

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

**Technological Innovation and Discrimination in Household
Finance**

Adair Morse and Karen Pence

2020-018

Please cite this paper as:

Morse, Adair, and Karen Pence (2020). “Technological Innovation and Discrimination in Household Finance,” Finance and Economics Discussion Series 2020-018. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2020.018>.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

TECHNOLOGICAL INNOVATION AND DISCRIMINATION IN HOUSEHOLD FINANCE

ADAIR MORSE

University of California, Berkeley and NBER

KAREN PENCE¹

Federal Reserve Board

February 2020

Abstract

Technology has changed how discrimination manifests itself in financial services. Replacing human discretion with algorithms in decision-making roles reduces taste-based discrimination, and new modeling techniques have expanded access to financial services to households who were previously excluded from these markets. However, algorithms can exhibit bias from human involvement in the development process, and their opacity and complexity can facilitate statistical discrimination inconsistent with antidiscrimination laws in several aspects of financial services provision, including advertising, pricing, and credit-risk assessment. In this chapter, we provide a new amalgamation and analysis of these developments, identifying five gateways whereby technology induces discrimination to creep into financial services. We also consider how these technological changes in finance intersect with existing discrimination and data privacy laws, leading to our contribution of four frontlines of regulation. Our analysis concludes that the net effect of innovation in technological finance on discrimination is ambiguous and depends on the future choices made by policymakers, the courts, and firms.

Keywords: discrimination, fair lending, statistical discrimination, FinTech, taste-based preferences, algorithmic decision-making, proxy variables, big data

¹ This paper was prepared as a chapter for the *Palgrave Handbook of Technological Finance*. The views in this chapter are not necessarily those of the Federal Reserve Board or its staff. Any errors are the sole responsibility of the authors. We are grateful to Katrina Blodgett, Carol Evans, and Varda Hussain for their careful review of the chapter and thank Bobby Bartlett, Tobias Berg, David Cross, Tim Lambert, Miles Larbey, David Palmer, Raghavendra Rau, Bradley Schnarr, Christopher Shelton, Luigi Zingales, and seminar participants at the University of Virginia for helpful comments and conversations.

I. Introduction

Technology has the potential to be a tremendous anti-discriminatory force in household finance. Technology can displace human discretion in decision-making—a traditional source of bias and discrimination—with more objective decisions based on data and algorithms. Technology can bring access to financial services to households who were previously excluded from these markets because of financial or geographic isolation; this exclusion itself can stem from discrimination. Technology can reduce the costs associated with extending credit and make it profitable to extend credit to households out of the mainstream majority community. Technology can increase access to information about the availability and pricing of financial products and services and thereby level the information playing field.

However, a positive outcome for the role of technology in mitigating discrimination is not pre-ordained. Technology inherently has no animus, but is not immune from being discriminatory. Some households have more access to and more facility with technology than others, such that even innovations may reinforce patterns of exclusion. Technology can allow firms to target advertising and product offers very precisely to consumers, raising the possibility that households have different information sets and even face different prices, sometimes in breach of fair lending, equality, public accommodation, and civil rights laws. Decision-making by financial services providers via algorithms incorporates thousands of variables and presents courts, policy makers, and regulators with complex questions about how to think about and detect discrimination. And although algorithms have no inherent bias, they can incorporate the biases embedded in the broader culture through the datasets used in their development and through the biases of their development teams.

This chapter brings together these benefits and friction points of how technology in finance can affect discrimination. The evidence indicates that technology is a powerful force for reducing discrimination stemming from human discretion (“taste-based” discrimination in economics parlance). But the net effect of technology as an abater of discrimination, especially looking into the future, is not obvious, and depends heavily on resolutions to legal and regulatory uncertainties surrounding the use of algorithms in what economists call “statistical” discrimination. If whether technology is net positive or negative for discrimination is the thread woven through our chapter,

the overall tapestry of our contribution lies in identifying the technological implementations that can lead to discrimination, particularly focusing on the interactions between financial service providers and households in human discretion, algorithmic decision-making, and innovation and inclusion. We identify five such gateways for discrimination: (i) human involvement in designing and coding algorithms, (ii) biases embedded in training datasets, (iii) practices of scoring customers for creditworthiness based on variables that proxy for membership in a protected class, especially through digital footprint and mobile data, (iv) practices of statistical discrimination for profiling shopping behavior, and (v) practices of technology-facilitated advertising, including ad targeting and ad delivery.

Within these implementations of technology, we further identify four regulatory “frontlines.” We use the term frontline to connote two sentiments—a situation of uncertainty (in particular, as to whether the regulatory status quo will remain) and a setting of potential conflict (as legal protections of individuals confront forces of business use of technology). How these legal and regulatory frontlines are resolved will affect whether technology is on net positive or negative in the long run for discrimination.

Our regulation frontlines focus on regulatory uncertainty concerning: (i) whether a variable is “correlated enough” with a protected class to be discriminatory itself, (ii) the use of input-based enforcement of large dataset algorithms versus output-based compliance, (iii) the extent to which privacy laws restrict algorithmic provision of financial services, and (iv) the applicability of public accommodation laws (also called equality laws) to disparities in access to online and mobile provision of financial services. These points of uncertainty affect not just legal tensions of how regulators and courts will act, but also the ability of financial service providers to innovate and the incidence of the benefits of innovation to consumers.

Our chapter builds heavily on the works of other scholars that examine various settings or specific aspects of discrimination. We highlight these works as we proceed. Our contribution is in the amalgamation and analysis of ideas toward understanding the gateways of discrimination entering technological finance, the frontlines of regulation, and the weights for and against technology as

an abater of discrimination in financial services. In the process, we gain the insight of just how dramatically technology has changed the way discrimination manifests itself in financial services.

