espeland and Stevens 1998, 2008; Starr 1987). This proves handy for many forms of modern social organization, including nation-states and bureaucracies (Espeland 1993; Timmermans and Epstein 2010; Weber 1946). The formality of case-based understanding helps authorities see and act on people "at a distance," both conceptually and literally, and make decisions consistently and in high volumes (Gandy 1993; Porter 1994; Scott 1998). Rendering people as cases also enables certain types of knowledge, most notably statistics (Daston 1989; Desrosières 1998; Porter 1995) but also less quantitative forms of expertise. For example, Heimer (2001, 2009) shows that doctors in neonatal intensive care units compare each new infant to the many others they have seen in the past, focusing on the abstract features (e.g., diseased body part, congenital malformation) that experience has taught matter the most (see also, Bowker and Star 2000:192; Desrosières 1998:83; Timmermans 1996).

This has epistemic implications because in construing people as cases, certain information is ignored or thrown away (Beniger 1986; Deleuze 1994; Espeland 1993). As real-life individuals are reimagined as abstractions with traits that can be catalogued and compared and as these abstractions are codified into organizational and institutional information infrastructures (Bowker 2005; Star and Ruhleder 1996), some aspects of lived experience are made more salient while others fade to the background. Subjective experience disappears in favor of that which is externally observable. Dynamism and temporal variation take a back seat to durable and stable characteristics. Ambiguity is expunged. And context of various sorts—from social connections to situational specificity—is for the most part left out. People can still be understood as wildly different from one another, but that difference (and any similarity) is itself homogenized. That is, people are different, but they are different in standardized ways.[2]

How else might we construe people? If not cases, then what? The main alternative is to understand people *narratively*, as actors in the unfolding series of events that make up their lives (Bruner 1986; Orbuch 1997; Somers 1992).[3] Narrative is still a representation of a real-life person, an approach to distilling "the chaos of lived experience" (Somers 1994:613), but as the expansive literature on narrative and storytelling shows, one of a very different sort. Rather than representing a person as a discrete entity with a particular set of attributes, narrative is a cultural form that renders a person in social context (a setting) with meaning arising from how things change over time (the plot) (Bruner 1990; Ricoeur [1981] 2016; White 1980, 1984). Actors themselves are among what changes. People grow up, turn over a new leaf, fall from grace, and so on. Importantly, narrative is well-suited to portraying people with mental states—emotions, intentions, and interpretations of events—and this is often what drives the action forward (Bruner 1986).[4]

Rendering people in narrative ensures a special role for temporality. Stories order things in time, and while they are not alone in doing this—cases can be arranged longitudinally, after all—narrative is distinct in that temporality is hardwired into the cultural form. In narrative, the passage of time is inextricably entangled with character and setting and variation in each, which, taken together, produces meaning and closure (Ewick and Silbey 1995; Ricoeur 1980; White 1980). Narrative does not convey simply what happened but also how we should make sense of it in a way that, say, a timeline (also arranged chronologically) never could (Orbuch 1997; Polletta et al. 2011). Narrative does this in part by revealing how characters think and feel about what is transpiring around them in addition to how they act as a result. Narrative at times also relies on the unsaid. The absence of detail is not "missing data" but often itself meaningful as listeners are moved to fill in the gaps (Steinmetz 1992). This lends narrative rhetorical power, but it is also one way stories can kick up multiple meanings; debates, for example, about what is "really" going on.

The process of rendering a person in narrative is further complicated by the fact that the temporal boundaries of stories are unstable. Wagner-Pacifici (2010) describes historical events as "restless," and we can say the same for events in people's lives. There is no intrinsic point at which to start a story or bring it to a close, and yet those bookends can be deeply consequential for the meaning conveyed (Tavory and Eliasoph 2013; Tavory and Timmermans 2009:249). Knowing what led Jean Valjean to steal the loaf of bread makes us evaluate him quite differently than if we start the tale with the theft itself. Moreover, while temporality in a story may be linear, the relationship between time and meaning is not. Events in the future, even just imagined ones, can have bearing on the meaning of the present and the past (Beckert 2013; Frye 2012; Mische 2009; Tavory and Wagner-Pacifici 2022). Stories are reflexive, and the meanings they produce about the people who populate them are emergent.

Understanding people narratively thus organizes our cognition about them quite differently. While turning a person into a case obscures connections between events, other actors, and people's thoughts and feelings, this information is not only preserved in narrative but also essential in making people coherent and whole (Bruner 1990; Polkinghorne 1988). The information narrative provides about a person is still just a subset of all the possible things that could be known (Ewick and Silbey 1995), but it is a qualitatively different subset than what we get with people construed as cases.

Now, none of this means that case and narrative are mutually exclusive cultural forms. Indeed, people are often rendered as cases within narratives (e.g., a tall, slender woman approaches our protagonist), and narrative helps assign people to the categories that define them as cases (e.g., the man did not mean to kill his neighbor, so he is a manslaughterer, not a murderer). Moreover, the two modes can play complementary roles in social action. Stevens (2007), for example, shows that at the start of college admissions, officers get through large numbers of applicants with case-based thinking (viewing students as particular combinations of test scores, grades, and extracurricular activities), but when it comes time for the final cut, decision-makers turn to narrative, rehashing students' life stories.

Nonetheless, in many situations, one construal is dominant thanks to the broader organizational and institutional environment. Actors operating in bureaucratic settings that rely on protocols, for instance, generally depend on people being construed as cases to enact scripted procedures, whereas those in less routinized settings are freer to understand people as unique individuals living out their own biographies (Heimer 2001, 2009; Lempert and Sanders 1986). The information infrastructures of large-scale modern entities—nation-states, corporations, and the like—record people as cases, which provides a particular perspective on those people, certain cuts of their information "flows" (Deleuze and Guattari 1987). We always have the ability to think about a person categorically or narratively, but often there is a path of least resistance.

## *Enter Algorithms*

That brings us to algorithms. At its most basic, an algorithm is simply a set of rules one follows to carry out an operation. Here, I focus on the sorts of algorithms that have most captured the attention of sociologists—those designed to allocate social and economic benefits and burdens, such as loans, jobs, police surveillance, educational opportunities, fraud investigations, and so on (Brayne and Christin 2021; Burrell and Fourcade 2021; Levy et al. 2021). While algorithms need not be computationally sophisticated, increasingly they are, harnessing techniques such as machine learning to draw on large amounts of personal data to detect patterns in human behavior and extrapolate those patterns into the future. These predictions then help people in positions of authority—loan officers, employers, police officers, admissions officers, fraud investigators—make decisions about who gets what.

For algorithms to work, people *must* be construed as cases. To an algorithm, a person is an entity with a set of features, and insight comes from comparing people with different sets of features to one another. In any given sample (or any sample useful to an algorithm), some people exhibit a behavior, and others do not. Some repay loans, and some default. Some succeed in college, and others drop out. An algorithm analyzes which patterns of characteristics have gone with various outcomes in the past and then, when new people come along, presumes their patterns will correlate with outcomes in the same way. A lender may, for example, use machine learning to identify a relationship between people moving money in and out of bank accounts and the likelihood of them paying their credit card bills on time. When courting a new customer, the lender examines a person's bank deposits and withdrawals to

see if the pattern is similar to those historically associated with nonpayment and if so, charges a higher interest rate or decides not to issue a card at all.

Behind all algorithms sit databases, which bring certain material constraints. Databases come in different architectural varieties—they may, for instance, be structured around hierarchies, networks, or attributes—but they all nonetheless depend on information rendered categorically (Bowker 2005; Couldry and Hepp 2016:132; Gillespie 2014). In the most commonly used "relational" databases, information is stored in sets of tables, with each column capturing a different type of information (e.g., name, age, occupation) and each row encoding a "relationship" between values for a particular "object" (e.g., person). Other types of databases use different organizational schemes, but the scheme itself is constant—common across the entire database. As Dourish (2017:108) explains, "arranging data into a common format makes each data item amenable to a common set of operations." Data items can be sorted, collated, grouped, summarized, and—crucially for the argument I present here—compared.

Indeed, the insights algorithms provide to decision-makers always depend on *comparison across cases*. How a person is predicted to behave is a function of how other, somehow similar and somehow different, people have behaved in the past (Hacking 1990; Schauer 2003). Even as data become more fine-grained and individualized, the mechanics of algorithms still demand decontextualization and standardization so that information flows from one person can be easily compared to those from the next (Guseva and Rona-Tas 2001). This continues to be true when algorithms themselves determine which sorts of information about people to attend to. So-called unsupervised methods still depend on the processes of categorization and classification even if those processes become opaque to their human designers (Fourcade and Healy 2017a).[5]

This is radically different from how one often uses narrative to understand a person and what the person is likely to do in the future. To make sense of cases, we gravitate toward comparing them, but with narrative, we can easily size up just a single person because that person moves through time and particular situations, demonstrating change and trajectory. For example, in Heimer's (2001, 2009) study of the neonatal ward, parents of ill children had typically never seen such a sick infant; they had no store of cases to draw on, as doctors did. Yet parents watched their own child hour by hour, day by day, witnessing incremental change and the circumstances that went with it. Their knowledge was intensive rather than extensive, and for that reason, they at times had insights about their child's health that doctors did not.

Now, that does not mean we *cannot* compare narratives. Indeed, we do all the time. Whose story is more believable, hers or his? How did this person succeed while that one failed?[6] What rendering people as cases uniquely enables is comparison *at scale*—figuring out similarity and difference not between two or three people but between dozens or hundreds or thousands or more. Comparing a small number of narratives is tractable because the basis of comparison can itself be emergent and reflexive, pivoting along with each story's narrative arc and the meaning it produces. Yet this is exactly what gums up the works if we try to compare stories many at a time.

