Which Data Fairly Differentiate? American Views on the Use of Personal Data in Two Market Settings

Barbara Kiviat

Stanford University

Abstract: Corporations increasingly use personal data to offer individuals different products and prices. I present first-of-its-kind evidence about how U.S. consumers assess the fairness of companies using personal information in this way. Drawing on a nationally representative survey that asks respondents to rate how fair or unfair it is for car insurers and lenders to use various sorts of information—from credit scores to web browser history to residential moves—I find that everyday Americans make strong moral distinctions among types of data, even when they are told data predict consumer behavior (insurance claims and loan defaults, respectively). Open-ended responses show that people adjudicate fairness by drawing on shared understandings of whether data are logically related to the predicted outcome and whether the categories companies use conflate morally distinct individuals. These findings demonstrate how dynamics long studied by economic sociologists manifest in legitimating a new and important mode of market allocation.

Keywords: personal data; fairness in markets; credit; insurance; cultural beliefs

COMPANIES increasingly use personal data to predict how consumers will behave and offer them different products, prices, and levels of service. Companies collect information about individuals from corporate and government databases, devices like cell phones, and in-person interactions and then circulate it, commodity-like, through networks of data brokers for use across organizations (Fourcade and Healy 2017b; Zuboff 2019). What a company offers a consumer and what it asks him or her to pay now may depend on any number of personal details, such as web browser history, driving patterns, debit card activity, exercise habits, marital status, drug prescriptions, and more. Companies justify this use of data with appeals to the virtues of personalization: people enter the market with different preferences, risk profiles, and abilities to pay, which makes it legitimate, desirable even, for people to receive different things in return (Cohen 2013; Rule 2009; Seaver 2015).

The use of personal data in determining market outcomes is reshaping the economic lives of everyday Americans—and yet we know little about the extent to which Americans think various sorts of personal data are fair to use in this way.[1] This article offers first-of-its-kind insight into this question, drawing on an original, nationally representative survey of 1,095 U.S. adults asked to evaluate the fairness of companies using 16 types of data in two market settings: car insurance and consumer lending. It is possible that Americans are broadly at ease with companies using all sorts of data to differentiate. In economic contexts, Americans often think it is fair for different people to get different things (Hochschild 1981; Kluegel and Smith 1986), an outlook reinforced by cultural celebration of markets as efficient

and good. That said, economic sociologists consistently show that people make moral distinctions in market settings and get offended when transactions mismatch the relationship at hand (Bandelj 2020; Wherry 2016; Zelizer 1996, 2012) or when the underlying categories markets depend on betray broader moral commitments (Fourcade and Healy 2007; Kiviat 2019b; Zelizer 1979). Despite industry's efforts at commodification, data may retain social and moral meaning, dictating how Americans think they ought to be used.

How everyday Americans morally evaluate the corporate use of personal data is an important question because these innovations in labeling and sorting consumers represent a powerful new mode of stratification. As Fourcade and Healy (2013) argue, markets no longer simply reflect inequalities tied to class structure and ascriptive traits like race and sex. Companies now have the power to create fresh positions of advantage and disadvantage based on whom they expect to be the most profitable customers. The ways companies categorize individuals give rise to distinctions that are materially and symbolically meaningful. People pay more or less for products, get better or worse loan terms, wait on longer or shorter telephone queues, encounter sympathy or suspicion when filing insurance claims, and so on. A key component of any system of stratification is a set of cultural beliefs that legitimate why some people wind up better or worse off (Della Fave 1980; Shepelak 1987); this article offers insight into the current state of such beliefs.

The results of the survey reveal that Americans make strong moral distinctions among the types of personal data used in market transactions. Quantitative findings show that people think about data fairness in market-specific ways and that newer forms of data that capture the ostensibly free actions of individuals are at times judged no more favorably than ascriptive traits such as race and sex. Within each market, Americans are largely unified in their beliefs that some sorts of data are permissible and others are proscribed, yet many other types of data fall in between. For these unsettled data, Americans often hold strong but opposing opinions. Open-ended responses reveal that behind this fragmentation sit different views about whether data are related to what companies are trying to predict and whether data group people in morally consistent ways. Taken together, these findings point to a deep disconnect between corporate practice and everyday Americans' moral intuitions—as well as to pathways by which Americans' opinions may ultimately be swayed as corporations work to institutionalize their version of fairness and policymakers increasingly push back.

The Corporate Use of Personal Data

Corporations today have more access to data about individuals than ever before, thanks to a boom in digital information and brokers that circulate it. Companies use these data to categorize and rank consumers in order to more profitably decide whom to offer what (Fourcade and Healy 2017a; Rule 2009; Zuboff 2019). The options and prices a person sees may now depend on a wide array of personal details, from cell phone use to retail purchase history to credit score and college major. Aside from a few verboten categories, such as race, companies rarely make normative distinctions among the data used for these purposes (Moor and Lury

2018; Williams, Brooks, and Shmargad 2018; Zwick and Knott 2009). What is important is whether data help predict which consumers are the valuable ones. The goal, then, is to collect as much data as possible, irrespective of an intuitive connection to the task at hand (Fourcade and Healy 2017b). As the CEO of a company using unconventional data for lending decisions put it, "all data is credit data" (Deville and van der Velden 2016:94). The remark points to a belief in the potential usefulness of any information that can be had as well as the permissibility of using it.

Economic sociologists show that markets depend on particular moral orientations, even when those orientations are naturalized and difficult to see (Fourcade and Healy 2007). In this new ecosystem of mass data use, personalization is assumed to be a moral good—not only the path to increased profit, but also the mark of a more rightful market. Companies portray individualization as a way to give people what they like and want, as well as what they deserve, by more precisely identifying who is creditworthy, insurable, price sensitive, and so on (Cohen 2013; Fourcade and Healy 2017b; Seaver 2015). Bolstering this ethos is that many new types of data are "behavioral" in that they capture traces of information people leave behind as they go about their day-to-day lives, the product of seemingly volitional decisions for which people can ostensibly be held accountable. The fact that data appear to be about people as individuals, rather than members of social groups, lends further legitimacy (Fourcade 2016; Krippner 2017; Starr 1992), although scholars underscore that data can act invisibly as a proxy for race, sex, and other traits people do not control (Barocas and Selbst 2016; Gandy 2009; Noble 2018).

Credit and insurance, the industries examined in this article, are forerunners of this new style of capitalism. Lenders and insurers have long collected information about individuals in order to predict how they will behave—whether they will, say, default on a loan or file an insurance claim—and then adjust offerings as a result. When companies predict that a person will be a costly customer, they charge that person a higher interest rate or larger premium or deny a loan or policy altogether. This represents a profit-making strategy as well as a moral understanding that people rightly bear the cost of their own risk (Baker and Simon 2002; Kiviat 2019a; Stone 1993). In the United States, the use of personal data in car insurance and consumer lending is institutionalized and well known.

Even so, recent decades have witnessed a series of public policy debates about which sorts of personal data lenders and insurers can fairly use in differentiating among consumers. Starting in the 1960s, policymakers at both the state and federal levels worked to prevent market decisions, including those about credit and insurance, from hinging on certain social markers, such as race, religion, and national origin (Avraham, Logue, and Schwarcz 2014; Capon 1982; Krippner 2017). In more recent years, as companies have expanded the information they use to slot and sort consumers, policymakers have occasionally raised alarm, questioning lenders' use of, among other things, unpaid medical debt, college major, utility bill payment, and social media connections (Consumer Financial Protection Bureau 2014; Task Force on Financial Technology 2019). In car insurance, policymakers have investigated the use of credit scores, web browser history, zip codes, education level, data from

devices that track driving in real time, and more (Banham 2015; Karapiperis et al. 2015; Kiviat 2019b; State of New Jersey Department of Banking and Insurance 2008).