II. Views of Discrimination: Lawyers and Economists

The United States has comprehensive federal laws prohibiting discrimination in lending, and a patchwork of state and federal laws that cover, less comprehensively, other financial services.² Evans (2017) provides a review of these laws with an emphasis on their implications for fintech firms. In the United Kingdom (UK) and Canada, discrimination in financial services is prohibited under broader anti-discrimination laws.³ The European Union (EU) recognizes non-discrimination as a fundamental right, but relegates specific legislation to member states, who in turn vary in their attention to discrimination legislation and enforcement. In practice, financial service providers have had more freedom in continental Europe to use protected characteristics for profiling, but this is changing, as the EU and UK take a leadership role in regulating the use of technology in finance.

Discrimination laws cover a varying set of protected classes. Individuals are usually safeguarded against discrimination based on race, ethnicity, religion, marital or family status, and disability, and sometimes on additional characteristics such as age, gender, national origin, sexual orientation, gender reassignment, political views, genetic or biometric information, veteran status, and use of social safety nets.

Discrimination laws in principle cover all the steps and practices involved in offering a financial service. However, the enforcement by which this principle is carried out varies from country to country on at least two dimensions. First, countries may differ in the intensity of the focus on the steps. The U.S. antidiscrimination laws and their implementing regulations generally have greater specificity as to the steps and practices covered; in fair lending, for instance, a lender in the U.S. cannot discriminate in advertising, credit risk assessment, and pricing of a loan. The UK's broad-based laws have some individual requirements and carve-outs but in general have less specificity

² The main fair lending laws in the United States are the Equal Credit Opportunity Act and the Fair Housing Act.

³ The Equality Act 2010 prohibits discrimination in the provision of services in the UK. See Hale (2018) for a discussion of equality under the law in the UK. The Canadian Human Rights Act prohibits discrimination at the federal level in Canada. Discrimination may also be regulated at the provincial level for financial service providers that operate in only one Canadian province or territory.

than the U.S.⁴ Second, countries may differ in the specificity of sectors to which antidiscrimination laws apply. Again, the U.S. code is more specifically written, delineating housing, credit, and employment as sectors with particularly detailed regulations. An advantage to specificity in preventing and enforcing is the attention to the particulars for the steps and sectors listed in the laws. An advantage to generality is in flexibility to expand to considering new steps and sectors as the provision of financial services expands and changes with technology. Of course, the question in the more general case—a question being played out across the different country jurisdictions in the EU—is whether a country will delve into discrimination compliance within the steps of provision if the law does not explicitly say so.

Finally, all discrimination laws speak to direct discrimination—treating individuals differently on the basis of protected characteristics such as race, ethnicity, or gender. This is called *disparate treatment* in the U.S. U.S. regulators distinguish between “overt evidence” of disparate treatment, when a lender openly discriminates on a prohibited basis, and “comparative evidence” of disparate treatment, when a lender treats an applicant differently based on a prohibited basis.⁵ Comparative evidence can encompass treating individuals differently on the basis of variables that are highly correlated with a prohibited basis. In the U.S., variables such as grey hair (for age in employment decisions) or targeted zones of zip codes (for minority neighborhoods in credit decisions) fall under this category and are generally considered direct discrimination, since these variables are masked versions of the original prohibited basis.⁶ Some discrimination laws also encompass indirect discrimination, which is when a policy or practice that on its face seems neutral disadvantages a protected group indirectly through other variables. The Australian Human Rights Commission provides the example of a public building being only accessible by stairs as representing indirect discrimination against people with disabilities. This example is *indirect* because the lack of accessibility of a building is presumably not done to preclude people with disabilities but rather to save on costs.⁷ In the U.S., the term *disparate impact* maps roughly to indirect discrimination.

⁴ This may be changing, as a new regulation took effect in June 2019 that forbids any advertising that includes gender stereotypes that are likely to cause harm (Safronova, 2019).

⁵ FDIC Consumer Compliance Examination Manual, Fair Lending Laws and Regulations, IV – 1.1-1.2, September 2015, <https://www.fdic.gov/regulations/compliance/manual/4/iv-1.1.pdf>.

⁶ Rothstein (2017) discusses the historical roots of redlining. Pop (2013) describes a similar debate that unfolded in Germany.

⁷ <https://www.humanrights.gov.au/quick-guide/12049>

In contrast to legal views of discrimination, economists view discrimination through the lens of whether it is “taste-based” or “statistical.” Under taste-based discrimination (Becker, 1957), decision-makers get utility from engaging in prejudice, and are willing to sacrifice other priorities—such as hiring the most productive workers possible—in order to satisfy their biases. The much-cited culmination of Becker’s theory is that taste-based discrimination cannot persist because other employers, who do not have a taste for prejudice, will hire workers based solely on their productivity. These non-discriminating firms will be more profitable than their prejudiced competitors, and the prejudiced firms will go out of business. As we discuss in section *III*, this culmination may not play out in practice.

Under statistical discrimination (Arrow, 1973; Phelps, 1972), discrimination results from the practice of using variables as statistical discriminants to uncover unobserved variables. There are two crucial differences between statistical and taste-based discrimination for our purposes.

First, statistical discrimination does not require employers or other decision-makers to have animus or negative taste toward a protected category (non-whites in the Beckerian formulation). Rather, decision-makers engage in statistical discrimination because they are missing information on a characteristic that is key to their decision, such as credit risk in the case of lending. In the formative theory models, a lender that lacks such information may try to recover proxies for credit risk by using the average credit risk of a group, where the group is defined by gender, race, ethnicity, or other characteristic. In practice, applying the averages by protected groups is illegal, but lenders may use other variables that correlate with a protected category to recover credit risk, and thereby implement statistical discrimination in a more general way than the original theories.