Importantly, narrative is not synonymous with singularity. At times, we care about the stories of individuals' lives precisely because they capture the experiences of other people in similar positions (Ewick and Silbey 1995; Polletta 1998; Steinmetz 1992). Protagonists can have proper names, or they can be archetypes (e.g., a single mother, a conservative voter, a Silicon Valley executive). Narrative thus works at two levels. Stories can be about specific individuals, or they can invoke cultural tropes about particular types of individuals.[7] Rather than a lack of generality, what makes understanding people in narrative distinct from

understanding them as cases is the way that narrative bakes in chronology and context and characters who must grapple with both—with meaning emerging from this entanglement. These features then carry important consequences.

The temporal flow of narrative, in which one event leads to the next, means stories are well-suited to exhibiting *why* things happen (Bruner 1991; White 1980). As Polkinghorne (1988:21) writes, if we want to understand a man buying life insurance and we have rendered him as a case, our explanation may go something like this: "because he is a [W]hite male, in the 40-to-50 age category, and those in this category are, in 70 percent of cases, also in the category of people who buy life insurance." The explanation is unsatisfyingly sparse, telling us precious little about *how* one state of being translates into another. It gives no hint—as narrative would—of what has transpired in the man's life (becoming a father, witnessing a colleague die suddenly) or motivated him to action (thinking of himself as a provider, receiving an insurance solicitation in the mail).

That does not mean people rendered as cases cannot be arranged temporally. Narrative, as discussed earlier, does not have a unique claim on one thing happening after the next. Yet to arrange cases longitudinally, they in fact must *be arranged*. As Manovich (1999:85) writes, a "database represents the world as a list of items which it refuses to order." Someone can come along and impose a temporal order on those items, but such sequencing—and the meaning that follows—comes from without, not within.

Moreover, because narrative often reveals (or at least suggests) a protagonist's inner mental life, including emotions and intentions, it helps assign responsibility, blame, and praise (Bruner 1986; Harding et al. 2017; Stone 1989; Wagner-Pacifici 2010). In the example of the man purchasing life insurance, we draw different conclusions when the buyer is worried about his family than when he is contemplating how he might fake his own death. Indeed, even small changes to the interplay between narrative arc, characters' actions, and the surrounding situation can lead to very different judgments about intentionality and control and, in turn, how we feel about and understand the person at the center of the story (Cerulo and Ruane 2014; Monroe and Malle 2017; Scott and Lyman 1968).

All of which is to say, construing people as cases and analyzing them through case comparison allows for a different sort of understanding than does making sense of people as actors in a narrative. Each mode of organizing cognition foregrounds different types of information, making people legible in different ways and enabling different styles of reasoning. While not all scholars of categorization and classification directly engage with narrative as the epistemic "other," they nonetheless tend to describe the processes they study as ones that render people as entities that are somehow flat, thin, stripped down, or otherwise incomplete—albeit incredibly well-suited to modern forms of bureaucratic, quantitative, and computational governance.

### *The Missing Link to Moral Evaluation*

What I argue has gone less theorized and that the comparison to narrative puts into sharp relief are the *moral* ramifications of rendering people as cases. That is not to say the literature on categorization and classification ignores the issue. In one sense, it is rife with moral matter. When powerful institutions, such as states or corporations, categorize people in some but not other ways, they make certain social experiences visible while obscuring others (Bowker and Star 2000; Lampland and Star 2009; Lengwiler 2009; Minow 1990; Starr 1992). This then patterns access to material and symbolic resources, at times reinforcing patterns of historical (dis)advantage, hence the high stakes of battles over which categories ought to be used (Bourdieu 1984; Starr 1992; Tilly 1998).[8] Categories can also themselves

be value-laden, or "categories of worth," whether they appear in official classifications or underpin broader cultural understandings of good and bad (Durkheim 1965; Lamont 2000; Mohr 2005; Zerubavel 1997). The moral valence of such categories makes them powerful in motivating social action—consider the politics of the "deserving" and "undeserving" poor (Somers and Block 2005; Steensland 2010).

That said, nearly all the scholarly focus is on how different values and moral understandings are embodied by one classification system versus another. There is much less exploration of what it morally means to categorize people in the first place, to construe them as entities that are properly differentiated and defined by their attributes (whatever those may be). The attention is on understanding what it means to render people as this sort of case or that sort of case—not on understanding the moral implications of the cultural form itself, of transforming a person who is not a case into one who is.

On this count, the quantification literature is more developed. Espeland and Stevens (2008:431–32), for example, talk at length about the "ethics of numbers," including the tendency of quantifying processes to "narrow our appraisal of value and relevance to what can be measured easily, at the expense of other ways of knowing." It is not simply that *the way* we quantify embodies particular values, although it does (Fourcade 2011; Salais 2016). It is also that the reduction of difference to that which can be unambiguously counted and otherwise mathematically manipulated is, in and of itself, a judgment about how certain types of (incommensurable) features do not matter (Espeland 1998; Espeland and Stevens 1998).

Yet quantification is not the same thing as categorization. Before we can count things up, we first must fit them into categories—categorization logically and technically precedes quantification (Fourcade 2016; Hacking 1986; Hirschman, Berrey, and Rose-Greenland 2016). This means we cannot necessarily assume lessons learned from the study of quantification apply to categorization because quantification is the subsequent process. The question of what moral stakes come with categorization and classification (i.e., the processes that render people as cases) as such remains largely unexplored.

## COMPARATIVE AND NONCOMPARATIVE JUSTICE AND WHAT IT TAKES TO ADJUDICATE EACH

In this section, I turn to an extant philosophical framework to more systematically theorize the moral implications of construing and acting on people as cases. There are multiple ways to divide up the universe of possible justice concerns, but one division that provides useful scaffolding for the argument offered here is the distinction between what Feinberg (1973, 1974) calls comparative and noncomparative justice.[9] "In all cases," Feinberg (1974:298) writes, "justice consists in giving a person his due, but in some cases one's due is determined independently of that of other people, while in other cases, a person's due is determinable only by reference to his relations to other persons." The former is noncomparative; the latter is comparative.

In other words, to judge whether some entity (algorithm or otherwise) allocates benefits and burdens fairly, we can look to see how people are treated relative to one another, or we can look to see if each person individually is held to the right standard irrespective of how anyone else fares. A concern about whether people are treated equally, for instance, speaks to comparative justice, whereas a concern about whether a meritorious person gets what she deserves speaks to noncomparative justice (Feldman 2016; Miller 2021). Importantly, comparative justice concerns can be about either how individuals compare or how *groups* of individuals compare (i.e., how individuals within those groups stack up to one another, on average). Moreover, comparative justice concerns can be about equality or proportionality:

It may not be necessary that two people or two groups of people be treated exactly the same way as long as the amount of advantage or disadvantage each receives sits in the right relation to the other (e.g., is more or less).

In many settings, people care about both comparative and noncomparative criteria. When it comes to criminal justice, for example, we generally want the guilty to be convicted and the innocent to go free. These are noncomparative standards in that we attend to each person independently—that is, to each person's guilt or innocence—to determine if the standards have been met or violated. At the same time, we tend to grow concerned when among the guilty, certain groups receive longer sentences than others for having committed the exact same crimes (as is the case with Black men in the United States). Here, the moral breach is comparative; the problem becomes visible only once we know how people are treated relative to one another. Noncomparative justice requires that individuals be treated in line with their rights and deserts, whereas comparative justice requires that people who are similar be treated alike and those who are dissimilar be treated differently (Feinberg 1973, 1974; Montague 1980).

There are many ways to fill in the details of both comparative and noncomparative justice. On the comparative side, philosophers and social theorists have been debating for centuries which sorts of equalities and proportionalities are the right ones for various situations (e.g., Jencks 1988; Rae 1979; Williams 1964). If justice requires that similar people be treated alike, then one must answer the very tricky question of alike *in which ways*. In other words, which forms of similarity and difference are the morally relevant ones (Boxill 1992; Hellman 1997; Minow 1990; Shelby 2016). On the noncomparative side, things are not any easier. If justice means allocating benefits and burdens according to appropriate rules and making sure people are held to account for the right things, then one must figure out what makes rules just and which standards go with which situations—for example, when we should allocate according to need versus merit versus right, and so on (Eckhoff 1974; Lucas 1980).

My purpose here is not to engage in such debates but, rather, to show that no matter the details, there is a key difference between comparative and noncomparative justice in that evaluating the success of each relies on a distinct information infrastructure. We need different sorts of information about people to adjudicate whether standards of comparative and noncomparative justice have been met.[10]

To assess comparative justice, we need to know about multiple people or groups of people. But we cannot know about them in just any way. We need to know them such that they can be systematically compared and at scale. Similarity and difference cannot be idiosyncratic but must, to a large degree, be standardized. Reductionist and unnuanced as racial classification may be, identifying criminal defendants as Black or White makes comparative injustice visible. To assess noncomparative justice, on the other hand, we need to know about only one person. And the sorts of things we need to know are often quite different. Here, nuance and specificity are not hindrances, but telling details. To adjudicate guilt or innocence requires understanding not only what happened but also why and what the defendant was aware of and thinking throughout—whether he planned the attack, knew the gun was loaded, feared for his life, and so on. The implication, which I expand on below, is that rendering people as cases facilitates adjudications of comparative justice, and rendering them in narrative facilitates adjudications of noncomparative justice.

A strong version of the argument, then, would be that when we render people as cases, we make moral evaluations comparatively, and when we render people in narrative, we make moral evaluations noncomparatively. Yet my actual argument is less deterministic than that; I am not saying cases are inherently comparable and narratives are not. In a moment, I return

to the key issue of how the different sorts of information provided by case and narrative facilitate different modes of moral evaluation. But first, I take a brief detour to better explain the nature of that facilitation.