To justify using a particular type of data, companies often point out that the data mathematically predict an outcome of legitimate interest, such as tendency to repay money on time. For example, when American Express was pilloried in the press for using data about where people shop to set credit card limits, a spokesperson defended the practice by arguing that "it's purely math" (Teegardin 2008; see also Board of Governors of the Federal Reserve System 2010). That data predict functions as a moral justification, albeit a thin one that is prone to challenge (Hellman 1997; Underwood 1979). Studying policy debate about the use of credit scores in car insurance pricing, Kiviat (2019b) shows that for policymakers to accept predictive data as fair, they also needed to be convinced of palatable causal connections between data and outcome and that using the data did not improperly group people with distinct moral standing (e.g., those with bad credit from irresponsible behavior vs. those with bad credit through no fault of their own).

Consumer Views on Economic Differentiation and Data Use

Although the literature offers insight into how corporate and policy elites make distinctions among fair and unfair data, we have little corresponding evidence about the beliefs of everyday Americans. It is possible that consumers take the same view as corporations, that markets fairly differentiate on the basis of all kinds of information. When it comes to economic matters, as opposed to, say, politics or home life, Americans tend to define fairness through differentiation, assuming that people are "different in ways that usually call for unequal allocations" (Hochschild 1981:51; Kelley and Evans 1993). Those who can pay more buy more, those who are more skilled earn more, and so on—a state of affairs generally seen as unproblematic in part because of widespread American belief in a person's ability to choose freely and control his or her own fate (Kluegel and Smith 1986; McCall 2013; Shepelak 1987). This, combined with general faith in markets, suggests that Americans may be broadly at ease with companies using personal data to give different people different things.

That said, two perennial lessons from economic sociology about how people make distinctions in market settings suggest that Americans may instead hold more discerning views about which data companies can fairly use to allocate resources. First, economic transactions are deeply relational and full of social meaning (Bandelj 2020; Wherry 2016; Zelizer 1996). In economic exchange, people seek out what Zelizer (2012:151) calls "viable matches," demonstrating that the appropriateness of a transaction and its constituent parts depends on who the exchange partner is and the nature of the relationship (e.g., Polletta and Tufail 2014; Velthuis 2005). When aspects of an exchange do not align, the response is often "anger, shock, or ridicule" (Zelizer [1994] 2017:19). Similarly, privacy scholars show that opinions about who rightly sees information are relational and context dependent (Anthony, Campos-Castillo, and Horne 2017; Martin and Nissenbaum 2016; Smith 2018).[2]

People provide medical history to doctors and credit history to mortgage brokers, but most would balk at switching the two. This suggests that Americans may make distinctions among types of personal data based on whether they perceive the data to match the transaction at hand.

The second lesson is that markets depend upon systems of moral categorization (Fourcade and Healy 2007; Massengill and Reynolds 2010). At the extreme, cultural notions about what counts as sacred preclude certain things from being traded in the market in the first place (Healy 2006; Quinn 2008; Zelizer 1979). Yet even if there is widespread agreement that it is appropriate for markets to allocate a given resource, broader understandings about right and wrong, about value and worth, shape what people take to be legitimate parameters of exchange (Almeling 2007; Anteby 2010; Spillman 1999). Notably, these ideas are socially constructed and change over time, reflecting cultural shifts and the evolving interests of powerful actors (Homans 1974; Thompson 1971). Once, a defensible argument for paying married men more was that they had families to support, but no longer. This suggests that as people adjudicate which data are fair for companies to use in deciding who gets what, they will draw on moral understandings that transcend the market and judge harshly when market categories undermine more generalized normative commitments.

In the U.S. context, there's reason to believe those more generalized normative commitments may include ideas about agency, individual responsibility, and deservingness. Writing about the legitimate bases of social classification, Starr (1992) argues that in a liberal state such as the United States, people are likely to be most at ease with rewards tied to voluntary and meritorious behavior and with disadvantage tied to blameworthy acts individuals could have taken steps to avoid—rather than, say, immutable traits or social class (see also, Rubinstein 1988; Underwood 1979). This pattern is visible in a range of debates about the legitimate distribution of economic resources, from the role of meritocracy in who gets jobs and large incomes (Sauder 2020; Sherman 2017) to the primacy of deservingness in the allocation of state benefits (Watkins-Hayes and Kovalsky 2016; Steensland 2006). As companies increasingly turn to "behavioral" data—those arising from specific individuals' actions—some scholars suggest that differentiated market outcomes will strike people as natural and right (Fourcade 2016; Fourcade and Healy 2017b), a perspective that makes sense given broader cultural belief in holding people accountable for their own actions.

Past research on data use in car insurance and lending offers some, but limited, insight into how everyday Americans make moral distinctions. Recent years have seen a series of one-off surveys, typically in conjunction with policy debates about controversial types of data. These surveys, often funded by industry or advocacy organizations, show, for example, that everyday Americans generally think the use of credit scores in car insurance pricing is unfair (Heller and Styczynski 2016; O'Leary, Richards, and Quinlan 2013).[3] Yet, to the best of my knowledge, no extant survey compares fairness judgments across data types across industry, sources of variation that are fruitful for disambiguating reactions to the use of data from the use of *particular data* from the use of particular data *in particular contexts*. Nor do these surveys specify that companies hold data to be mathematically predictive, a potentially important qualifier for moral evaluations. In these ways, the survey

used in this article, described in the next section, goes further and promises to offer a more nuanced theoretical understanding of everyday Americans' beliefs about corporate data use.

Data and Methods

This article draws on a survey of 1,095 respondents designed to be representative of the U.S. adult population.[4] Each respondent saw two scenarios, presented in random order. In the first scenario, a car insurance company planned to use various sorts of personal data to predict who would file insurance claims and then charge those people higher prices for car insurance or not sell them insurance at all. In the second scenario, a lender planned to use various sorts of data to predict who would fail to repay a loan on time and then charge those people higher interest rates or not lend to them at all. In each case, respondents were told that the companies had done statistical analysis and had said that each type of information helped to predict insurance claims or loan nonrepayment. The full prompts appear in the online supplement.

After seeing each scenario, respondents were asked to rate the fairness of using each sort of data and, for a subset of data types, to explain their answers. The closed-ended scale consisted of five options: Very Fair (5), Somewhat Fair (4), Neither Fair nor Unfair (3), Somewhat Unfair (2), and Very Unfair (1).[5] Respondents evaluated one type of data at a time, presented in random order. The open-ended prompts reminded respondents of their answers and then asked for an explanation; for example, "Earlier in this survey, you said that in deciding how much to charge for car insurance, it would be Somewhat Fair for a car insurance company to use data from a device in the person's car that tracks what time of day or night they drive. Please explain your thinking."[6] The survey thus captured both quick-response moral intuitions and more deliberative moral reasoning (DiMaggio 1997; Lizardo et al. 2016; Vaisey 2009).

For each scenario, respondents rated 16 types of data, most of which are currently used by car insurers and/or lenders or are highly sought-after data assets in the broader "big data" economy. Table 1 lists each type of data, described the way respondents saw it. I included data according to four criteria. First, I included data that could be easily understood as either behavioral or nonbehavioral (e.g., speeding tickets and grocery store purchases vs. race and sex). Second, I included data that might be construed as either matched or mismatched to each industry (e.g., accident history vs. a person's Facebook posts). Third, I included data that are banned by law (e.g., race/ethnicity) to establish comparison with what would presumably be some of the least favorably rated data. And fourth, I included data that are the subject of current public policy debate (e.g., the payment of various sorts of bills, which factor prominently in discussions about "alternative" credit data).