Second, statistical discrimination, unlike taste-based discrimination, is profit-maximizing for the financial service providers. This finding implies that the target for using statistical determinants from the firm (and economist) perspective is profits, not uncovering the unobserved component of credit risk. We discuss in section *IVa* how this economists’ concept of statistical discrimination sometimes misaligns with the legal view. As a preview here, we note that the possibility that

financial service providers could be illegally discriminating while profit maximizing is an uncomfortable juxtaposition of the economist's view of discrimination with the law.

III. Human Decisions and Discretion

III.a. Discriminatory Discretion Ameliorated by Technology

Historically, lenders exhibit patterns in providing financial services that appear consistent with taste-based discrimination against certain types of individuals, even when acting on these biases has resulted in lower profits (Charles and Hurst, 2002; Bayer, Ferreira, and Ross, 2017; Alesina, Lotti, and Mistrulli, 2013; Deku, Kara, and Molyneux 2016; Dobbie, Liberman, Paravisini, and Pathania, 2018; and Bartlett, Morse, Stanton, and Wallace, 2019). In Becker's theory, taste-based discrimination is competed away by market forces. However, if the market is not fully competitive, or if the foregone profits associated with employees who discriminate are fairly small, taste-based discrimination can persist. This type of discrimination is particularly likely to emerge in settings where decision-makers have discretion.

Technology has the potential to limit discretionary discrimination by providing information about financial services more broadly and at lower cost and by limiting the face-to-face interactions that appear to facilitate discrimination. Scott Morton, Zettelmeyer, and Silva-Russo (2003), for example, found that Black and Latinx car purchasers paid more than white purchasers when the sales negotiation took place in person but not when it occurred on the Internet.

When humans are removed fully from negotiations, the decision-making becomes algorithmic, which has been found to reduce costly discriminatory discretion in many settings. For instance, Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan (2018) show that a machine learning algorithm outperforms human judges in predicting which defendants will skip their next court appearance or commit crimes while out on bail, and does so without increasing racial disparities in the probability of being released on bail. In the realm of lending, Dobbie, Liberman, Paravisini, and Pathania (2018) show that a high-cost lender in the UK would increase profits and reduce bias if it used a machine-learning based algorithm to make lending decisions instead of relying on the judgment of lending examiners. Bartlett et al. (2019) find that the discrepancy between the rates

charged to white and Black/Latinx borrowers is lower for algorithmic lenders than conventional lenders, and that algorithmic lenders show no disparities in mortgage rejection rates.

III.b. Discriminatory Discretion Enabled by Technology

Yet technology does not always remove discretion, because humans remain involved in technology processes—either through the role of peer input in platforms or through coding. The new modes of human involvement in technology-provided financial services shift some of the fault lines of discrimination from banking and loan officer discretion to coding and the crowd.

Peer-to-peer platforms provide a setting where humans remain involved in a technology-induced market. Just as in brick-and-mortar lending, when online platforms are information portals for human decision-making (by the peer investors) and include pictures and names, discretion biases decisions toward taste preferences and in-group biases. For instance, Edelman, Luca, and Sverisky (2017) find that AirBnB applicants with distinctively African-American names are less likely to be approved. In the lending context, Ravina (2019) studies the platform prosper.com’s use of the photos of borrower. She finds that attractive borrowers are more likely to get a loan and more likely to default. Online platforms have little incentive to rein in this behavior because providing pictures and names appears to give consumers more trust in participating in transactions with strangers (Edelman and Luca, 2014, Doleac and Stein, 2013) and because the Communications Decency Act may shield the providers from liability (Edelman and Stemler, 2018).

Coding and data inputs are another avenue through which human discretion can lead to discrimination in technological finance. Algorithms are ultimately designed by humans, even if the mechanics are handled by artificial intelligence. How programmers set up the optimization problem, classify the data, and choose the training data sets can influence the output of the algorithm (Barocas and Selbst, 2016). In 2017, for example, Amazon scrapped an artificial intelligence tool for reviewing resumes because it systematically discriminated against women. The tool was trained on the resumes of past job applicants, most of which were men, and so it learned to penalize words such as “women’s” (Dastin, 2018). As is seen in this example, training data sets can facilitate discrimination either because they embed the existing prejudices in society or because they are not representative of the broader population. The stark under-representation of

Blacks, Latinx, and women among programmers means that algorithms are unlikely to be designed by diverse teams that might spot some of these issues (Lee, 2018).

Discrimination Gateway #1: Human involvement in designing and coding algorithms can lead to discrimination.

Caliskan, Bryson, and Narayanan (2017) provide a systematic study of the prejudice that can result from biased data. They find that word-association algorithms trained on the text from the Internet were more likely to associate European-American names with pleasant attributes such as “health” and “honest” and African-American names with unpleasant attributes such as “poverty” and “failure.” The algorithms were also more likely to associate male names with words about careers and female names with words about families, and to associate certain types of occupations with gender. They note, for example, that at the time of their paper, Google translated “O bir doktor. O bir hemşire.” as “He is a doctor. She is a nurse.” despite the fact that Turkish pronoun *O* has no gender.

Other algorithms may perform poorly because they are trained on non-representative data. Buolomwini and Gehr (2018) show that facial recognition software, which is developed using machine learning techniques, is less accurate in classifying the gender of females and of dark-skinned individuals. The three classifiers that they examined had error rates of less than 1 percent for light-skinned male faces and from 20 percent to 34 percent for dark-skinned female faces. The authors note that darker-skinned faces appear to be less represented in some datasets that are used to train the software. They also cite the finding of Roth (2009) that default camera settings are set to expose lighter-skinned faces. Similarly, Blodgett and O’Connor (2017) show that language-identification software is less accurate in identifying as English shorter-length Tweets written in the dialect of American English referred to in their paper as African-American English.

Discrimination Gateway #2: Biases can be embedded in training datasets.