## *Affordances and the Work They Do*

I posit that rendering a person as a case or in narrative *affords* different forms of moral reasoning. Drawing on ecological psychology and design studies, I define affordances as features of an environment or object that steer attention and behavior. A chair, with its horizontal platform at knee level, affords sitting. A door with a handle that juts out in the shape of a sideways "U" affords pulling, whereas one with a plate that sits flush against the glass affords pushing. One certainly *can* stand on a chair (or throw it across a room) or push a U-shaped handle forward, but these are the less likely acts.

Working on problems of visual perception, Gibson (1979) was the first to popularize the concept of affordance as a way to link an animal's environment to its actions while downplaying the intermediary step of mental representation.[11] Norman (2002) later picked up the term, applying it to everyday objects and technologies in part to explain why people often instinctually use such items incorrectly—because their designers have created them with the wrong affordances. Affordance means slightly different things in each tradition (Chong and Proctor 2020; Davis 2020), but what is consistent is a particular unthinking quality. Through affordances, the world around us guides us to certain actions even in the absence of deliberate reasoning and active decision-making.

Affordances thus have real-world effects. In one example, Norman (2002:9) notes that when British transportation authorities built bus shelters out of glass, vandals smashed them, and when they built shelters out of wood, vandals drew on them with graffiti. Those setting out to commit crime did so in different ways because each material invited a different set of actions. Yet, importantly, affordances are nondeterministic. In fact, the ability of an affordance to shape behavior is often culturally and institutionally mediated (Davis and Chouinard 2016). Door handles afford opening doors, but in certain times and places, women still do not, waiting instead for the man to do so.

This, I argue, is also the nature of the relationship between cultural forms for rendering people (as cases or actors in narrative) and types of moral reasoning. Systems of modern governance—including but not limited to algorithms—depend on capturing people in reduced, rationalized, unambiguous formats, such as in rows of spreadsheets with *n* columns of characteristics. This, in effect, is a design decision, which makes it easier to think in some directions than in others.[12] This is akin to what information scholars Dourish and Mazmanian (2013:101) call "the material consequences of representational practices." As they write, "the particular forms that information takes…shape the questions that can easily be asked of it, the kinds of manipulations and analyses it supports, and how it can be used to understand the world" (Dourish and Mazmanian 2013:100).

When people are rendered in a way that affords comparison, decision-makers are more likely to compare them for the purposes of knowing things about them—and, I add, for normative evaluation. This does not mean that standards of comparative justice are the only ones applied, just that they are the more salient ones, the ones people more easily reach for. Affordance is not fate, and there is still plenty of room, even with people construed as cases, for concerns about noncomparative justice to surface. But the information infrastructure itself does not facilitate this.

## *Moral Reasoning and the Cultural Form People Take*

To return, then, to the main issue: How does rendering people as cases afford one mode of moral reasoning while posing informational difficulties for another?

Rendering people as cases reduces the information available about them, but in a specific way. Through categorization and classification, people retain only attributes that are common across individuals. When a bureaucratic worker logs clients into a database, the fields may take on different values, but the fields themselves do not change. A person may be tagged as married, single, or divorced, but marital status as a category of knowing is held in common. People are still, ontologically speaking, idiosyncratic, but in decision-making systems that capture them as cases, axes of similarity and difference are standardized. This homogenization makes it easy to compare people to one another—including when adjudicating whether standards of comparative justice, such as equality and proportionality, have been met. If we want to know whether men and women are treated equally, it helps to have each person flagged as either a man or a woman. If we want to know whether low- and moderate-income taxpayers are audited at higher rates than rich ones, it is useful to have individuals assigned to discrete income brackets.

Trying to get a handle on a comparative justice issue like social equality with people rendered in narrative would be much trickier. The modularity of cases paves the way for comparison at scale, but the entanglement of narrative's component parts (plot, character, setting) complicates matters, as does how a story produces meaning—in emergent, reflexive, and nonlinear ways. A story is more than the sum of its parts (even omitted information can be critical), and it functions as an integrated whole. As Somers (1994:616) writes, unlike categorical thinking, "narrativity precludes sense-making of a singular isolated phenomenon" (see also Ricoeur 1980). To reduce a story to a set of standardized dimensions along which it might be compared with others would be to destroy the essence of narrative. Moreover, narrative's complicated relationship with temporality undermines *simultaneous* understanding (Wagner-Pacifici 2010), which is key for comparison.

By contrast, for many standards of noncomparative justice, such as desert and merit, understanding people as cases informationally falls short. Desert, for example, is a standard by which people rightly receive benefits or burdens according to the moral worth of their past actions (Feinberg 1970; Feldman and Skow 2019; Rachels 1991; Sher 1987). Students who master the material deserve good grades, and people who commit violent acts deserve prison sentences. Adjudicating whether such standards have been properly applied often requires knowing the sorts of things case-based renderings are bad at conveying, such as what led to observable behavior. Students who ace an exam because they stole the answer key actually do not deserve good grades, and people who hurt others in desperate acts of self-defense (arguably) do not deserve prison sentences, or at least not long ones. Construing people as cases strips away information about intention and motivation, context about the constraints people face, and connections between disparate events—all of which is often necessary for determining what people can be fairly held accountable for in a noncomparative sense.

This type of information is much more readily supplied by narrative, given the way plot, character, and setting are hardwired into the cultural form. Indeed, when faced with people construed as cases (e.g., as they exist in administrative documents and other records), decision-makers often actively reach beyond the written page and turn to narrative to fill in informational gaps—whether that means soliciting stories from the individuals they are evaluating or simply imagining what those stories are likely to be (Espeland 2015; Lynch 2017). Through narrative, it is easier to see and appreciate the role of social structure, differences in subjective understanding, and life trajectory—what in moral evaluation adds up to consideration for extenuating circumstances, personal redemption, and so on.

Key to this process is how narrative conveys both how *and why* things happen. In adjudicating noncomparative justice, an important question is often how much of a role a person has had in bringing about events, whether the person did so intentionally, and if the person

could have acted otherwise (Lempert and Sanders 1986). Narrative excels at providing thick causal understanding, the sort that refers back to both proximate and ultimate causes, and integrates not only observable events but also people's thoughts and feelings—stories come with a *point of view*. Indeed, some scholars of narrative argue that the establishment of cause and effect is inherent to the form, part of what makes narrative different from other kinds of temporal ordering and what gives stories a sense of meaning and closure (Polkinghorne 1988; Polletta et al. 2011; White 1980; see also, Knight and Reed 2019:248). In this way, narrative is well-suited to the assignment of responsibility, not just in a mechanical sense but also a moral one.

Saying that rendering people in narrative affords adjudications of noncomparative justice is not to say that every noncomparative situation requires narrative's fullest trappings. Some noncomparative standards, such as rights, may depend on only the sparsest of biographical detail, such as having been born in a particular country. Again, we are in the world of affordance, not one-to-one correspondence. That said, in practice, narrative is often what helps decision-makers get noncomparative justice exactly right rather than only approximately so. Consider, for example, the students above who do well on an exam. At first pass, understanding the students as decontextualized, detemporalized cases may appear to work just fine. All we need to know is that students who fall into the category "all answers correct" deserve an A. But this is not true for the cheaters. The stories of how people wind up in the categories they do often reveal moral distinctions that classification obfuscates (Kiviat 2021). Indeed, this is where the instability of narrative's temporal boundaries can come in handy. Starting the story a bit earlier or letting it play a bit longer can cast important new light on the person in question.

None of this is to say that narrative is a better tool for adjudicating justice claims overall. For identifying comparative (in)justice, construing people as cases clearly works better. Indeed, narrative's relative informational inattention to points of comparison across different people can mean that through storytelling, we reintroduce bias. Enmeshed in narrative particulars, it can be hard to see the forest for the trees. Moreover, narrative's showcasing of mental states, including motive, can help decision-makers claim that what seems to be a comparative injustice actually is not (e.g., an employer did not intentionally discriminate, and therefore, the resulting allocation of jobs is not problematic) (Lempert and Sanders 1986:390). The goal is not to dismiss thinking about people as cases but, rather, to show that for a full normative evaluation of the allocation of benefits and burdens, we need both cultural forms.

## EXAMPLES OF THE RELATION BETWEEN ALGORITHMIC ALLOCATION AND MORAL EVALUATION

In this section, I present two examples to demonstrate how construing people as cases affords adjudications of comparative justice and how when social actors undertaking moral evaluation want to bring in noncomparative justice concerns, they turn to narrative. This entails going beyond understanding people as they are rendered in administrative records and databases—where they exist as cases—and reembedding them in the situational and temporal context of their unfolding lives.

Both examples involve groups of professionals who explicitly set out to determine whether algorithms are allocating benefits and burdens fairly. In the first example, computer scientists and other scholars tackle this question by contributing to the burgeoning literature on "algorithmic fairness," which touches on many topics but frequently returns to algorithms used in criminal justice and lending. The audience for this literature is not only fellow academics but

also organizations that use algorithms and want them to be understood as fair. In the second example, policymakers and adjacent stakeholders publicly debate the fairness of credit scores and the algorithms that produce them. At times, judgments of unfairness result in legal restrictions. In both examples, actors are preoccupied with notions of distributive justice—for example, whether loans, bail decisions, and so on are allocated fairly. This is a scope condition for what follows. Distributive justice includes both comparative and noncomparative standards (equality is an example of the first, merit of the second), but other dimensions of moral concern that surface in broader conversations about data use and algorithms (e.g., privacy) are peripheral to these examples.

In each milieu, actors undertake similar projects of moral interrogation, yet there are important differences. Key among these is that the normative back-and-forth that takes shape in the algorithmic fairness literature belongs to a social world which is deeply aligned with the production of algorithms and, as a consequence, is committed to construing people as cases. I elaborate on this in the following section and suggest how it relates to the fact that in this example, considerations of comparative justice dominate. Then, I turn to public policy debates about the fairness of credit scoring algorithms. In this example, the moral back-and-forth transpires in a more generalized social world, where actors external to the world of algorithm construction repeatedly raise concerns about noncomparative justice. Crucially, they do so by reaching beyond the data to invoke narratives of people's lives.