Many of the data types respondents evaluated carry substantial weight in car insurance and lending decisions. Major factors in how car insurers price policies include accident history, speeding tickets, sex, credit score, and zip code (National Association of Insurance Commissioners 2011), while consumer lenders rely heavily

Table 1: Types of data respondents rated as fair or unfair to use

Car Insurance

A person's accident history
How many speeding tickets a person gets
A person's connections, posts, and "likes" on social media sites like Facebook
The number of addresses a person has lived at in the past 5 years
Data from a device in the person's car that tracks how much they slam on the brakes, accelerate hard, and turn sharply while driving
A person's credit report or credit score
How much money a person makes
A person's level of education (e.g., high school, college)
A person's sex/gender
A person's race/ethnicity
Which web sites a person visits
Data from a device in the person's car that tracks where they drive
A record of what the person buys at the grocery store
Whether a person rents or owns their home
Data from a device in the person's car that tracks what time of day or night they drive
The zip code a person lives in

Consumer Lending

The number of addresses a person has lived at in the past 5 years
A person's connections, posts, and "likes" on social media sites like Facebook
How often a person pays the cable TV bill on time
A person's credit report or credit score
How much money a person makes
How often the person pays the utility bill on time
A person's sex/gender
A person's race/ethnicity
Which web sites a person visits
What subjects the person studied in college (i.e., a person's major)
How often the person pays the rent on time
A record of what the person buys at the grocery store
How many speeding tickets a person gets
The zip code a person lives in
How often the person pays the childcare bill on time
Whether a person smokes

Note: Respondents saw data types one at a time and in random order.

on credit history and income. Other data, such as education level in car insurance and residential mobility in lending, are used less universally but can still substantially change what consumers pay (Boyle 2016; Florida Office of Insurance Regulation 2007). Less traditional data increasingly hold influence over decisions, as well. Large, established firms are quickly adopting some sorts of information (e.g., rent and utility bill payment in lending and real-time telematics driving data in car insurance), while start-ups trying to gain a competitive foothold have focused

on others (e.g., college major and social media data in lending) (Experian 2019; Karapiperis et al. 2015; Robinson and Yu 2014).

To analyze the results of the survey, I drew on both the quantitative, closed-ended questions and qualitative responses to the open-ended prompts. For the quantitative results, I applied weights designed to make the survey nationally representative.[7] For the free-text responses, I began by reading answers to each open-ended question and writing a memo to capture common and theoretically interesting answers. I then returned to the data for a second round of reading and memo-writing in order to identify similarities and differences across data types. Finally, I assigned codes to each open-ended response, which helped me more precisely see how responses clustered by evaluations of data being fair versus unfair. In this part of the analysis, after considering each response category on its own, I collapsed together Very Fair and Somewhat Fair, and Very Unfair and Somewhat Unfair, to make the major differences between fairness and unfairness more salient.
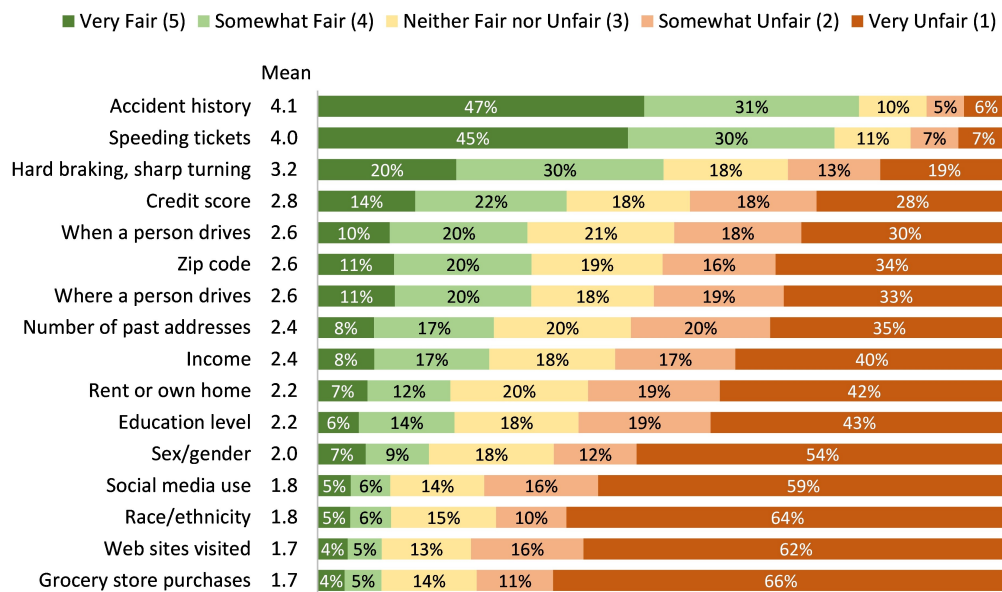
The framing of the scenarios used in this survey leads to three important scope conditions for the findings that follow. First, I specified that the car insurer and lender wanted to use the data for reasons generally understood as legitimate: predicting insurance claims and loan defaults. To the extent the goal of using data is itself morally suspect, Americans may make blunter, negative judgments about data use.[8] Second, I specified that companies wanted to use the data because statistical analysis showed each type of information to be mathematically predictive of the outcome of interest. This likely biased respondents in the direction of evaluating data use as fair. Finally, the scenarios said nothing about the effect using data would have on who received car insurance and loans or how much they would pay. The distribution of resources that results from data use is often fodder for moral claims (e.g., that data use expands access to markets or disproportionately disadvantages racial minorities). I return to the relevance of such arguments in the discussion section below.

Findings

Figures 1 and 2 present the quantitative results. Figure 1 presents findings for car insurance and Figure 2 for lending decisions. The overarching takeaway is that everyday Americans make sharp moral distinctions among the types of data companies use to differentiate. The figures show that Americans, taken as a whole, perceive some sorts of data as fair to use (those toward the top), whereas they view other sorts of data as overwhelmingly unfair (those toward the bottom). Although Americans may broadly endorse differentiation in market contexts, that does not mean all forms of differentiation pass moral muster. People believe that the market rightly differentiates on some—but far from all—grounds. (Tables A1 and A2 in the online supplement show standard errors for the means presented in Figures 1 and 2.)

Notably, data that Americans see as fair to use in one market domain, they often see as unfair to use in another. In some cases, the difference is quite large. Whereas 75 percent of respondents consider car insurers' use of speeding tickets

■ Very Fair (5)   ■ Somewhat Fair (4)   ■ Neither Fair nor Unfair (3)   ■ Somewhat Unfair (2)   ■ Very Unfair (1)

	Mean					
Accident history	4.1	47%	31%	10%	5%	6%
Speeding tickets	4.0	45%	30%	11%	7%	7%
Hard braking, sharp turning	3.2	20%	30%	18%	13%	19%
Credit score	2.8	14%	22%	18%	18%	28%
When a person drives	2.6	10%	20%	21%	18%	30%
Zip code	2.6	11%	20%	19%	16%	34%
Where a person drives	2.6	11%	20%	18%	19%	33%
Number of past addresses	2.4	8%	17%	20%	20%	35%
Income	2.4	8%	17%	18%	17%	40%
Rent or own home	2.2	7%	12%	20%	19%	42%
Education level	2.2	6%	14%	18%	19%	43%
Sex/gender	2.0	7%	9%	18%	12%	54%
Social media use	1.8	5%	6%	14%	16%	59%
Race/ethnicity	1.8	5%	6%	15%	10%	64%
Web sites visited	1.7	4%	5%	13%	16%	62%
Grocery store purchases	1.7	4%	5%	14%	11%	66%

Figure 1: How Americans rate the fairness of companies using various types of data in car insurance decisions. *Notes*: Survey conducted by YouGov for the author February 11 to 14, 2019. $N = 1,095$. Values weighted to be nationally representative.

to be somewhat or very fair, only 33 percent of respondents say the same about lenders using these data. Similarly, whereas 68 percent of respondents think it is somewhat or very fair for lenders to consider a person's credit report or score, only 36 percent say it is somewhat or very fair for car insurance companies to do so. Americans do not think about the fairness of companies using personal data in a generalized way. Rather, fairness dictates that certain data can be used for some, but not other, purposes.

Within each figure, the data can be interpreted as falling into three clusters. Data in the rows near the top and bottom of each figure reflect that Americans broadly agree that some data are *permissible* and other data are *proscribed*. For example, three-quarters of respondents consider it somewhat or very fair for car insurers to use accident history and for lenders to use rent payment history; about just as large a share of respondents judge it somewhat or very *un*fair for car insurers to use grocery store purchases and for lenders to use race and ethnicity.