IV. Algorithmic Decisions

By reducing the role of human discretion, algorithmic decision-making has the potential to decrease discrimination in financial services. But algorithmic decision-making may also amplify discrimination because the opacity and complexity of the algorithms may mask the fact that the

use of variables can cause disparities in outcomes against protected groups that is illegal under some antidiscrimination laws. In this section, we focus on three interactions between households and financial service providers—credit risk assessment, pricing of financial services, and advertising—where the introduction of new algorithmic modeling techniques may cause disparities not justified in antidiscrimination laws.

IV.a. Statistical Discrimination

In this section, we bridge the economists’ concept of statistical discrimination with the legal framework. Statistical discrimination, in the eyes of economists, is a solution to a signal extraction problem, where the signal provides data on a fundamental skill or attribute that is otherwise unobservable. In the eyes of the law, businesses are allowed to use proxies for these unobserved factors if it can be justified as a legitimate business necessity and if the use of these proxy variables does not have a disproportionate effect on individuals in a protected category.

Statistical discrimination can help firms model unobserved variables that are key to their decision-making. For example, in the context of credit risk scoring, recent evidence in Pope and Sydnor (2011b) and Ravina (2019) demonstrates that unobserved fundamental credit risk is correlated with race even controlling for modern-day, sophisticated observable measures of credit risk. This finding suggests that if proxy variables correlated with hidden credit-risk fundamentals were available, statistical discrimination might enable lenders to do a better job in modelling credit risk.

Yet depending on the legal setting, profiling individuals, even without taste-based intent to discriminate, with proxy variables that results in disparate outcomes by protected categories is illegal under two conditions and may be ruled in breach (depending on jurisdiction) in another three situations.⁸ We refer to these as the *illegal practices* and *frontiers* in the discussion below.

Practices Aimed at Statistical Discrimination that are Illegal

First, the use of a protected category (in most jurisdictions) or a variable highly correlated with the protected category (in many jurisdictions) as a statistical determinant is illegal, even if the variable

⁸ We use the term statistical discrimination for these activities or outcomes, but some prefer to use this term only for legal use of variables to profile individuals according to the law.

is also correlated with the unobservable target of the statistical discrimination application. (The term “target” refers to the business necessity motive—creditworthiness in finance, flight risk in bail setting, or productivity in labor decisions.) This rule is clear enough on the use of the protected category variable, but the threshold of what is “highly correlated” is ambiguous and defined in a context-specific application by courts or regulators. For example, in February 2019, the New York City Commission on Human Rights issued a prohibition on businesses using hairstyles to sort individuals for decisions, because of a high correlation with race (Stowe, 2019). Furthermore, the concept of what is highly correlated becomes increasingly complex as the number of variables used in an algorithm expands with technology. Because of this ongoing tension, we present this concern as our first regulatory uncertainty frontline:

Regulatory Frontline #1: Uncertainty exists in the extent to which courts, regulators, and policy makers will expand the list of highly correlated variables for exclusion in statistical discrimination in an era where it is easier both for financial service providers to amass such variables and for compliance officers and plaintiffs to calculate statistically significant correlations.

Second, practices aimed at statistical discrimination are not legal in the U.S. if the target motivating the statistical discrimination is not defined as a court-justified motive. For lending, U.S. courts have been explicit in ruling that the target is credit risk assessment and that profit motives beyond credit risk are not legal reasons for statistical discrimination. In particular, Bartlett et al. (2019) highlight three cases stating this precedent:

In *A.B. & S. Auto Service, Inc. v. South Shore Bank of Chicago*, 962 F. Supp. 1056 (N.D. Ill. 1997), the ruling reads: “...the defendant-lender must demonstrate that any policy, procedure, or practice has a manifest relationship to the creditworthiness of the applicant...”. This language again appears in *Lewis v. ACB Business Services, Inc.*, 135 F.3d 389, 406 (6th Cir. 1998): “The [ECOA] was only intended to prohibit credit determinations based on ‘characteristics unrelated to creditworthiness,’” and in *Miller v. Countrywide Bank, NA*, 571 F.Supp.2d 251, 258 (D. Mass 2008): “[rejecting arguments] “based on subjective criteria beyond creditworthiness.”

These directives by the court matter because lenders might rationally price loans strategically to maximize profits, not just to profile individuals on credit risk. Using profits as the target may not

be legitimate under the law if lenders take higher profit margins above costs from a protected category group. We return to this point when talking about pricing of financial services.

In this U.S. setting that focuses on the business necessity motive for statistical discrimination, the determination by courts, regulators and compliance officers that certain practices are illegal is based on determining whether a specific input, such as a variable, dataset, or process, has led to a disparity in outcomes against a protected group. Recently, however, researchers have begun to advocate for approaches other than this traditional approach of simply analyzing an algorithm's decision process through consideration of its inputs. Pope and Sydnor (2011a) offer an approach that is based on constructing *corrections* to disparities created against protected groups by input variables, as opposed to actions to exclude such inputs.⁹ Cowgill and Tucker (2019) go further, arguing that “regulations focusing on *outcomes* exhibit more flexibility, fewer loopholes, greater efficiency and stronger incentives for innovation.” Yet Bartlett, Morse, Stanton, and Wallace (2020) contend that U.S. law demands an approach that first considers the justifiability of business necessity criteria in the input use, which means an output-assessment defense to a discrimination claim will not hold up in court. In a similar vein, Kleinberg, Ludwig, Mullainathan, and Sunstein (2018) argue that antidiscrimination prevention and enforcement should focus on the training procedure¹⁰ with an emphasis on the transparency¹⁰ of all code and datasets used to determine how inputs are used in algorithmic screening.¹¹

Regulatory Frontline #2: Algorithmic decision makers are increasingly advocating for the use of output-testing or correction approaches in compliance and court defenses, creating a fault line tension as to whether input-based compliance, which the law demands, can be challenged.