Taken together, these examples illustrate two things. First, when people are rendered as cases, it is possible to have a robust debate about fairness while failing to see that not all standards of distributive justice are comparative in nature. Second, when social actors do introduce noncomparative standards to such debates, their mode of doing so is by telling stories that necessarily depend on information that sits outside of records and databases.

## *Theorizing Algorithmic Fairness*

As algorithmic decision-making has spread far and wide, so have ideas about what it means for algorithms to allocate fairly. Organizations increasingly use computational predictions about how people will behave to assign benefits and burdens. These practices and their rhetorical defenses have been shaped by academic ideas about how to identify injustice. There is now a large body of scholarship on "algorithmic fairness," which chiefly sits in computer science but also features the work of legal experts and philosophers. Here, I take this scholarship as an empirical phenomenon of interest in its own right—as an instance of a group of people embarking on a task of moral evaluation.[13]

The algorithmic fairness literature largely arises from the discipline of computer science, and it bears that field's hallmarks. The goal is generally to find a solution to a problem. Each member in a population experiences an outcome (e.g., a loan is granted, bail is denied) and the task is to figure out if by a particular definition of fairness, those outcomes have been allocated fairly. In approaching this problem, computer scientists use their usual tools and methods, including abstraction, formal logic, and calculation. Computer scientists inhabit an intellectual world where the information infrastructure must accommodate computation (Dourish and Mazmanian 2013; Malik 2020; Selbst et al. 2019:59). This means people necessarily get rendered as cases.

In this setting, conceptions of comparative fairness dominate. The nearly universal starting point is the Aristotelian notion of treating like cases alike. In one influential article, Dwork et al. (2012:215) write: "We capture fairness by the principle that any two individuals who are similar with respect to a particular task should be classified similarly." In another article, Friedler, Scheidegger, and Venkatasubramanian (2016:2) point to the field's "guiding

informal understanding of fairness" as the precept that "similar people should be treated similarly." Canvassing dozens of fair machine learning experts in academia, industry, and government, Chouldechova and Roth (2018) identify two main types of fairness definitions: those that require parity between individuals and those that require parity between groups (e.g., that members of a protected class fare no worse than others, on average). In all instances, equivalence is key. Fairness is defined by how some person or group of people is treated relative to another. The definitions on offer unilaterally reflect comparative justice.

Applying these principles is no easy task, and there is substantial debate about *which* equivalencies are the important ones. A survey of leading machine learning outlets returned 20 prominent standards, including statistical parity (equal outcomes across groups), predictive parity (equal chances of being correctly assigned to the good outcome), and balance for false negatives (equal chances of being incorrectly assigned to the bad outcome) (Verma and Rubin 2018; see also, Berk et al. 2021; Mehrabi et al. 2022; Mitchell et al. 2021).[14] In practice, the stakes can be high. In one often-cited case, the makers of a bail-setting algorithm touted their product as racially unbiased because Black and White defendants showed equivalence on one outcome (i.e., one definition of fairness)—even though, as others later showed, Black defendants fared unequally in different ways (Angwin et al. 2016; Dieterich, Mendoza, and Brennan 2016). There is thus great dynamism in the algorithmic fairness literature and all the more so given that different metrics often cannot be satisfied simultaneously (Corbett-Davies and Goel 2018; Hellman 2020; Kleinberg, Mullainathan, and Raghavan 2017).

And yet, much is missing. Notions of noncomparative justice—of ensuring each individual gets what they deserve or have a right to—are practically nonexistent. If judges send all jaywalkers to prison for 10 years, then they are treating everyone equally and yet arguably also unfairly given how greatly the punishment outweighs the crime. The algorithmic fairness literature has no metric to detect such an injustice. Nor does it have a way to articulate what might be unfair about an algorithm that uses information on preexisting medical conditions to allocate health insurance, one that draws on social media data to determine who should get a loan, or one that uses occupation and education to price car insurance—all contentious practices, commonly called "unfair" in the broader public sphere. Algorithmic fairness, as currently theorized, is thus blind to what Fazelpour and Lipton (2020:8) call "the non-comparative harm of making decisions on irrelevant grounds."[15]

Moreover, the literature lacks a way to parse which predictions are themselves unjust. Consider, for example, universities that allocate scholarship money based on algorithmic predictions about the minimum amount of aid it will take to entice a particular student to enroll (Aulck, Nambi, and West 2019; Engler 2021). Were we to intuit that price sensitivity is an unfair standard for allocating resources in the domain of higher education, there would be no way to formalize this in the algorithmic fairness literature. Nor would there be a way to articulate that a prediction *of any sort* may be an unjust way to distribute a benefit or burden. In linking allocations to behavior, decision-makers can look to the past or to the future; to evidence of how people have behaved or to predictions about how they are likely to do so going forward (Jencks 1988). Many noncomparative standards of justice, such as merit and desert, rely on how people have behaved (e.g., whether they worked hard, acted with intent, achieved high marks), while beliefs about *future* behavior are irrelevant (Underwood 1979). In the algorithmic fairness literature, this goes entirely uncommented on.

So, why? Why exclude noncomparative justice concerns in favor of comparative ones? I propose that the answer, or at least an important part of the answer, is precisely because people have been construed as cases.[16] For algorithms to work mechanically, it is necessary to render the people they are designed to act on as cases. Scholars seeking to build tools to test the fairness of how these algorithms work therefore begin with this view of people—as

database-bound objects with attributes that can be sorted and slotted according to similarity and difference (Dourish 2017). This information infrastructure then affords comparison both analytically *and morally*. Scholars of algorithmic fairness solve the problems they are able to with the tools they have, and the tools they have make comparative forms of justice the more tractable ones. The "material consequences of representational practices" are not only epistemological in nature (Dourish and Mazmanian 2013:101) but normative as well.

This is not to say that computer scientists and their colleagues do not have intuitions about noncomparative (in)justice. It is possible to think outside the form in which data arrive. Indeed, in the next section, I show an example of just that—of a group that deliberately reaches beyond records and databases to get at noncomparative standards such as desert. The fact that the algorithmic fairness literature does not do this or even comment on its omission is telling. Davis and Chouinard (2016) underscore that affordances are culturally and institutionally mediated, and most scholars of algorithmic fairness inhabit an intellectual world where the reduction of complex social phenomena, such as lived human experience, to discrete classification and formal logic is taken for granted. It is easy to be blind to standards of noncomparative justice when one is entrenched in institutions dependent on understanding people as cases and comfortable with the assumption that case-based renderings are sufficient representations. Affordance, as noted previously, is not fate, but it may be more so in the hands of some users.

Next I turn to a different domain, that of public policy, where the reduction of social phenomena to cases is not taken for granted and where questions about the fairness of algorithmic allocation at times take a very different form.

## *Debating Credit Scoring Algorithms*

Scholars often hold up credit scores as the paradigmatic algorithmic prediction (e.g., Fourcade and Healy 2017b; Pasquale 2015). In their classic iteration, credit scores are constructed from information in consumer credit files—records that construe people as cases defined by their debts and other money-related attributes—and are designed to predict the likelihood that a person will repay borrowed funds. Today, organizations besides lenders, such as insurance companies and landlords, also consult credit scores to make allocative decisions (e.g., what premium to charge, whether to take on a person as a tenant), and scoring algorithms increasingly include data from beyond credit files (Consumer Financial Protection Bureau 2012). Regulators, lawmakers, and other public policy actors have watched these developments with interest, especially with the progression from easy-to-understand scorecards to more computationally sophisticated and opaque algorithms. The fairness of credit scoring has been debated in public policy venues on countless occasions (Kiviat 2019b; Krippner 2017; Krippner and Hirschman 2022).

In these debates, there is plenty of concern for comparative justice. Policymakers, consumer advocates, and others routinely raise the issue that even though credit scores do not directly rely on attributes such as sex, race, ethnicity, and geography, using such scores in decisions about lending, insurance, and housing might nonetheless lead to problematic differences across groups defined by these and other socially meaningful traits.[17] Arguments about disparate impact, proxy discrimination, and so on rest on comparisons between people understood to be different sorts of cases (e.g., men vs. women). The comparative fairness questions that preoccupy the algorithmic fairness literature reflect real-world concerns, and the materiality of credit scoring algorithms and the credit bureau databases they depend on lend themselves to answering those questions. With borrowers rendered as cases, it is possible to compare millions of them at a time, revealing patterned social inequalities. Credit records and algorithms may be proprietary, but credit bureaus have shared enough information over the years for policy actors to raise alarm at such differences.

Yet public policy debates about credit scores also revolve around noncomparative fairness. This often takes the form of contestation over the sorts of information that can rightly be fed into algorithms—that is, the aspects of a person's life that organizations should use in shaping the distribution of socially important resources. Consider, for example, long-standing pushback—from state attorney generals, the Consumer Financial Protection Bureau, and others—against including medical debt in credit scoring algorithms (Best, Bosco, and Wu 2019; Carrns 2022; Consumer Financial Protection Bureau 2022). Here, the issue is noncomparative in that policymakers question whether having (unpaid) medical debt is a legitimate standard for making a decision about what sort of credit, insurance, or housing a person receives. Irrespective of how it may help predict future behavior, medical debt is seen as inappropriate grounds for making such decisions—a perfect example of the sort of "noncomparative harm" Fazelpour and Lipton (2020:8) find missing in the algorithmic fairness literature.