Yet for many other sorts of data, Americans as a group are far from unified in their opinions. For these *unsettled* data, Americans at times hold strong, but conflicting views. Approximately one-third (36 percent) of respondents say that it is somewhat or very fair for lenders to use the number of addresses a person has had, whereas approximately one-third (37 percent) say such data is somewhat or very *un*fair to use. Other data are even more polarizing. For example, 20 percent of respondents say it is very fair for a car insurer to use data from a device that tracks how a person drives, whereas 19 percent of respondents say using such data is very unfair.[9] Figure 3 plots the variance of responses for each sort of data
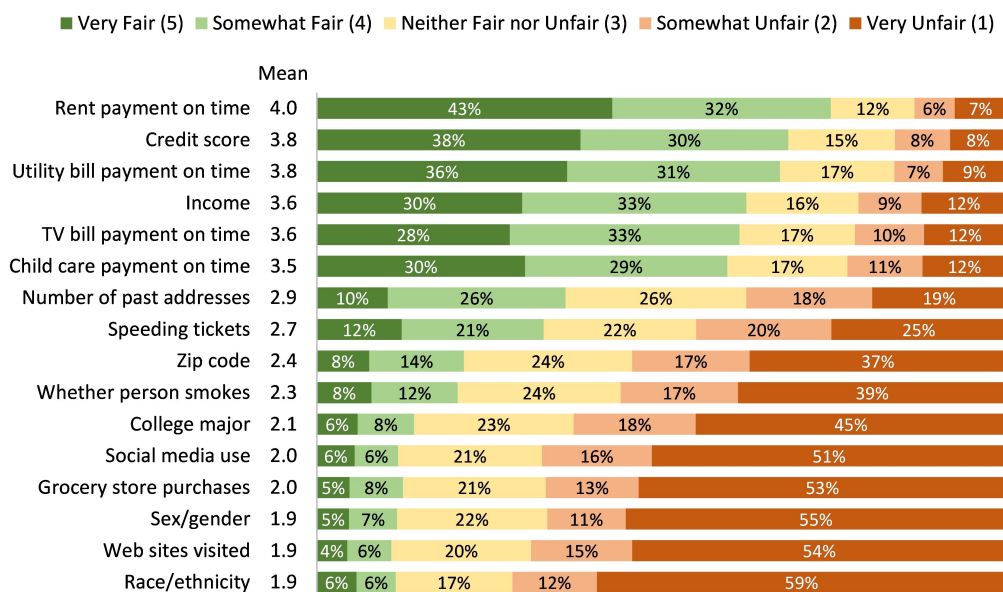
Figure 2: How Americans rate the fairness of companies using various types of data in lending decisions. *Notes*: Survey conducted by YouGov for the author February 11 to 14, 2019. $N = 1,095$. Values weighted to be nationally representative.

in the car insurance question (collapsing together Somewhat and Very Fair, and Somewhat and Very Unfair), and Figure 4 does the same for each sort of data in the lending question. The large variances that correspond to the data types appearing in the middle of each list demonstrate a lack of consensus among respondents about whether these data are fair to use or not.

Importantly, data seeming "behavioral" is not a consistent moral demarcation. Data that capture the ostensibly deliberate actions of individuals can register as either permissible, proscribed, or unsettled. In both Figures 1 and 2, three archetypical examples of behavioral data—web browser history, social media use, and retail (grocery) purchases—appear at the very bottom, alongside race and sex, ascriptive traits that since the 1970s have been largely taboo in allocative decisions and banned from most market transactions.[10] Yet other behavioral markers—smoking, timely bill payment, speeding tickets—appear significantly higher in the figures, with considerably greater proportions of respondents deeming the data as fair to use.

Especially telling is variation across three types of telematics driving data, which are collected from a device in a person's car. Respondents may have generalized concerns about electronic monitoring and the collection of data from personally intimate spaces, but that does not stop them from deeming some data collected in this way as fairer to use than others. Whereas 50 percent of respondents think it is somewhat or very fair for a car insurer to use data about *how* a person drives (whether they slam on the brakes, turn sharply, etc.), only about 30 percent think it is somewhat or very fair for a car insurer to use data about *where* or *when* a person
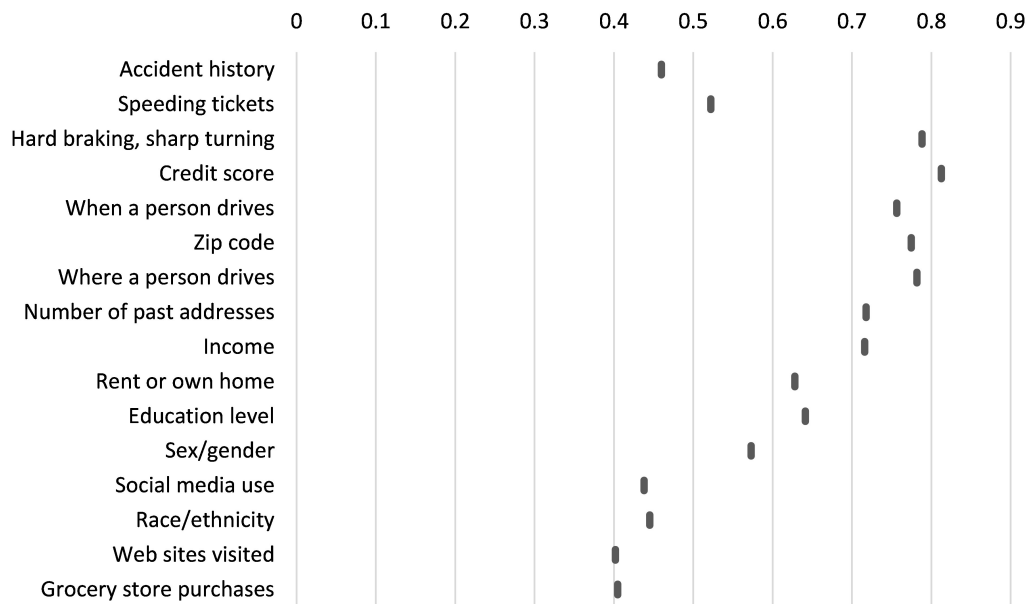
Figure 3: Variances of responses for each sort of data in the car insurance question. *Note*: Collapsed response categories used to calculate variances (Somewhat/Very Unfair; Neither Fair nor Unfair; Somewhat/Very Fair).

drives. Americans think it is fair for companies to hold people accountable for some, but not other, behaviors.

To better understand these patterns in the quantitative data, I turn to respondents' free-text explanations of why they rated certain types of data as fair or unfair to use. In the next three sections of this article, I describe the dominant explanation offered by respondents rating data use as fair; the major hang-up of respondents rating data use as unfair; and how these two dynamics, taken together, account for why so many data types remain unsettled.

Relatedness and Judging Data as Fair to Use

Respondents explained rating data use as fair in a variety of ways, but one explanation appeared much more frequently than the rest: that the data were related (or relevant) to the transaction at hand.[11] To a large extent, this meant that using the data would help the company achieve its goal of predicting either insurance claims or loan defaults. Here, it is important to remember that in both car insurance and lending, using personal data to make predictions about individuals is firmly institutionalized and that these two predictions in particular are broadly taken to be legitimate ones to make. Under this scope condition, data were adjudicated as fair when they were taken to be accurate indicators of the outcome they were meant to predict.