Frontiers of the Legality of Statistical Discrimination

In Europe, the landmark European General Data Protection Regulation (GDPR), which took effect in May 2018, creates challenges to the use of direct and indirect variables in profiling individuals for financial services. The GDPR applies to any company that processes the personal data of EU

⁹ See Altenburger and Ho (2018) for implementation frictions with this method.

¹⁰ A training procedure in credit screening uses a dataset of inputs, lending decisions, and loan success outcomes (repayment and profits) to determine the algorithm's parameters for screening new applicants.

¹¹ Further interesting reading on the application of statistical discrimination is in Bohren, Haggag, Imas, and Pope (2019), who provide a context for how behaviors relegate statistical discrimination practices into accurate and inaccurate sortings.

residents, regardless of where that processing takes place. The law forbids processing “special category” data (race, ethnicity, sexual orientation, religion, genetics and biometrics, etc.) unless the individual has given consent or such data collection is in the public interest. This provision makes direct use of these special-category variables illegal in all countries in the EU without consent. (An interesting question is whether the probability of consenting to such data collection itself varies systematically across groups.) The law also limits the ability of firms to use only an algorithm to make decisions that affect humans significantly, such as job offers or extensions of credit, without explicit consent. Firms must disclose to consumers the types of data and the logic used by the algorithm. Further, if a consumer is denied credit or other benefits on the basis of an algorithmic decision, firms must provide consumers the option to appeal the decision to a human decision-maker. It is yet to be seen whether individuals’ power over consent has leverage in preventing discrimination. Goodman and Flaxman (2019) discuss the frictions of the implementation of the GDPR for algorithmic decision-making and machine learning, and comment on how discrimination cannot be eliminated by excluding protected categories data because firms can still engage in statistical discrimination via variables that are correlated with protected classes.

Regulatory Frontline #3: Considerable uncertainty surrounds the eventual scope and implementation of the GDPR, especially as it is interpreted to encompass statistical discrimination on variables correlated with special-category data.

The second frontier challenge is being played out in the U.S., where technology has brought out new interest in the language about “unfair or deceptive acts or practices” regarding data from the Federal Trade Commission Act and the Dodd-Frank Wall Street Reform and Consumer Protection Act. A number of compliance risks associated with these laws in the context of data inputs and algorithmic decision-making are found in Evans (2017) and Federal Trade Commission (2016), including discussions of recent actions against companies for failing to disclose that borrowers’ credit limits could be reduced based on a behavioral scoring model, for misrepresenting how data collected on-line was used, and for selling data to customers that a company had reason to know would use the data for fraudulent purposes.

The third frontier challenge is playing out concerning what are called public accommodation laws (in the United States) or equality laws (in Europe, the United Kingdom, and Canada). These

statutes govern the equal rights of individuals to access business establishments and services. In the United States, these laws are primarily at the state level and vary in the types of individuals that are protected and in how comprehensive the laws are.¹² For example, the Unruh Act in California states that

All persons within the jurisdiction of this state are free and equal, and no matter what their sex, race, color, religion, ancestry, national origin, disability, medical condition, genetic information, marital status, sexual orientation, citizenship, primary language, or immigration status are entitled to the full and equal accommodations, advantages, facilities, privileges, or services in all business establishments of every kind whatsoever.

California Civil Code § 51(b).

The use of these state laws to police discrimination in online platforms is new and is rapidly evolving. Initially, legal actions against online service providers under the Unruh Act focused primarily on whether websites were accessible to people with disabilities.¹³ In 2019, the California Supreme Court established that the Unruh Act governs discrimination against people who intend to transact with websites, regardless of whether they actually do so.¹⁴ The most recent focus is discrimination in advertising. In July 2019, Governor Cuomo asked the New York Department of Financial Services to look into Facebook’s policy of allowing advertisers to target ads by protected classes, and in November 2019, a lawsuit filed in California alleged that this same Facebook policy violates the Unruh Act (*Opiotennione v. Facebook Inc.*, hereafter *Opiotennione*).¹⁵ These developments mark a possible expansion of the scope of public accommodation laws to discrimination by algorithms. It also marks an expansion of the venues available for redress in the U.S. for discrimination in financial services other than lending. As noted earlier, federal discrimination regulations are more comprehensive for lending than for other financial services.

Regulatory Frontline #4: It is yet to be seen as to whether jurisdictions will rule that public accommodation and equality laws apply to the policing of algorithmic information provision and access to mobile and online financial services.

¹² According to the National Conference of State Legislatures: “Five states... do not have a public accommodation law for nondisabled individuals. All states with a public accommodation law prohibit discrimination on the grounds of race, gender, ancestry and religion. In addition, 18 jurisdictions prohibit discrimination based on marital status, 25 prohibit discrimination based on sexual orientation and 21 prohibit discrimination based on gender identity. Nineteen jurisdictions also prohibit age-based discrimination in areas of public accommodation.” <https://www.ncsl.org/research/civil-and-criminal-justice/state-public-accommodation-laws.aspx>

¹³ *Thurston v. Midvale Corp.*, 39 Cal. App 5th 634 – Cal: Court of Appeal, 2nd Appellate Dist., 8th Div. 2019.

¹⁴ *White v. Square*, 7 Cal. 5th 1019, 1025 (2019).

¹⁵ Complaint, *Opiotennione v. Facebook Inc.*, No. 3:19-cv-07185 (N.D. Cal. filed Oct. 31, 2019).

IV.b. Credit Risk

Lenders may have a profit-maximizing motive to discriminate in loan underwriting inasmuch as membership in a protected class could be correlated with a credit risk factor that is difficult for the lender to observe, even after conditioning on observable measures of credit risk. Such unobserved risk factors are likely themselves inextricably tied with past or present discrimination. For example, because of long-standing discrimination in the labor market or in the education system, family resources (an often hidden fundamental in credit risk) might be lower for some groups than others. We emphasize that the possible existence of such a profit-maximizing motive does not mean that such profiling is legal.