Policymakers invariably express this by reaching beyond credit records and bringing narrative into the debate. In explanations across domains, proponents of eliminating medical debt from credit scores focus on how this debt is rarely incurred intentionally. A person becomes ill or has an accident and faces no choice but to receive high-priced medical treatment—in some situations, such as being rushed to the hospital, a person may not even be conscious. Arguments against using medical debt also highlight the role of social structure (e.g., a lack of universal health care) and other actors, such as insurance companies that are responsible for medical costs but may fail to pay bills fast enough to keep bad debt off credit reports. In other words, *how* a person comes to be represented as a particular sort of case (e.g., one with delinquent medical debt) is key to understanding the moral valence of being categorized that way. Notably, none of this relies on making comparisons between individuals or groups. The claim of injustice arises from a person—even just one, single person—being held accountable for something that is not their fault.[18]

Narrative works so well in articulating such noncomparative qualms because narrative captures how events have unfolded (or may have unfolded or may unfold) in a person's life. As a cultural form, narrative foregrounds factors such as a person's agency in making events happen and the role of external forces and actors, which can be key details in deciding whether a person deserves the treatment they receive (Feinberg 1970; Rachels 1991; Sher 1987). Moreover, the ambiguity of narrative's temporal boundaries allows us to peer into both the past and the future. Credit bureau records can be arranged chronologically, to be sure, but doing so is an exercise in sequencing, not sensemaking. Only with narrative's entanglement of character, setting, and plot does meaning emerge.

Indeed, stories about what gives rise to high or low credit scores pervade debates about scoring algorithms. For example, in hearings about the use of credit scores in insurance pricing, legislators routinely testify that it is unfair for rates to go up for senior citizens whose low scores come from decades of prudently avoiding borrowing (Kiviat 2019c). In telling the story (of either a specific senior or seniors writ large), a sparse credit history takes on new meaning: It becomes the mark of a life virtuously led. Likewise, consumer advocates question the increasing inclusion of utility bill payments in credit scores, concerned that in the winter, poor people may let skyrocketing heating bills temporarily slide and later be penalized for that decision (Wu 2012). Here, it is not the past but the imagined future that confers meaning on the present. A person who *intends* to repay ought to be given special moral status, or at least that is the normative position a certain narrative lets us consider.

In these situations, people use narrative to inject information about actors' motivations and intentions, future plans (e.g., making the utility company whole once spring comes), and external constraints and structural forces. These details—the sort stripped away when people

are rendered as cases—then shed light on whether algorithms disadvantage people when they do not deserve to be. In other words, *narrative helps sort out whether the standards people are held to are fair*.

This is not to say that filling in the details of noncomparative justice is easy. There can be much disagreement about which standards are the right ones. For example, in debates about credit scores in car insurance pricing, two standards perpetually clash: one that looks to the future and says drivers predicted to file insurance claims rightly pay more and one that also looks to the past and says among those predicted to be costly, only those who have demonstrated irresponsible behavior ought to face a surcharge (Kiviat 2019c). For those backing the second standard, high rates should not follow if a low credit score is not a person's fault or if the link between scores and claims is driven by socioeconomic standing—for example, poor people both struggle to keep good credit and consistently file claims when bad things happen because they cannot pay for claims out of pocket. Different actors in the debate mobilize different stories, but whatever story they tell, the point is that it *is* a story.[19] The discursive terrain on which moral evaluation takes places is one of narrative.

Yet to underscore an earlier point, this does not mean narrative is a better tool for moral evaluation overall. Rendering people in narrative lets us see some moral fault lines, and rendering them as cases allows us to see others. In fact, at times, the two cultural forms work in concert. Consider, for example, that simply wanting allocative systems to treat people equally says nothing about *which* people ought to experience equivalent outcomes. Why, for instance, is it so important in the United States to test credit scoring (and other) algorithms for different outcomes by race? Because of the history (i.e., the *story*) of hundreds of years of subjugation and marginalization, in which powerful people and institutions limited the agency and life chances of Black Americans. Rendering people as cases may help make inequalities visible, but narrative is what deciphers which inequalities are the important ones to nullify in the first place.

## DISCUSSION

In their classic *Sorting Things out*, Bowker and Star (2000:5–6) write that decisions about how to categorize and classify the social world—including the people who inhabit it—are deeply moral ones in that "each category valorizes some point of view and silences another." When we categorize people in some ways but not others, we lay the groundwork for particular forms of inclusion and exclusion, advantage and disadvantage, and other morally relevant distinctions. The literature on categorization and classification echoes this insight many times over (e.g., Bourdieu 1984; Starr 1992; Zerubavel 1997).

What I have added here is that the choice to turn to categories, as opposed to narrative, in the first place is itself a morally consequential one. Reducing people to cases—to instances of particular, categorically defined types—means some aspects of lived human experience become more salient than others. The same is true when we reduce lived experience to narrative, but the reduction is of a fundamentally different sort. Moral consequences follow because different understandings of fairness draw on different types of information about people, some more readily provided by case-based renderings and others by storytelling. More specifically, understanding people as cases affords moral evaluations that look for injustices *across individuals or groups*, whereas understanding people as actors in narrative affords moral evaluations that look for injustices *across time*—through the unfolding events of a person's life.

This matters greatly because modernity means living in a case-based world, in a society prone to category-centric rationality, one that reduces people—and other types of entities—to constellations of decontextualized attributes. Rendering people as cases facilitates

bureaucracy, quantification, and other hallmarks of contemporary social organization (Lampland and Star 2009; Lipsky 1980; Prottas 1979; Scott 1998; Weber 1946). It also makes some sorts of (comparative) injustice more visible than they may be otherwise. When similarity and difference are homogenized, it is easier to compare people, and when it is easier to compare people, we more readily see the ways they compare unequally. The flip side is that rendering people as cases obscures noncomparative forms of justice. It is perhaps unsurprising, then, that when social actors are faced with systems that render people as cases, they often turn to narrative to fill in the epistemic and moral gaps (see Kiviat 2019a; Stevens 2007).

Consider, for instance, federal sentencing guidelines. Instituted in the 1980s as a way to achieve more uniform sentences for people committing similar crimes, the guidelines involve judges applying an elaborate rubric that includes classifying both the crime at hand and a defendant's criminal record (Espeland and Vannebo 2007; Schauer 2003). That is, the guidelines present defendants as cases: as combinations of discrete, regularized, and atemporal categories. With people rendered this way, it is easier to detect inequalities in how they are treated. Yet in practice, there is substantial tension. Judges, attorneys, and other court officials constantly push back on how guidelines capture and act on defendants, calling into question whether individual outcomes truly are deserved. And as they do so, they rely on narrative. As Lynch (2017) shows, courtroom actors tell stories about defendants' life trajectories, including where they have come from and where they are going, focusing on arcs of redemption, incorrigibility, moral agency, and so forth, all to make the case for a different sort of fairness in outcome, one we would call noncomparative.

This tendency to reach beyond the record reflects that an information infrastructure based entirely on categories and classes does not give actors everything they need to morally reason. In many settings, people care about both comparative and noncomparative notions of justice. Injecting narrative into discussions in such settings is not simply an act of discretion or an effort to reassert professional authority but also an implicit acknowledgment that without understanding the stories linked to cases, something *moral* goes missing. This implies that for organizational gatekeepers, policymakers, or others in positions of power to make full moral sense of how decision-making systems—algorithmic or otherwise—allocate benefits and burdens, they must understand people as *both* cases *and* actors in narrative. Each rendering affords a different sort of moral evaluation. Those who would dismiss either—cases as antiseptic and heartless or narrative as sentimental and unscientific—are limiting their ability to see the full landscape of morally relevant features. This brings an important implication for studies of algorithmic fairness, in particular. The tools that create algorithms, which depend on rendering people as cases, are not on their own sufficient for morally evaluating the outcomes they produce.

More broadly, this argument has demonstrated the value in being attentive to the ways that particular cultural forms (e.g., categories, stories) carry moral consequences. It is not just the *content* of the categories we use, the stories we tell, and so on but also the actual forms themselves and the scaffolding on which they rest (here, the information infrastructure) that funnels normative evaluation in one direction or another. The mere act of categorization affords particular kinds of moral reasoning. By marrying an extant philosophical framework, that of comparative and noncomparative justice, with the sociological study of categorization and classification, I have called attention to how seemingly innocuous decisions about how we represent our social world can afford some forms of moral thinking while making others more difficult.

At the end of the day, the world in which we live is only becoming more quantitative and computational. Many scholars, not to mention everyday observers, have proffered warnings

about algorithmic allocation, pointing to the increasingly opaque (i.e., "black-box") programs that shape who gets what. Yet as I have shown, long before we get to the fancy mechanics of machine learning and the like, there is a defining moral turn in the road these systems take, one that is as old as formal rationality itself—one that assumes people can be fully known through static, atemporal, and decontextualized attributes and that fundamentally different ways of knowing people and making decisions about them can be unproblematically left aside.

## ACKNOWLEDGMENTS

For insightful comments and advice, the author thanks Carly Knight, Mark Granovetter, Karen Levy, Momin Malik, Laura Adler, Kim Pernell, and the reviewers. The author is grateful for feedback received at the Radcliffe Institute's 2020 Ethics of Technology conference and as part of Australian National University's Humanising Machine Intelligence public online seminar series.