Respondents established relatedness in two ways. The first was mathematical. The survey prompt told respondents that each type of data was statistically linked
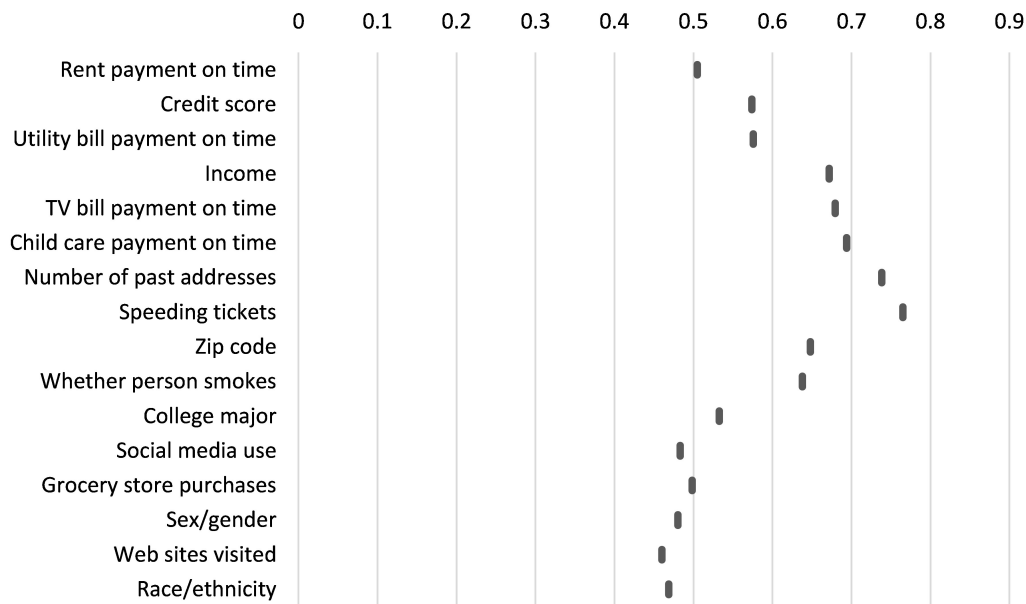
Figure 4: Variances of responses for each sort of data in the lending question. *Note*: Collapsed response categories used to calculate variances (Somewhat/Very Unfair; Neither Fair nor Unfair; Somewhat/Very Fair).

to the outcome of interest, and some respondents referred back to this in justifying data use as fair. For example, one respondent, explaining why it would be fair for a car insurer to use how many addresses a person has had, wrote, "If there is a statistical correlation between how often you move and accident history, this seems like a fair thing to consider." Or, as another wrote about a lender using the same information, "If that's a factor [in] how likely they are to pay money back it's fair."

Yet invocations of mathematical relatedness paled in frequency compared with explanations of *logical* relatedness, by which I mean connections people draw through reasoning about how the world works. Prior research has shown that moral claims rooted in statistical relationships are fragile and can lose legitimacy when not backed up by more intuitive explanations of why two things are related (Kiviat 2019b; Underwood 1979). That dynamic was on strong display as respondents offered up reasons why it made logical sense that various data types would have bearing on insurance claims and loan defaults.

Such logical relatedness took two forms. First, respondents assumed that people would act similarly in new situations that were analogous to ones in which they had already been observed. For example, when asked about a lender using television bill payments, one respondent wrote, "If you don't pay other bills on time it is probably a good indicator if you would pay off a loan. I believe any bill would be fair to use." Bill payment here is seen as similar to loan payment (*other* bills), which makes it logical to expect the same behavior and fair to hold the person accountable.

In the second form of logical relatedness, respondents assumed that past behavior was indicative of a person's internal disposition or character and that this

would carry over to a new situation. For example, one respondent, explaining why it would be fair for a car insurer to consider a person's credit history, wrote, "If a person is responsible in their credit, they're more likely to be a responsible driver." Or, as a respondent explaining the fairness of a lender using number of past addresses wrote, "It show[s] if you are stable or not. A person that do[es] not move a lot is more likely to pay their loan back." In these examples, respondents took data as signals of people's inherent qualities—marks not simply of what people do but of who they *are*.

Two aspects of this are worth noting. First, respondents who assumed that data shed light on a person's disposition often did so with moralized language, using words such as responsible, reliable, trustworthy, reckless, transient, unstable, dangerous, and other morally thick concepts. This is significant because rendering people in these terms inculcates notions of fault, blame, and merit, which lend moral legitimacy to people winding up better or worse off at the hands of the market.

The second aspect worth noting is that a focus on personal disposition, rather than situational similarly, enables data to reach further, into a broader range of new situations. Utility bill and loan payments are similar enough that most respondents simply said that paying one would indicate paying the other. Yet for other sorts of data that were less alike—number of past address, for example—respondents were more likely to turn to justifications about individuals' presumably unchanging personal traits. Using data to essentialize people helps make those data seem fair to use in more distant market settings.

Morally Heterogeneous Categories and Judging Data Use as Unfair

One way that respondents justified evaluating data use as *un*fair was to deem the data unrelated or irrelevant: the mirror image of the response described above. Sometimes respondents pushed back against mathematical relatedness, denying that a correlation actually existed. At other times, respondents challenged that the relationship was logical. As one respondent, discussing the use of credit scores by car insurers, wrote, "If I miss a payment to a credit card, or even have a house foreclosure it doesn't [a]ffect my ability to drive or how safely I drive." Yet, although common, questioning relatedness in broad strokes was not the dominant way respondents justified judging data use as unfair.

Rather, what permeated responses was the sense that data improperly conflated morally distinct situations and behaviors.[12] For example, many respondents who said it would be unfair for lenders or car insurers to consider number of past addresses pointed out that although moving *can* be a red flag (as in the case of eviction), many moves are perfectly legitimate—as when one moves for a better job, to be close to family, to attend college, because of a military transfer, and so on. As one respondent explained, "There are way too many reasons why a person would move (or not) over time... Sometimes moving is good, sometimes not. Since the number of houses a person lived in is not a reliable indicator by itself, I consider it to be unfair." Respondents saw moving as a *morally heterogenous* act, and so simply knowing that a person had changed addresses, with no additional context, could not be fairly linked to market outcomes.

Across data types, one piece of context some respondents felt to be problematically missing was whether a person had been in control of the situation that gave rise to the data. Respondents pointed out that medical bills can lead to lower credit scores, spiking rents can drive people to move, ex-spouses can delay childcare payments, freezing winters can make utility bills unaffordable, and so on. Although data may be construed as reflecting people's choices, respondents acknowledged that choices can be constrained to the point of not really being choices at all, and therefore (they argued) not a fair basis for increasing prices.

One interpretation of concern about morally heterogenous data is that it reflects an inherent tension of statistical prediction: that predictions are about what is true on average for a group of people, and yet those who receive treatment are specific individuals (Barry and Charpentier 2020; Gandy 2009; Schauer 2003). Indeed, at times respondents referred to the fact that although predictions might "work" in the aggregate, they were nonetheless unfair to certain individuals. As one respondent wrote, "[T]he number of times a person has moved may statistically work for the lender (people in general), [but] for the individual it may be totally irrelevant and be grossly unfair. In my youth I moved a lot, but I have never failed [to] pay a loan in my life." Or, as another respondent explained about lenders using the timeliness of childcare payments, "They are surmising your total character based upon one event. While statistically this may work for them, for the individual this may not work at all and be totally unfair."

Yet to conclude that Americans think it unfair to hold individuals accountable for group averages as a general rule elides an important distinction. The problem is not that people have been grouped but that they have been *inappropriately* grouped. When the boundaries of a group capture too many different sorts of people and behavior, holding each individual responsible for the average is likely to assign blame where it does not belong, a sort of moral ecological fallacy.

Consider, for example, differences in how respondents reacted to two types of telematics driving data: information about how often a person slams on the brakes, turns sharply, and accelerates quickly, and information about what time of day or night a person drives. Respondents only occasionally saw moral distinctions in the former (e.g., driving recklessly vs. slamming on the brakes to avoid an accident) but frequently did in the latter. Driving late at night, which respondents presumed to be the "bad" behavior insurers were interested in, was routinely described as potentially legitimate, given the demands of work, elder care, and so on. As one respondent colorfully explained, "A person might have to work 3rd shift or be a drug dealer. Both work at night and shouldn't be judged on the same criteria." What respondents saw as unfair was that people who were meaningfully different might wind up looking the same in the data and therefore receive the same treatment.