The advent of technology has changed this dynamic in two ways. First, data advances have made it easier to measure underlying creditworthiness and thus reduced lenders' need to rely on flawed proxies. For example, lenders now extract information from transaction-account data to create measures of expected cash flow.¹⁶

Second, technology has amplified the practice of using correlated variables as proxies for missing fundamentals on credit risk, sometimes resulting in statistical discrimination against protected classes of households. What distinguishes credit risk assessment today from a long history of lenders using proxies is the sheer number of variables involved and the complexity and opacity of the algorithms. The use of new variables is not necessarily problematic. Upstart Network, for example, asserted in a submission with the Consumer Financial Protection Bureau that its use of education and employment history as underwriting variables expanded access to credit without displaying any disparities that required further fair lending analysis (Ficklin and Watkins, 2019).

This is not always the case, but the evidence is yet thin. What we do know comes from implementations in countries with fewer restrictions on the use of proxy variables as it relates to discrimination. For example, Berg, Burg, Gombovic, and Puri (2019) showed that “digital footprints” such as the type of device (tablet, computer, phone), operating system (Windows, iOS,

¹⁶ One evaluation of these cash-flow measures found that their use in underwriting appears to expand access to credit without creating a disparate impact on protected groups (FinRegLab, 2019).

Android), and email provider predicted default rates among the customers of a German lender. It is possible that credit risk assessment based on digital footprints would fail discrimination suits, as the case could be made that such variables disproportionately affect a protected group beyond any effect that operates through creditworthiness. Likewise, BJORKEGREN and GRISSIN (2019) generated measures of creditworthiness from mobile phone usage data in a South American country and showed that these measures predicted default. Again, mobile phone usage might be additionally correlated with a protected category, even beyond its correlation with credit risk fundamentals.

Discrimination Gateway #3: Practices of scoring customers for creditworthiness based on variables (including the now-pervasive digital footprint and mobile data) that correlate with membership in a protected class are at risk to lead to discrimination.

IV.c. Pricing

Pricing is fertile ground for discrimination. There is wide variation in the amount that individuals pay for financial services, even for the same loan originated in the same market on the same day (BHUTTA, FUSTER, and HIZMO, 2019; ALEXANDROV and KOULAYEV, 2018). Within that variation, women and Black/Latinx men appear to pay higher rates on loans than white men (WOODWARD and HALL, 2012 (and references therein in online appendix A), BARTLETT et al., 2019). Class-action suits and enforcement actions brought against mortgage lenders in the aftermath of the financial crisis provided considerable evidence of discretionary pricing policies that discriminated against Black and Latinx borrowers (AYRES, KLEIN, and WEST, 2017).¹⁷

This pricing dispersion persists, in part, because many individuals do little comparison-shopping for financial services products and because some markets are more competitive than others. ALEXANDROV and KOULAYEV (2018) noted that about half of mortgage borrowers do not shop before taking out a mortgage. BHUTTA, FUSTER, and HIZMO (2019) find that paying a lower mortgage rate is associated with whether the borrower considered multiple lenders or consulted mortgage websites. In their paper, this shopping behavior partly explained why borrowers with lower

¹⁷ Bhutta and Hizmo (2020) show that Black and Latinx borrowers who obtained FHA mortgages in 2014 and 2015 paid higher mortgage rates than white and Asian borrowers, but that this interest-rate differential is offset by Black and Latinx borrowers paying fewer points. However, although the choice set of (interest rate, point) contracts does not appear to vary by race, Black and Latinx borrowers appear to end up in contracts that are more profitable for the lender, which could be evidence of discrimination if it results from a specific lender practice.

incomes, wealth, and credit scores paid higher mortgage rates. Although these particular characteristics are not protected classes, Woodward and Hall (2012) show that limited shopping is one of the reasons why Black and Latinx borrowers pay higher mortgage broker charges.¹⁸

Bartlett et al. (2019) provide supporting evidence that even in a setting in which lenders bear no credit risk, Black and Latinx borrowers pay higher prices for mortgages issued by both traditional and FinTech lenders. Their contribution does not provide direct evidence tying rates to shopping behavior as in Woodward and Hall (2012). However, because Bartlett et al. can rule out omitted credit risk variables that could confound the interpretation of other studies, the disparate pricing result must come from either profiling for shopping or the competition environment, and thus is a complement to Woodward and Hall (2012).

Indeed, algorithms can infer the propensity of an individual to shop around, as well as other individual-specific factors that affect a consumer's willingness to pay for a product. Thus technology allows firms to make these inferences with far greater precision than was available previously and to use that information in a pricing strategy. Donnelly, Ruiz, Blei, and Athey (2019) (also see Fuster, Goldsmith-Pinkham, Ramadorai, and Walther, 2018) present a machine-learning model, for example, that identifies which consumers are most price sensitive in their demand for a given product and allows for personally targeted price discounts. While this behavior is profit-maximizing for the firm, U.S. courts have ruled that profit motives beyond credit risk are not legal reasons for statistical discrimination in pricing, as discussed in section *IV.a.* and in Bartlett et al. (2020).

We also note that the economist concept of “profit-maximizing for the firm” can be observationally equivalent to the consumer advocate’s concept of “profiteering from vulnerable groups.” The connection between pricing and shopping behavior is complicated, for example, because shopping behavior also reflects differences in search costs across groups. Ayers (1991) describes some reasons why women and nonwhite men might face higher search costs in the auto retail market.