## ORCID ID

Barbara Kiviat https://orcid.org/0000-0002-2142-5871

## NOTES

1. In this article, I use the unmarked terms "justice" and "fairness" in their broad senses to indicate conceptions of the good, right, and morally desirable.
2. The processes of categorization and classification are sufficient for turning people into cases, although quantification does even more to homogenize similarity and difference.
3. Scholars across disciplines make this and similar distinctions between what psychologist Bruner (1986) calls the logico-scientific and narrative modes of thought. Notable examples in sociology include Heimer's (2001) discussion of case versus biography and Somers's (1994) examination of social science explanation versus the narrative "other."
4. I focus on narrative as a cultural form through which individuals and organizations make sense of people, but narrative can also be a strategic tool used in attempts to convince and influence (as in social movements).
5. Some algorithmic predictions may need data from only one person. For example, an algorithm designer might be able to use a single person's Fitbit data from the past year to predict what time she will wake up tomorrow. But this works only because waking up is a high-frequency event. A life insurer, on the other hand, could never predict when a particular person will die without looking at data from other people. Most outcomes sociologists care about are experienced by any given person infrequently (e.g., loan default, rearrest, completion of a college degree).
6. In this way, we also compare characters in narratives, perhaps especially so when they are types—for example, noting that the way a protagonist acts is not the way a person *like her* normally would.
7. Narratives themselves are also not immune to grouping. Consider genres, in which narratives are reduced to categorically defined types (e.g., a Western, a mystery, a romantic comedy).
8. Some critical theorists would go even further in articulating the moral danger. For example, a key part of Adorno's (1973) critique of modern society is its dependence on identity-thinking, or the constant reduction of specific and unique individuals to examples of abstract categories.
9. Other common distinctions include those between procedural and substantive justice, corrective and distributive justice, and conservative and ideal justice (Miller 2021).
10. Sen (1990) makes a similar point in his discussion of the "informational basis of justice," although he focuses on less concrete types of "personal features," including subjective utilities and commodity bundles enjoyed.
11. Gibson drew heavily on the work of Gestalt psychologists, including Kurt Koffka and Kurt Lewin.
12. This focus on the built environment means I follow in the tradition of Norman. Taking a more Gibsonian approach, which emphasizes aspects of the natural and biological environment, Keane (2016:27) develops the idea of "ethical affordances," or "aspects of people's experiences and perceptions that

they might draw on in the process of making ethical evaluations and decisions, whether consciously or not." One example is theory of mind: humans' ability to imagine what it is like to see through other people's eyes.

13. To analyze this literature, I searched for journal articles and conference proceedings featuring the phrase "algorithmic fairness" and turned to a series of highly cited review articles.
14. The subliterature on causality (e.g., Kilbertus et al. 2017; Loftus et al. 2018) may at first seem to contradict the focus on equivalence. In fact, this part of the literature is not so different: It uses formal, case-based approaches to causal reasoning (e.g., path diagrams) to detect inequalities.
15. These noncomparative harms could become comparative harms if they affect different groups in different ways; for example, if using education level in insurance pricing disproportionately disadvantages Black Americans because they have less education, on average. In this way, a noncomparative harm (the inequity of allocating insurance based on education level) can be compounded by a comparative one (Feinberg 1974:310).
16. Another part of the answer may be legal precedent (e.g., disparate impact law). That said, the law includes both comparative and noncomparative concerns (Levine and Pannier 2005).
17. Since the 1960s, various laws have restricted allocative decisions based directly on these and similar traits.
18. Injustice might also arise from using information that compounds prior injustice. For instance, even if victims of domestic abuse make worse insurance risks or prisoners who were abused as children are more likely to recidivate, Hellman (1997, 2019) considers it wrong to use such information in insurance and prison release decisions, respectively.
19. Which stories actors turn to likely reflects social location. Both narratives and category systems reflect what Haraway (1988) calls "situated knowledges."

## REFERENCES

Abbott, Andrew. 2001. *Time Matters: On Theory and Method*. Chicago, IL: University of Chicago Press.

Adorno, Thedor W. 1973. *Negative Dialectics*. New York, NY: Seabury Press.

Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. "Machine Bias." *ProPublica*, May 23. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

Aulck, Lovenoor, Dev Nambi, and Jevin West. 2019. "Using Machine Learning and Genetic Algorithms to Optimize Scholarship Allocation for Student Yield." Paper presented at SIGKDD '19: ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Anchorage, AK August 4–8.

Beckert, Jens. 2013. "Imagined Futures: Fictional Expectations in the Economy." *Theory and Society* 42(3):219–40.

Berk, Richard, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. 2021. "Fairness in Criminal Justice Risk Assessments: The State of the Art." *Sociological Methods & Research* 50(1):3–44.

Beniger, James R. 1986. *The Control Revolution: Technological and Economic Origins of the Information Society*. Cambridge, MA: Harvard University Press.

Best, Michael, Jenifer Bosco, and Chi Chi Wu. 2019. *Don*'t *Add Insult to Injury: Medical Debt & Credit Reports*. Boston, MA: National Consumer Law Center.

Bourdieu, Pierre. 1984. *Distinction: A Social Critique of the Judgement of Taste*. Cambridge, MA: Harvard University Press.

Bowker, Geoffrey C. 2005. *Memory Practices in the Sciences*. Cambridge, MA: MIT Press.

Bowker, Geoffrey C., and Susan Leigh Star. 2000. *Sorting Things out: Classification and Its Consequences*. Cambridge, MA: MIT Press.

Boxill, Bernard R. 1992. *Blacks and Social Justice*. Lanham, MD: Rowman and Littlefield.

Brayne, Sarah, and Angèle Christin. 2021. "Technologies of Crime Prediction: The Reception of Algorithms in Policing and Criminal Courts." *Social Problems* 68(3):608–24.

Bruner, Jerome. 1986. *Actual Minds, Possible Worlds*. Cambridge, MA: Harvard University Press.

Bruner, Jerome. 1990. *Acts of Meaning*. Cambridge, MA: Harvard University Press.

Bruner, Jerome. 1991. "The Narrative Construction of Reality." *Critical Inquiry* 18(1):1–21.

Burrell, Jenna, and Marion Fourcade. 2021. "The Society of Algorithms." *Annual Review of Sociology* 47:213–37.

Carrns, Ann. 2022. “A Health Issue for Your Credit Report.” *The New York Times*, March 5, B6.
Cerulo, Karen A., and Janet M. Ruane. 2014. “Apologies of the Rich and Famous: Cultural, Cognitive, and Social Explanations of Why We Care and Why We Forgive.” *Social Psychology Quarterly* 771(2):123–49.
Chong, Isis, and Robert W. Proctor. 2020. “On the Evolution of a Radical Concept: Affordances According to Gibson and Their Subsequent Use and Development.” *Perspectives on Psychological Science* 15(1):117–32.
Chouldechova, Alexandra, and Aaron Roth. 2018. “The Frontiers of Fairness in Machine Learning.” arXiv. https://arxiv.org/abs/1810.08810.
Consumer Financial Protection Bureau. 2012. “Key Dimensions and Processes in the U.S. Credit Reporting System: A Review of How the Nation’s Largest Credit Bureaus Manage Consumer Data.” https://files.consumerfinance.gov/f/201212_cfpb_credit-reporting-white-paper.pdf.
Consumer Financial Protection Bureau. 2022. “Medical Debt Burden in the United States.” https://files.consumerfinance.gov/f/documents/cfpb_medical-debt-burden-in-the-united-states_report_2022-03.pdf.
Corbett-Davies, Sam, and Sharad Goel. 2018. “The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning.*” arXiv*. https://arxiv.org/abs/1808.00023.
Couldry, Nick, and Andreas Hepp. 2016. *The Mediated Construction of Reality*. Malden, MA: Polity Press.
Daston, Lorraine. 1989. *Classical Probability in the Enlightenment*. Princeton, NJ: Princeton University Press.
Davis, Jenny L. 2020. *How Artifacts Afford: The Power and Politics of Everyday Things*. Cambridge, MA: MIT Press.
Davis, Jenny L., and James B. Chouinard. 2016. “Theorizing Affordances: From Request to Refuse.” *Bulletin of Science, Technology & Society* 36(4):241–48.
Deleuze, Gilles. 1994. *Difference and Repetition*. New York, NY: Columbia University Press.
Deleuze, Gilles, and Felix Guattari. 1987. *A Thousand Plateaus: Capitalism and Schizophrenia*. Minneapolis: University of Minnesota Press.
Desrosières, Alain. 1998. *The Politics of Large Numbers: A History of Statistical Reasoning*. Translated by C. Naish. Cambridge, MA: Harvard University Press.
Dieterich, William, Christina Mendoza, and Tim Brennan. 2016. “COMPAS Risk Scales: Demonstrating Accuracy Equity and Predictive Parity.” Working paper, Northpointe Inc., Traverse City, MI. http://go.volarisgroup.com/rs/430-MBX-989/images/ProPublica_Commentary_Final_070616.pdf.
Dourish, Paul. 2017. *The Stuff of Bits: An Essay on the Materialities of Information*. Cambridge, MA: MIT Press.
Dourish, Paul, and Melissa Mazmanian. 2013. “Media as Material: Information Representations as Material Foundations for Organizational Practice.” Pp. 92–118 in *How Matter Matters: Objects, Artifacts, and Materiality in Organization Studies*, edited by P. R. Carlile, D. Nicolini, A. Langley, and H. Tsoukas. Oxford, UK: Oxford University Press.
Durkheim, Emile. 1965. *The Elementary Forms of the Religious Life*. New York, NY: Free Press.
Dwork, Cynthia, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. “Fairness through Awareness.” Pp. 214–26 in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS ‘12)*. New York, NY: Association for Computing Machinery.
Eckhoff, Torstein. 1974. *Justice: Its Determinants in Social Interaction*. Rotterdam, the Netherlands: Rotterdam University Press.
Engler, Alex. 2021. “Enrollment Algorithms are Contributing to the Crisis of Higher Education.” Brookings Institution. https://www.brookings.edu/research/enrollment-algorithms-are-contributing-to-the-crises-of-higher-education/.
Espeland, Wendy. 1993. “Power, Policy and Paperwork: The Bureaucratic Representation of Interests.” *Qualitative Sociology* 16(3):297–317.
Espeland, Wendy N. 2015. “Narrating Numbers.” Pp. 56–75 in *The World of Indicators: The Making of Governmental Knowledge through Quantification*, edited by R. Rottenburg, S. Merry, S. Park, and J. Mugler. Cambridge, UK: Cambridge University Press.
Espeland, Wendy Nelson. 1998. *The Struggle for Water: Politics, Rationality, and Identity in the American Southwest*. Chicago, IL: University of Chicago Press.
Espeland, Wendy Nelson, and Mitchell L. Stevens. 1998. “Commensuration as a Social Process.” *Annual Review of Sociology* 24:313–43.
Espeland, Wendy Nelson, and Mitchell L. Stevens. 2008. “A Sociology of Quantification.” *European Journal of Sociology* 49(3):401–36.