Competing Cultural Meanings as the Key to Unsettled Data

Taken all together, these findings show that people assign social and moral meaning to personal data in market settings. In both rhetoric and action, companies treat data as commodities—interchangeable bits that are unproblematically transferred from one situation to the next. Yet to everyday Americans, personal information is not

so easily stripped from broader context. Just as people find ways to imbue money and other seemingly fungible aspects of markets with social and moral significance (Zelizer [1994] 2017), respondents worked to establish what data *meant* (cf. Levy 2013). This is noteworthy because respondents knew the data were mathematically predictive; that is, they had at their disposal a way to morally reason that did not require digging deeper into data's meaning. Nonetheless, that is what most of them did.

This, then, brings us back to Figures 1 and 2. Looking again at the data in the topmost and bottommost rows, those which Americans broadly consider *permissible* or *proscribed*, we can take this consistency of opinion to mean that people generally agree about what the data signify. Accident history is "obviously" related to car insurance claims in a way that what one buys at the grocery store is "obviously" not. In one case, there is a culturally shared understanding for construing accidents as logically and morally connected to claims, and in the other case the toolkit is empty: there is no schema to grab for, and so using the data is overwhelmingly judged as unfair.

The most interesting parts of the figures may then be the middle rows, those in which Americans as a group hold strong but opposing viewpoints. What makes data *unsettled*? The answer, I suggest, is that the broader culture provides fodder for competing social meanings, in terms of either relatedness or moral categorization. A person who moves around a lot may be running from financial obligations or working their way up the corporate ladder. Someone who drives late at night may be bar hopping or heading to the second of two jobs to keep food on the table. When people don't agree on what data mean, then neither will they agree on whether those data are fair to use.

Importantly, social meanings vary by market, and therefore so do fairness assessments. Recall, for example, that respondents think it much fairer for car insurers than lenders to use speeding tickets. This makes sense given that there is a logical connection between tickets and insurance claims—that is to say, a dominant cultural narrative that holds car accidents are caused by individuals being reckless behind the wheel.[13] To justify the fairness of using speeding tickets for lending, respondents must reach further, relying on an essentializing schema that says people who speed are reckless and that a reckless disposition will translate into blowing off loan payments. The connection is more of a stretch, but still thinkable. Not everyone buys into the analogy, but enough do to show Americans remarkably split on whether the data are fair to use.

Discussion

As corporations harness untold amounts of personal data, this article shows that everyday Americans make strong moral distinctions among the types of information firms use to tailor their treatment of individuals. People largely make these distinctions according to whether they see data as *logically related* to the behaviors companies are trying to predict and whether data sort individuals in *morally consistent* ways. This article thereby illustrates how two dynamics long studied by economic sociologists—relational matching and moral categorization—take hold

in justifying and contesting the fairness of a new, increasingly important mode of market allocation. At the same time, perceptions about the fair use of many types of data remain unsettled, with competing ideas about what data mean pulling Americans' moral evaluations in different directions.

Companies may take data to be commodities—portable, interchangeable, and stripped of meaning beyond instrumental utility—but that is far from how consumers see things. This suggests that as the personal data economy continues to develop, discursive battles are likely on the horizon. If the moral validity of this new mode of stratification partly hangs on the social meanings people attach to data, then companies have incentive to work to institutionalize socially acceptable meanings. Indeed, in the policy debate about car insurers using credit scores, that is exactly what happened. When arguments about mathematical relatedness failed to fully convince policymakers that credit scores were fair to use, members of industry repeatedly explained that credit scores were a signal of personal responsibility, which is why they predicted insurance claims and could be legitimately used for that purpose (Kiviat 2019b). Although prior scholarship has suggested that reifying data as behavioral brings an aura of fairness, this article suggests that a clearer path is casting data as capturing morally laden aspects of individuals' dispositions.

While this article offers an important first look at the extent to which everyday Americans think various sorts of personal data can be fairly used to determine market outcomes, future research might usefully include other types of justice concerns. The survey used here offered no information about the procedures used to collect data (e.g., whether individuals would have to give consent), nor did it specify what the distributional outcomes of data use would be (e.g., whether certain groups of people would systematically receive higher or lower prices). A major concern with personal data use, occluded here, is the potential to reinforce patterns of racial disadvantage (Barocas and Selbst 2016; Gandy 2009). One question for future research is how Americans weigh various moral standards: for example, the extent to which logically related, morally homogenous data lose their luster if they are known to lead to unequal outcomes by race or sex—or, in the other direction, if logically *un*related, morally *hetero*genous data don't seem so bad if using them promises to expand the market to previously excluded individuals.

Finally, an important direction for future research is to explore what happens to moral evaluations when data are used for less institutionalized purposes. As companies use increasing amounts of personal data, they also make an increasingly wide range of predictions. Companies want to know who is price sensitive, unlikely to complain when products disappoint, likely to rack up large penalty fees, and so on. Insurers, for instance, want to know who will remain loyal in the face of price increases, and lenders score consumers for profitability in addition to repayment (Jeanningros and McFall 2020; Kiviat 2019a). Many novel predictions prove controversial when the public finds out about them. Given that the predictions discussed in this article are generally seen as legitimate ones to make, the findings may, if anything, overstate how sanguine everyday Americans are on corporate data use.

Notes

1 I define personal data as information about an identifiable person. I use the words *information* and *data* interchangeably.

2 Data privacy and data use are related but distinct. Privacy-unconcerned individuals may not care who reads their social media streams, sees their credit scores, or knows they are divorced but nevertheless think it unfair for such information to be used in deciding whether they receive jobs, get apartments, or are charged more for products.

3 Other surveys ask people what it would take for them to share data with companies (e.g., level of price discount). I do not include those here because participating in a market and thinking it fair are distinct phenomena (cf. Turow, Hennessy, and Draper 2015).

4 The survey research firm YouGov conducted the survey on the author's behalf (February 11 to 14, 2019).

5 "Fair" is one of many moral judgments a person may make. I use it here because fairness is often the main moral claim in discussions of economic practices, including those about car insurance pricing and consumer lending.

6 Each respondent answered three open-ended questions. All respondents were asked about a lender and a car insurer's use of the number of addresses a person has lived at in the past five years. Respondents were also randomly assigned to explain their answer to one of six additional data types. These were, for car insurance, credit score ($n = 182$); data about slamming on the brakes, turning sharply, and accelerating quickly ($n = 183$); and time of day or night one drives ($n = 194$) and, for lending, on-time payment of bills for utilities ($n = 174$), cable TV ($n = 184$), and childcare ($n = 178$).

7 The survey was sampled to be representative of the U.S. adult population (based on the 2016 American Community Survey) by age, gender, race, and education and was weighted to be representative by age, gender, race, educational attainment, region, and political orientation using propensity score weighting.

8 For example, regulators have investigated whether car insurers use personal data to predict price sensitivity rather than risk (Banham 2015).

9 I take inspiration from Swidler's (1986) use of the word unsettled, but use it in a different way, referring not to moments of social transformation but rather to objects around which a population does not hold a consensus view in terms of cultural meaning.

10 In many states, car insurers are allowed to use sex to set prices, although the practice is controversial.

11 Other, less-frequent explanations of why data were fair to use included the information being in the public domain, the data being objective, the company having a right to use the data, the belief that the data signaled particularly bad behavior, the possibility that using the data would incentivize better behavior, and the assumption that certain consumers would financially benefit. Some respondents wrote that as long as consumers consented to data use, it would be fair.

12 Other, less-frequent explanations of why data were unfair to use included privacy concerns, the possibility for data use to disadvantage certain groups of people such as minorities or the poor, potential mistakes in the data, the risk of making insurance unaffordable for certain individuals, the potential for a person to suffer disadvantage twice from the same event, the fact that there are fairer data to use, and the belief that people should be free to make life choices as they see fit without those decisions constraining access to important resources.

13 Gusfield (1981) nicely shows that the idea that delinquent individuals cause car accidents, rather than, say, poor road design or automobile safety features, is a social construction dating back to the 1920s.