¹⁸ Tabulations from the 2016 National Survey of Mortgage Originations, however, suggest that Hispanics and nonwhites are a bit more likely than non-Hispanic whites to consider more than one lender (Avery, Bilinski, Clement, Critchfield, Frumkin, Keith, Mohamed, Pafenberg, Patrabansh, and Schultz, 2018, Table 10).

Shopping is also linked to advertising, and as we discuss in the next section, there is some evidence that Black and Latinx individuals are less likely to be shown ads that provide them with useful information. While firms may not have much control over shopping behavior per se, they do have control over the information that consumers receive and the way choices are framed. We turn next to a discussion of how technology has changed that dynamic.

Discrimination Gateway #4: Practices of statistical discrimination that profile shopping behavior can lead to discrimination, even if the firm's statistical discrimination is motivated purely by the economics of profit maximization.

IV.d. Advertising and Information

Companies always have had an incentive to tailor their advertising content, targeting, and delivery to different markets. In this section we discuss how technology is changing this playing field and how these changes interact with the existing laws on discrimination. Evans and Miller (2019) provide a fuller treatment of the fair lending implications of these issues.

Evidence from mortgage advertising underscores the existing differences in information delivered to target audiences, even before technology. Perry and Motley (2009) show that prime borrowers were more likely to be shown advertisements with detailed information that helped them make better financial decisions, whereas subprime borrowers were more likely to be shown information that played on their fears. Some of the ads in Perry and Motley (2009) were explicitly aimed at minority borrowers.¹⁹ This Perry and Motley finding is consistent with a lawsuit filed by the Attorney General of Illinois that alleged that Countrywide Financial Corporation aggressively marketed subprime mortgages to Black and Latinx borrowers.²⁰

¹⁹ Perry and Motley (2009) documented the differences between prime and subprime mortgage advertisements in a sample of ads placed on television, radio, and print from 2005 to 2007. They found that 50 percent of ads targeted to prime borrowers contained specific loan details, compared with 21 percent of ads targeted to subprime borrowers. Mortgage ads that positioned the lender as a solution for borrowers “drowning in debt” were placed almost exclusively in publications targeting Black and Latinx households. These ads contained very little information on the actual loan terms and thereby increased the borrower’s dependence on the mortgage lender.

²⁰ Complaint, People of the State of Ill. v. Countrywide Fin. Corp., No. 10-ch-27929 (Cir. Ct. of Cook Cty., Ill. filed June 29, 2010). Another example of the discrepancy between prime and subprime advertising campaigns can be seen by contrasting the findings of Grundl and Kim (2018) and Gurun, Matvos, and Seru (2016). The first paper showed that the direct-mail mortgage refinance advertisements to prime borrowers appeared to improve their welfare by prodding them to refinance when it was in their interest. The second paper looked at subprime mortgage lenders and found that the lenders that advertised more heavily in print media and direct mail also tended to provide less favorable

Although it is unclear whether technology affects firms' incentives to provide helpful or deceptive information, particularly in ways that differ by protected class, technology clearly amplifies the ability of firms to target their advertising. As Athey and Luca (2019) note, "Older media, such as print and television, do not allow for showing different advertisements or tracking behavior at the individual consumer level" (p. 219). In contrast, in Internet and mobile advertising, advertisers create an advertisement and specify their target audience. Often multiple advertisers are looking for the same audiences, and an algorithmic auction ensues based on the relevance of the ad and the advertisers' relative willingness-to-pay.²¹ All of these mechanisms result in systematic differences in the ads that different demographic groups see.

In online marketing, "ad targeting" is the term that applies to when advertisers choose to have their ads shown only to certain groups. For ads that concern housing, employment, or credit, ad targeting to a protected class is illegal under the same federal laws that govern discrimination in other aspects of these markets.²² Tech companies are increasingly being held accountable for facilitating such targeting.²³ Open questions are threefold. First, is this targeting legal for advertisements for financial services other than credit? The current test case, mentioned in section *IV.a.*, concerns whether the court will decide that public accommodations law prohibit discrimination in advertising for financial services other than credit and has been brought under California's Unruh Act. *Opiotennione* alleges that Facebook denied older and female users "the opportunity to learn about and obtain financial services...over the past three years (or longer) due to Facebook's discriminatory advertising and business practices and its aiding and abetting of financial services companies' discriminatory advertising and business practices." The second open question concerns how ad targeting will interact with the European GDPR rules on consent; the third

terms to borrowers. This pattern was concentrated in zip codes with more minority, low-educated, or low-income residents.

²¹ See Ali et al. (2019) for a detailed description of Facebook's advertising platform.

²² The relevant laws are the Fair Housing Act, the Equal Credit Opportunity Act (ECOA), Title VII of the Civil Rights Act, and the Age Discrimination in Employment Act. A creditor is allowed to engage in affirmative advertising to solicit or encourage traditionally disadvantaged groups to apply for credit (staff commentary to Regulation B, which implements ECOA).

²³ In March 2019, Facebook settled five lawsuits that alleged that its advertising platform allowed companies to illegally target ads (Murphy, 2019), and in July 2019 the Equal Employment Opportunity Commission found that several employers violated federal law when they targeted their job ads on Facebook only to younger or male individuals.

concerns how fair lending laws will consider ad targeting on variables that are highly correlated with protected classes but not the protected classes themselves.

“Ad delivery” is the term that explains the mechanical auction process by which algorithms allocate advertising space. The fact that “female eyeballs are more expensive” (women both control a greater share of household expenditures and have higher viewing-to-purchase rates (Cowgill and Tucker, 2019; Lambrecht and Tucker, 2019)) implies that ads from other retailers, such as employment opportunities in STEM fields in the Lambrecht and Tucker paper and presumably financial services as well, will be crowded out (Lambrecht and Tucker, 2019).²⁴ Ali, Sapiezynski, Bogen, Korolova, Mislove, and Rieke (2019) conducted experiments that deocumented that Facebook was more likely to direct ads that included stereotypical male and female pictures to male and female audiences, respectively, even though the authors did not target their ads based on gender.²⁵ As Lambrecht and Tucker note, one way to counteract these distortionary effects of ad delivery is to target ads deliberately to different groups. This remedy, however, has the potential itself to be judged illegal discrimination for certain types of ads.