Espeland, Wendy Nelson, and Berit Irene Vannebo. 2007. "Accountability, Quantification, and Law." *Annual Review of Law and Social Science* 3:21–43.
Eubanks, Virginia. 2017. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York, NY: St. Martin's Press.
Ewick, Patricia, and Susan S. Silbey. 1995. "Subversive Stories and Hegemonic Tales: Toward a Sociology of Narrative." *Law & Society Review* 29(2):197–226.
Fazelpour, Sina, and Zachary C. Lipton. 2020. "Algorithmic Fairness from a Non-ideal Perspective." Paper presented at the AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society, New York, NY, February 7–8.
Feinberg, Joel. 1970. *Doing and Deserving: Essays in the Theory of Responsibility*. Princeton, NJ: Princeton University Press.
Feinberg, Joel. 1973. *Social Philosophy*. Englewood Cliffs, NJ: Prentice-Hall.
Feinberg, Joel. 1974. "Noncomparative Justice." *The Philosophical Review* 83(3):297–338.
Feldman, Fred. 2016. "Feinberg on Comparative and Noncomparative Justice." Pp. 199–216 *in Distributive Justice: Getting What We Deserve from Our Country*, edited by F. Feldman. Oxford, UK: Oxford University Press.
Feldman, Fred, and Brad Skow. 2019. "Desert." In *The Stanford Encyclopedia of Philosophy*, edited by Zalta, E. N.. https://plato.stanford.edu/archives/win2019/entries/desert/.
Fourcade, Marion. 2011. "Cents and Sensibility: Economic Valuation and the Nature of 'Nature.'" *American Journal of Sociology* 116(6):1721–77.
Fourcade, Marion. 2016. "Ordinalization: Lewis A. Coser Memorial Award for Theoretical Agenda Setting 2014." *Sociological Theory* 34(3):175–95.
Fourcade, Marion, and Kieran Healy. 2017a. "Categories All the Way down." *Historical Social Research* 43(1):286–96.
Fourcade, Marion, and Kieran Healy. 2017b. "Seeing Like a Market." *Socio-Economic Review* 15(1):9–29.
Friedler, Sorelle A., Carlos Scheidegger, and Suresh Venkatasubramanian. 2016. "On the (Im)possibility of Fairness." arXiv. https://arxiv.org/abs/1609.07236.
Frye, Margaret. 2012. "Bright Futures in Malawi's New Dawn: Educational Aspirations as Assertions of Identity." *American Journal of Sociology* 117(6):1565–624.
Gandy, Oscar H. 1993. *The Panoptic Sort: A Political Economy of Personal Information*. Boulder, CO: Westview Press.
Gibson, James J. 1979. *The Ecological Approach to Visual Perception*. Boston, MA: Houghton, Mifflin and Company.
Gillespie, Tarleton. 2014. "The Relevance of Algorithms." Pp. 167–93 in *Media Technologies: Essays on Communication, Materiality, and Society*, edited by T. Gillespie, P. J. Boczkowski, and K. A. Foot. Cambridge, MA: MIT Press.
Guseva, Alya, and Akos Rona-Tas. 2001. "Uncertainty, Risk, and Trust: Russian and American Credit Card Markets Compared." *American Sociological Review* 66(5):623–46.
Hacking, Ian. 1986. "Making Up People." Pp. 161–72 in *Reconstructing Individualism: Autonomy, Individuality, and the Self in Western Thought*, edited by Heller, T. C., M. Sosna, and D. E. Wellbery. Stanford, CA: Stanford University Press.
Hacking, Ian. 1990. *The Taming of Chance*. Cambridge, UK: Cambridge University Press.
Haraway, Donna. 1988. "Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective." *Feminist Studies* 14(3):575–99.
Harding, David J., Cheyney C. Dobson, Jessica J. B. Wyse, and Jeffrey D. Morenoff. 2017. "Narrative Change, Narrative Stability, and Structural Constraint: The Case of Prisoner Reentry Narratives." *American Journal of Cultural Sociology* 5(1):261–304.
Heimer, Carol A. 2001. "Case and Biographies: An Essay on Routinization and the Nature of Comparison." *Annual Review of Sociology* 27:47–76.
Heimer, Carol A. 2009. "Conceiving Children: How Documents Support Case versus Biographical Analysis." Pp. 95–126 in *Documents: Artifacts of Modern Knowledge*, edited by A. Riles. Ann Arbor: University of Michigan Press.
Hellman, Deborah. 2019. "Indirect Discrimination and the Duty to Avoid Compounding Injustice." Pp. 105–21 in *Foundations of Indirect Discrimination*, edited by H. Collins and T. Khaitan. London, UK: Hart Publishing.

Hellman, Deborah. 2020. “Measuring Algorithmic Fairness.” *Virginia Law Review* 106(4):811–66.
Hellman, Deborah S. 1997. “Is Actuarially Fair Insurance Pricing Actually Fair? A Case Study in Insuring Battered Women.” *Harvard Civil Rights-Civil Liberties Law Review* 32:355–412.
Hirschman, Daniel, Ellen Berrey, and Fiona Rose-Greenland. 2016. “Dequantifying Diversity: Affirmative Action and Admissions at the University of Michigan.” *Theory and Society* 45(3):265–301.
Jencks, Christopher. 1988. “Whom Must We Treat Equally for Educational Opportunity to Be Equal?” *Ethics* 98(3):518–33.
Keane, Webb. 2016. *Ethical Life: Its Natural and Social Histories*. Princeton, NJ: Princeton University Press.
Kilbertus, Niki, Mateo Rojas-Carulla, Giambattista Parascandolo, Moritz Hardt, Dominik Janzing, and Bernhard Schölkopf. 2017. “Avoiding Discrimination through Causal Reasoning.” Pp. 656–66 in *NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing Systems*. New York, NY: Association for Computing Machinery.
Kiviat, Barbara. 2019a. “The Art of Deciding with Data: Evidence from How Employers Translate Credit Reports into Hiring Decisions.” *Socio-Economic Review* 17(2):283–309.
Kiviat, Barbara. 2019b. “Credit Scoring in the United States.” *Economic Sociology: The European Electronic Newsletter* 21(1):33–42.
Kiviat, Barbara. 2019c. “The Moral Limits of Predictive Practices: The Case of Credit-Based Insurance Scores.” *American Sociological Review* 84(6):1134–58.
Kiviat, Barbara. 2021. “Which Data Fairly Differentiate? American Views on the Use of Personal Data in Two Market Settings.” *Sociological Science* 8:26–47.
Kleinberg, Jon, Sendhil Mullainathan, and Manish Raghavan. 2017. “Inherent Trade-Offs in the Fair Determination of Risk Scores.” Pp. 43:1–43:23 in *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)(Leibniz International Proceedings in Informatics (LIPIcs))*, Vol. 67, edited by Christos H. Papadimitriou. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany. doi:10.4230/LIPIcs.ITCS.2017.43.
Knight, Carly, and Isaac Ariail Reed. 2019. “Meaning and Modularity: The Multivalence of ‘Mechanism’ in Sociological Explanation.” *Sociological Theory* 37(3):234–56.
Krippner, Greta R. 2017. “Democracy of Credit: Ownership and Politics of Credit Access in Late Twentieth-Century America.” *American Journal of Sociology* 123(1):1–47.
Krippner, Greta R., and Daniel Hirschman. 2022. “The Person of the Category: The Pricing of Risk and the Politics of Classification in Insurance and Credit.” *Theory and Society* 51(5):685–27.
Lamont, Michèle. 2000. The Dignity of Working Men: Morality and the Boundaries of Race, Class, and Immigration. New York, NY: Russell Sage Foundation.
Lampland, Martha, and Susan Leigh Star. 2009. “Reckoning with Standards.” Pp. 3–24 in *Standards and Their Stories: How Quantifying, Classifying, and Formalizing Practices Shape Everyday Life*, edited by M. Lampland and S. Star. Ithaca, NY: Cornell University Press.
Lempert, Richard, and Joseph Sanders. 1986. *An Invitation to Law and Social Science: Desert, Disputes, and Distribution*. New York, NY: Longman.
Lengwiler, Martin. 2009. “Double Standards: The History of Standardizing Humans in Modern Life Insurance.” Pp. 95–113 in *Standards and Their Stories: How Quantifying, Classifying, and Formalizing Practices Shape Everyday Life*, edited by Lampland, M., and S. Star. Ithaca, NY: Cornell University Press.
Levine, Raleigh Hannah, and Russell Pannier. 2005. “Comparative and Noncomparative Justice: Some Guidelines for Constitutional Adjudication.” *William & Mary Bill of Rights Journal* 14(1):141–92.
Levy, Karen, Kyla E. Chasalow, and Sarah Riley. 2021. “Algorithms and Decision-Making in the Public Sector.” *Annual Review of Law and Social Science* 17:309–34.
Lipsky, Michael. 1980. *Street-Level Bureaucracy: The Dilemmas of the Individual in Public Service*. New York, NY: Russell Sage Foundation.
Loftus, Joshua R., Chris Russell, Matt J. Kusner, and Ricardo Silva. 2018. “Causal Reasoning for Algorithmic Fairness.” arXiv. https://arxiv.org/abs/1805.05859.
Lucas, J. R. 1980. *On Justice*. New York, NY: Oxford University Press.
Lynch, Mona. 2017. “The Narrative of the Number: Quantification in Criminal Court.” *Law & Social Inquiry* 44(1):31–57.