References

Almeling, Rene. 2007. "Selling Genes, Selling Gender: Egg Agencies, Sperm Banks, and the Medical Market in Genetic Material." *American Sociological Review* 72(3):319–40. https://doi.org/10.1177/000312240707200301.

Anteby, Michel. 2010. "Markets, Morals, and Practices of Trade: Jurisdictional Disputes in the U.S. Commerce in Cadavers." *Administrative Science Quarterly* 55(4):606–38. https://doi.org/10.2189/asqu.2010.55.4.606.

Anthony, Denise, Celeste Campos-Castillo, and Christine Horne. "Toward a Sociology of Privacy." *Annual Review of Sociology* 43:249–69. https://doi.org/10.1146/annurev-soc-060116-053643.

Avraham, Ronen, Kyle D. Logue, and Daniel Schwarcz. 2014. "Understanding Insurance Antidiscrimination Laws." *Southern California Law Review* 87(2):195–274.

Baker, Tom, and Jonathan Simon. 2002. "Embracing Risk." Pp. 1–26 in *Embracing Risk: The Changing Culture of Insurance and Responsibility*, edited by Tom Baker and Jonathan Simon. University of Chicago Press.

Bandelj, Nina. 2020. "Relational Work in the Economy." *Annual Review of Sociology* 46:15.1–15.22. https://doi.org/10.1146/annurev-soc-121919-054719.

Banham, Russ. 2015. "Price Optimization or Price Discrimination? Regulators Weigh In." *Carrier Management*, May 17, 2015. Accessed September 15, 2020. https://www.carriermanagement.com/features/2015/05/17/139913.htm.

Barocas, Solon, and Andrew D. Selbst. 2016. "Big Data's Disparate Impact." *California Law Review* 104:671–732. https://doi.org/10.15779/Z38BG31.

Barry, Laurence, and Arthur Charpentier. 2020. "Personalization as a Promise: Can Big Data Change the Practice of Insurance?" *Big Data & Society* 7(1). https://doi.org/10.1177/2053951720935143.

Board of Governors of the Federal Reserve System. 2010. *Report to the Congress on Reductions of Consumer Credit Limits Based on Certain Information as to Experience or Transactions of the Consumer*. Washington, DC: Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/BoardDocs/RptCongress/creditcard/2009/consumercreditreductions.pdf.

Boyle, Ryan James. 2016. "Can Address Stability Predict Credit Risk?" *Banking Exchange*, December 9, 2016. Accessed November 19, 2020. https://m.bankingexchange.com/credit-risk/item/6601-can-address-stability-predict-credit-risk.

Capon, Noel. 1982. "Credit Scoring Systems: A Critical Analysis." *Journal of Marketing* 46(2):82–91. https://doi.org/10.2307/3203343.

Cohen, Julie. 2013. "What Privacy Is For." *Harvard Law Review* 126(7):1904–33. https://harvardlawreview.org/2013/05/what-privacy-is-for/.

Consumer Financial Protection Bureau. 2014. *Consumer Credit Reports: A Study of Medical and Non-medical Collections*. Washington, DC: Consumer Financial Protection Bureau.

Della Fave, L. Richard. 1980. "The Meek Shall Not Inherit the Earth: Self-Evaluation and the Legitimacy of Stratification." *American Sociological Review* 45(6):955–71. https://doi.org/10.2307/2094912.

Deville, Joe, and Lonneke van der Velden. 2016. "Seeing the Invisible Algorithm: The Practical Politics of Tracking the Credit Trackers." *Algorithmic Life: Calculative Devices in the Age of Big Data*, edited by Louise Amoore and Volha Piotukh. London: Routledge.

DiMaggio, Paul. 1997. "Culture and Cognition." *Annual Review of Sociology* 23(1):263–87. https://doi.org/10.1146/annurev.soc.23.1.263.

Experian. 2019. *State of Alternative Credit Data*. White paper, Experian. Accessed November 19, 2020. https://www.experian.com/blogs/insights/2019/05/state-of-alternative-credit-data/.

Florida Office of Insurance Regulation. 2007. *The Use of Occupation and Education as Underwriting/Rating Factors for Private Passenger Automobile Insurance*. Tallahassee, FL: Florida Office of Insurance Regulation.

Fourcade, Marion. 2016. "Ordinalization: Lewis A. Coser Memorial Award for Theoretical Agenda Setting 2014." *Sociological Theory* 34(3):175–95. https://doi.org/10.1177/0735275116665876.

Fourcade, Marion, and Kieran Healy. 2007. "Moral Views of Market Society." *Annual Review of Sociology* 33:285–311. https://doi.org/10.1146/annurev.soc.33.040406.131642.

Fourcade, Marion, and Kieran Healy. 2013. "Classification Situations: Life-Chances in the Neoliberal Era." *Accounting, Organizations and Society* 38(8):559–72. https://doi.org/10.1016/j.aos.2013.11.002.

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. https://doi.org/10.1093/ser/mww033.

Gandy, Oscar H., Jr. 2009. *Coming to Terms with Chance: Engaging Rational Discrimination and Cumulative Disadvantage*. Burlington, VT: Ashgate Publishing Company.

Gusfield, Joseph R. 1981. *The Culture of Public Problems: Drinking-Driving and the Symbolic Order*. University of Chicago Press.

Healy, Kieran. 2006. *Last Best Gifts: Altruism and the Market for Human Blood and Organs*. University of Chicago Press.

Heller, Douglas, and Michelle Styczynski. 2016. *Major Auto Insurers Raise Rates Based on Economic Factors: Low- and Moderate-Income Drivers Charged Higher Premiums*. Washington, DC: Consumer Federation of America.

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(2):355–412.

Hochschild, Jennifer L. 1981. *What's Fair? American Beliefs about Distributive Justice.* Cambridge, MA: Harvard University Press.

Homans, George C. 1974 [1961]. *Social Behavior: Its Elementary Forms*. Revised edition. New York: Harcourt Brace Jovanovich.

Jeanningros, Hugo, and Liz McFall. 2020. "The Value of Sharing: Branding and Behaviour in a Life and Health Insurance Company." *Big Data & Society* 7(2). https://doi.org/10.1177/2053951720950350.

Karapiperis, Dimitris, Birny Birnbaum, Aaron Brandenburg, Sandra Castagna, Allen Greenberg, Robin Harbage, and Anne Obersteadt. 2015. *Usage-Based Insurance and Vehicle Telematics: Insurance Market and Regulatory Implications*. Kansas City, MO: Center for Insurance Policy and Research, National Association of Insurance Commissioners.

Kelley, Jonathan, and M. D. R. Evans. 1993. "The Legitimation of Inequality: Occupational Earnings in Nine Nations." *American Journal of Sociology* 99(1):75–125. `https://doi.org/10.1086/230230`.

Kiviat, Barbara. 2019a. "Credit Scoring in the United States." *Economic Sociology: The European Electronic Newsletter* 21(1):33–42.

Kiviat, Barbara. 2019b. "The Moral Limits of Predictive Practices: The Case of Credit-Based Insurance Scores." *American Sociological Review* 84(6):1134–58. `https://doi.org/10.1177/0003122419884917`.

Kluegel, James R., and Eliot R. Smith. 1986. *Beliefs about Inequality: Americans' Views of What Is and What Ought to Be*. New York: Aldine de Gruyter.

Krippner, Greta R. 2017. "Democracy of Credit: Ownership and the Politics of Credit Access in Late Twentieth-Century America." *American Journal of Sociology* 123(1):1–47. `https://doi.org/10.1086/692274`.

Levy, Karen E. C. 2013. "Relational Big Data." *Stanford Law Review* 66:73–9. `https://www.stanfordlawreview.org/online/privacy-and-big-data-relational-big-data/`.

Lizardo, Omar, Robert Mowry, Brandon Sepulvado, Dustin S. Stoltz, Marshall A. Taylor, Justin Van Ness, and Michael Wood. 2016. "What are Dual Process Models? Implications for Cultural Analysis in Sociology." *Sociological Theory* 34(4):287–310. `https://doi.org/10.1177/0735275116675900`.