Malik, Momin M. 2020. "A Hierarchy of Limitations in Machine Learning." arXiv. https://arxiv.org/abs/2002.05193.
Manovich, Lev. 1999. "Database as Symbolic Form." *Convergence* 5(2):80–99.
Mehrabi, Ninareh, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2022. "A Survey on Bias and Fairness in Machine Learning." *ACM Computing Surveys* 54(6):1–35.
Miller, David. 2021. "Justice." In *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta. https://plato.stanford.edu/archives/fall2021/entries/justice/.
Minow, Martha. 1990. *Making All the Difference: Inclusion, Exclusion, and American Law*. Ithaca, NY: Cornell University Press.
Mische, Ann. 2009. "Projects and Possibilities: Researching Futures in Action." *Sociological Forum* 24(3):694–704.
Mitchell, Shira, Eric Potash, Solon Barocas, Alexander D'Amour, and Kristian Lum. 2021. "Algorithmic Fairness: Choices, Assumptions, and Definitions." *Annual Review of Statistics and Its Application* 8:141–63.
Mohr, John. 2005. "The Discourses of Welfare and Welfare Reform." Pp. 346–63 in *The Blackwell Companion to the Sociology of Culture*, edited by M. D. Jacobs and N. Weiss Hanrahan. Malden, MA: Blackwell Publishing Ltd.
Monroe, Andrew E., and Bertram F. Malle. 2017. "Two Paths to Blame: Intentionality Directs Moral Information Processing Along Two Distinct Tracks." *Journal of Experimental Psychology* 146(1):123–33.
Montague, Phillip. 1980. "Comparative and Non-comparative Justice." *The Philosophical Quarterly* 30(119):131–40.
Norman, Donald A. 2002. *The Design of Everyday Things*. New York, NY: Basic Books.
O'Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York, NY: Random House.
Orbuch, Terri L. 1997. "People's Accounts Count: The Sociology of Accounts." *Annual Review of Sociology* 23:455–78.
Pasquale, Frank. 2015. *The Black Box Society: The Secret algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press.
Polkinghorne, Donald. 1988. *Narrative Knowing and the Human Sciences*. Albany, NY: SUNY Press.
Polletta, Francesca. 1998. "Contending Stories: Narrative in Social Movements." *Qualitative Sociology* 21(4):419–46.
Polletta, Francesca, Pang Chin Bobby Chen, Beth Gharrity Gardner, and Alice Motes. 2011. "The Sociology of Storytelling." *Annual Review of Sociology* 37:109–30.
Porter, Theodore M. 1994. "Information, Power and the View from Nowhere." Pp. 217–30 in *Information Acumen: The Understanding and Use of Knowledge in Modern Business*, edited by L. Bud-Frierman. London, UK: Routledge.
Porter, Theodore M. 1995. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton, NJ: Princeton University Press.
Prottas, Jeffrey Manditch. 1979. *People-Processing: The Street-Level Bureaucrat in Public Service Bureaucracies*. Lexington, MA: Lexington Books.
Rachels, James. 1991. "What People Deserve." Pp. 136–48 in *Justice and Economic Distribution*, edited by J. Arthur and W. H. Shaw. Englewood Cliffs, NJ: Prentice Hall.
Rae, Douglas. 1979. "The Egalitarian State: Notes on a Contradictory System of Ideals." *Daedalus* 108(4):37–54.
Ricoeur, Paul. 1980. "Narrative Time." *Critical Inquiry* 7(1):169–90.
Ricoeur, Paul. [1981] 2016. "The Narrative Function." Pp. 236–58 in *Hermeneutics and the Human Sciences: Essays on Language, Action and Interpretation*, edited by J. B. Thompson. Cambridge, UK: Cambridge University Press.
Salais, Robert. 2016. "Quantification and Objectivity: From Statistical Conventions to Social Conventions." *Historical Social Research* 41(2):118–34.
Schauer, Frederick. 2003. *Profiles, Probabilities, and Stereotypes*. Cambridge, MA: Harvard University Press.
Scott, James C. 1998. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven, CT: Yale University Press.

Scott, Marvin B., and Stanford M. Lyman. 1968. "Accounts." *American Sociological Review* 33(1):46–62.
Selbst, Andrew D., danah boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. "Fairness and Abstraction in Sociotechnical Systems." Paper presented at FAT* '19: Conference on Fairness, Accountability, and Transparency (FAT* '19), Atlanta, GA, January 29–31.
Sen, Amartya. 1990. "Justice: Means versus Freedoms." *Philosophy & Public Affairs* 19(2):111–21.
Shelby, Tommie. 2016. *Dark Ghettos: Injustice, Dissent, and Reform*. Cambridge, MA: Harvard University Press.
Sher, George. 1987. *Desert*. Princeton, NJ: Princeton University Press.
Somers, Margaret R. 1992. "Narrativity, Narrative Identity, and Social Action: Rethinking English Working-Class Formation." *Social Science History* 16(4):591–630.
Somers, Margaret R. 1994. "The Narrative Constitution of Identity: A Relational and Network Approach." *Theory and Society* 23(5):605–49.
Somers, Margaret R., and Fred Block. 2005. "From Poverty to Perversity: Ideas, Markets, and Institutions over 200 Years of Welfare Debate." *American Sociological Review* 70(2):260–87.
Star, Susan Leigh, and Karen Ruhleder. 1996. "Steps toward an Ecology of Infrastructure: Design and Access for Large Information Spaces." *Information Systems Research* 7(1):111–34.
Starr, Paul. 1987. "The Sociology of Official Statistics." Pp. 7–58 in *The Politics of Numbers*, edited by W. Alonso and P. Starr. New York, NY: Russell Sage Foundation.
Starr, Paul. 1992. "Social Categories and Claims in the Liberal State." *Social Research* 59(2):263–95.
Steensland, Brian. 2010. "Moral Classification and Social Policy." Pp. 455–68 in *Handbook of the Sociology of Morality*, edited by S. Hitlin and S. Vaisey. New York, NY: Springer.
Steinmetz, George. 1992. "Reflections on the Role of Social Narratives in Working-Class Formation: Narrative Theory in the Social Sciences." *Social Science History* 16(3):489–516.
Stevens, Mitchell L. 2007. *Creating a Class: College Admissions and the Education of Elites*. Cambridge, MA: Harvard University Press.
Stone, Deborah A. 1989. "Causal Stories and the Formation of Policy Agendas." *Political Science Quarterly* 104(2):281–300.
Tavory, Iddo, and Nina Eliasoph. 2013. "Coordinating Futures: Toward a Theory of Anticipation." *American Journal of Sociology* 118(4):908–42.
Tavory, Iddo, and Stefan Timmermans. 2009. "Two Cases of Ethnography: Grounded Theory and the Extended Case Method." *Ethnography* 10(3):243–63.
Tavory, Iddo, and Robin Wagner-Pacifici. 2022. "Climate Change as an Event." *Poetics* 93(A):101600. doi:10.1016/j.poetic.2021.101600.
Tilly, Charles. 1998. *Durable Inequality*. Berkeley: University of California Press.
Timmermans, Stefan. 1996. "Saving Lives or Saving Multiple Identities?: The Double Dynamic of Resuscitation Scripts." *Social Studies of Science* 26(4):767–97.
Timmermans, Stefan, and Steven Epstein. 2010. "A World of Standards but not a Standard World: Toward a Sociology of Standards and Standardization." *Annual Review of Sociology* 36:69–89.
Underwood, Barbara D. 1979. "Law and the Crystal Ball: Predicting Behavior with Statistical Inference and Individualized Judgment." *Yale Law Journal* 88(7):1408–48.
Verma, Sahil, and Julia Rubin. 2018. "Fairness Definitions Explained." Pp. 1–7 in Proceedings of the International Workshop on Software Fairness. New York, NY: Association for Computing Machinery.
Wagner-Pacifici, Robin. 2010. "Theorizing the Restlessness of Events." *American Journal of Sociology* 115(5):1351–86.
Weber, Max. 1946. "Bureaucracy." Pp. 196–244 in *From Max Weber: Essays in Sociology*, edited by H. H. Gerth and C. Wright Mills. New York, NY: Oxford University Press.
White, Hayden. 1980. "The Value of Narrativity in the Representation of Reality." *Critical Inquiry* 7(1):5–27.
White, Hayden. 1984. "The Question of Narrative in Contemporary Historical Theory." *History and Theory* 23(1):1–33.
Williams, Bernard. 1964. "The Idea of Equality." Pp. 110–31 in *Philosophy, Politics and Society*, edited by P. Laslett and W. G. Runciman. Oxford, UK: Basil Blackwell.
Wu, Chi Chi. 2012. *Testimony before House Subcommittee on Financial Institutions and Consumer Credit of the Committee on Financial Services. Examining the Uses of Consumer Credit Data. Hearing held September 13. 112th Congress, 2nd session*. Washington, DC: U.S. Government Printing Office.

Zerubavel, Eviatar. 1996. "Lumping and Splitting: Notes on Social Classification." *Sociological Forum* 11(3):421–33.

Zerubavel, Eviatar. 1997. *Social Mindscapes: An Invitation to Cognitive Sociology*. Cambridge, MA: Harvard University Press.

Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York, NY: Public Affairs.

## AUTHOR BIOGRAPHY

**Barbara Kiviat** is an assistant professor of sociology at Stanford University. She is an economic sociologist who studies how moral beliefs and other cultural understandings shape markets and justify the inequalities they produce. Her recent research focuses on the corporate use of personal data in the pricing and allocation of socially important resources, such as insurance, housing, and jobs. Her work has appeared in the *American Sociological Review, Socio-Economic Review, Sociological Science*, and *Social Psychology Quarterly*, among other outlets.