Massengill, Rebekah P., and Amy Reynolds. 2010. "Moral Discourse in Economic Contexts." Pp. 485–501 in *Handbook of the Sociology of Morality*, edited by Steven Hitlin and Stephen Vaisey. New York: Springer.

Martin, Kirsten, and Helen Nissenbaum. 2016. "Measuring Privacy: An Empirical Test Using Context to Expose Confounding Variables." *Columbia Science & Technology Law Review* 18(1):176–218. `http://dx.doi.org/10.2139/ssrn.2709584`.

McCall, Leslie. 2013. *The Undeserving Rich: American Beliefs about Inequality, Opportunity, and Redistribution*. Cambridge University Press.

Moor, Liz, and Celia Lury. 2018. "Price and the Person: Markets, Discrimination, and Personhood." *Journal of Cultural Economy* 11(6):501–13. `https://doi.org/10.1080/17530350.2018.1481878`.

National Association of Insurance Commissioners. 2011. *A Consumer's Guide to Auto Insurance*. Kansas City, MO: National Association of Insurance Commissioners. `URLhere`

Noble, Safiya Umoja. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York University Press.

O'Leary, Patrick F., Randy R. Richards, and Thomas J. Quinlan, Jr. 2013. "Credit-Based Insurance Scoring." *International Journal of Business, Humanities and Technology* 3(7):11–20.

Polletta, Francesca, and Zaibu Tufail. 2014. "The Moral Obligations of Some Debts." *Sociological Forum* 29(1):1–28.

Quinn, Sarah. 2008. "The Transformation of Morals in Markets: Death, Benefits, and the Exchange of Life Insurance Policies." *American Journal of Sociology* 114(3):738–80. `https://doi.org/10.1086/592861`.

Robinson, David, and Harlan Yu. 2014. *Knowing the Score: New Data, Underwriting, and Marketing in the Consumer Credit Marketplace. A Guide for Financial Inclusion Stakeholders*. Washington, DC: Upturn. Accessed November 19, 2020. `https://www.upturn.org/static/files/Knowing_the_Score_Oct_2014_v1_1.pdf`.

Rubinstein, David. 1988. "The Concept of Justice in Sociology." *Theory and Society* 17(4):527–50. `https://doi.org/10.1007/BF00158887`.

Rule, James B. 2009. *Privacy in Peril: How We Are Sacrificing a Fundamental Right in Exchange for Security and Convenience.* Oxford University Press.

Sauder, Michael. 2020. "A Sociology of Luck." *Sociological Theory* 38(3):193–216. https://doi.org/10.1177/0735275120941178.

Schauer, Frederick. 2003. *Profiles, Probabilities, and Stereotypes.* Cambridge, MA: Harvard University Press.

Seaver, Nick. 2015. "The Nice Thing about Context Is That Everyone Has It." *Media, Culture and Society* 37(7):1101–09. https://doi.org/10.1177/0163443715594102.

Shepelak, Norma J. 1987. "The Role of Self-Explanations and Self-Evaluations in Legitimating Inequality." *American Sociological Review* 52(4):495–503. https://doi.org/10.2307/2095294.

Sherman, Rachel. 2017. *Uneasy Street: The Anxieties of Affluence.* Princeton University Press.

Smith, Aaron. 2018. "Public Attitudes toward Computer Algorithms." Pew Research Center. Accessed November 12, 2019. https://www.pewresearch.org/internet/2018/11/16/public-attitudes-toward-computer-algorithms/.

Spillman, Lyn. 1999. "Enriching Exchange: Cultural Dimensions of Markets." *American Journal of Economics and Sociology* 58(4):1047–71. https://doi.org/10.1111/j.1536-7150.1999.tb03407.x.

Starr, Paul. 1992. "Social Categories and Claims in the Liberal State." *Social Research* 59(2):263–95.

State of New Jersey Department of Banking and Insurance. 2008. *The Use of Occupation and Education Factors in Automobile Insurance.* Trenton, NJ: Department of Banking and Insurance, State of New Jersey. Accessed September 15, 2020. http://www.state.nj.us/dobi/division_insurance/pdfs/ed_occ_april2008.pdf.

Steensland, Brian. 2006. "Cultural Categories and the American Welfare State: The Case of Guaranteed Income Policy." *American Journal of Sociology* 111(5):1273–326. https://doi.org/10.1086/499508.

Swidler, Ann. 1986. "Culture in Action: Symbols and Strategies." *American Sociological Review* 51(2):273–86. https://doi.org/10.2307/2095521.

Stone, Deborah A. 1993. "The Struggle for the Soul of Health Insurance." *Journal of Health Politics, Policy and Law* 18(2):287–317. https://doi.org/10.1215/03616878-18-2-287.

Task Force on Financial Technology. 2019. *Examining the Use of Alternative Data in Underwriting and Credit Scoring to Expand Access to Credit.* Hearing before the House Financial Services Comm., 116th Congr., 1st Sess.

Teegardin, Carrie. 2008. "Credit Crisis: Confusing Restrictions: Consumer Profiling? His Card's Limit Was Cut over Where He Shopped." *Atlanta-Journal Constitution*, December 21.

Thompson, E. P. 1971. "The Moral Economy of the English Crowd in the Eighteenth Century." *Past and Present* 50(1):76–136.

Turow, Joseph, Michael Hennessy, and Nora Draper. 2015. "The Trade-Off Fallacy: How Marketers are Misrepresenting American Consumers and Opening Them Up to Exploitation." Annenberg School for Communication. Accessed September 16, 2020. https://www.asc.upenn.edu/sites/default/files/TradeoffFallacy_1.pdf.

Underwood, Barbara D. 1979. "Law and the Crystal Ball: Predicting Behavior with Statistical Inference and Individualized Judgment." *Yale Law Journal* 88(7):1408–48.

Vaisey, Stephen. 2009. "Motivation and Justification: A Dual-Process Model of Culture in Action." *American Journal of Sociology* 114(6):1675–715. https://doi.org/10.1086/597179.

Velthuis, Olav. 2005. *Talking Prices: Symbolic Meanings of Prices on the Market for Contemporary Art*. Princeton University Press.

Watkins-Hayes, Celeste, and Elyse Kovalsky. 2016. "The Discourse of Deservingness: Morality and Dilemmas of Poverty Relief in Debate and Practice." Pp. 193–220 in *The Oxford Handbook of the Social Science of Poverty*, edited by David Brady and Linda M. Burton. Oxford University Press.

Wherry, Frederick F. 2016. "Relational Accounting: A Cultural Approach." *American Journal of Cultural Sociology* 4(2):131–56. https://doi.org/10.1057/ajcs.2016.1.

Williams, Betsy Anne, Catherine F. Brooks, and Yotam Shmargad. 2018. "How Algorithms Discriminate Based on Data They Lack: Challenges, Solutions, and Policy Implications." *Journal of Information Policy* 8:78–115. https://doi.org/10.5325/jinfopoli.8.2018.0078.

Zelizer, Viviana A. 1979. *Morals and Markets: The Development of Life Insurance in the United States*. New York: Columbia University Press.

Zelizer, Viviana A. 1996. "Payments and Social Ties." *Sociological Forum* 11(3):481–95. https://doi.org/10.1007/BF02408389.

Zelizer, Viviana A. 2012. "How I Became a Relational Economic Sociologist and What Does That Mean?" *Politics and Society* 40(2):145–74. https://doi.org/10.1177/0032329212441591.

Zelizer, Viviana A. [1994] 2017. *The Social Meaning of Money*. New York: Basic Books.

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

Zwick, Detlev, and Janice Denegri Knott. 2009. "Manufacturing Consumers: The Database as New Means of Production." *Journal of Consumer Culture* 9(2):221–47. https://doi.org/10.1177/1469540509104375.

Barbara Kiviat: Department of Sociology, Stanford University.
E-mail: bkiviat@stanford.edu.