Discrimination Gateway #5: Practices of technology-facilitated advertising, including ad targeting and ad delivery, can lead to discrimination.

V. Innovation and Inclusion

The prior sections considered the role of technology in mitigating or facilitating discrimination resulting from human and algorithmic decisions. In this final, brief section, we step back from decisions and consider a final point that innovation itself is not immune from discrimination frontiers.

Worldwide, disparities exist in access to and inclusiveness of financial services. For example, in the U.S., younger, Black, and Latinx individuals are less likely to have a banking relationship (FDIC, 2017), and more likely to have a limited or insufficient set of credit history information such that access to credit is inhibited (Brevoort, Grimm, and Kambara, 2015); Black and Latinx

²⁴Ali et al. (2019) subsequently verified the Lambrecht and Tucker finding by running a series of advertisements on Facebook and varying the maximum bid amount. The share of women shown the ad rose with the daily ad budget.

²⁵ Employment ads for lumberjacks were delivered primarily to white men, whereas black women were more likely to see employment ads for janitor positions.

individuals are more likely to have their mortgage applications rejected (Dietrich, Liu, Skhirtladze, Davies, Jo, and Candilis, 2019). Technology can reduce disparities in financial access and inclusion, such as through new modes of financial service delivery.²⁶ It can also increase these disparities, such as when more privileged households are more likely to benefit from new technology.²⁷ The common perception (although untested and unclear) is that on net, technology reduces disparities by expanding access.

Yet the impact on disparities is not equivalent to the impact on discrimination. Instead, each new innovation hits the market with the possibility that *even if the innovation reduces disparities*, it must be provided without any of the forms of discrimination discussed in this chapter. This can be a tall order, given the complex issues discussed in this chapter, and reaching those previously excluded from markets can be costly.

This creates a bit of an ironic outcome. Ex ante, discrimination may have caused disparities in access to financial services for some groups. However, innovators cannot simply use technology tools that remedy disparities without considering whether those tools discriminate against some groups. Regulators are keenly aware of this issue, and are increasingly providing support to help firms understand the fair-lending implications of their innovations. For example, the Financial Conduct Authority (FCA) in the United Kingdom and the Consumer Financial Protection Bureau (CFPB) and the Office of the Comptroller of the Currency in the United States have introduced offices of innovation. The FCA and the CFPB have created “sandboxes” that allow companies to obtain regulatory relief for a limited period of time while the companies test financial innovations.

²⁶ The widespread adoption of smartphones has increased access to financial services for individuals who live in communities underserved by traditional financial institutions. In the U.S., where 80 percent of adults own a mobile phone, the share of individuals owning a mobile phone does not vary by race (Pew Research Center, 2019), implying that innovations in mobile financial services have the potential to reach all racial groups equally. In 2016, U.S. regulators acknowledged the importance of access to technology for financial inclusion when they finalized rules that included “improving broadband access” as among the activities that might count toward a bank’s assessment under the Community Reinvestment Act. (See https://www.ffiec.gov/cra/pdf/2016_QA_Federal_Register_Notice.pdf.)

²⁷ In the U.S., borrowers from peer-to-peer lenders appear to be more affluent than borrowers overall (Morse, 2015; DiMaggio and Yao, 2019). Likewise, mortgage borrowers who obtain their loans from FinTech firms tend to be more educated than mortgage borrowers overall (Fuster, Plosser, Schnabl, and Vickery, 2019). Dettling, Goodman, and Smith (2018) find that the spread of broadband technology led to students performing better on the SAT and applying to more colleges, presumably because broadband reduced the effort and informational costs of college applications. However, the effects were strongest for students from families of higher socioeconomic status, “suggesting that the technology may have increased pre-existing inequalities.”

The companies typically supply data in exchange and are monitored closely during the test period. In the United Kingdom, initial data suggests that the FCA's Innovate program has reduced the time it takes for innovations to reach the market (Financial Conduct Authority, 2019).

VI. Conclusion

The punchline of our analysis is that technological innovation has changed the way that discrimination manifests itself in financial services markets, with an ambiguous effect on the overall level of discrimination going forward.

The fact that technology has changed the manifestation of discrimination may not be a surprise. Yet once one starts to amalgamate all the channels of removing discretionary discrimination and increasing access to financial services while potentially introducing new discrimination through algorithmic process and human involvement in data and coding, the landscape shift becomes profound. Taste-based discrimination surely is less of a factor when decision-making in financial services markets is guided by algorithms rather than humans. However, opaque and complex algorithms drawing upon thousands of variables are increasingly governing how the creditworthiness of consumers is assessed, which advertisements consumers see, and which prices they are quoted. These processes have the potential to engender significant illegal statistical discrimination.

Regulation is evolving with technology, leading to the regulatory frontlines that we draw out in this chapter. Policymakers, regulators, and the courts must figure out how to modify the existing fair-lending (and other equality law) infrastructure to a world where bias is embedded in data sets and the inner workings of algorithms. They must consider whether the traditional input-based framework for detecting and policing discrimination can be adapted to a world where the inputs have become so complex. They must judge when a variable is correlated enough with a protected class to be discriminatory itself, a task made more complicated by the sheer number of variables under consideration in algorithmic decision-making. And these decision-makers must also discern how the existing laws interact with new laws, such as the GDPR and other data privacy laws, as well as with public accommodation laws that have not historically been used to police

discrimination in online spaces. The implications of these laws for technological innovation by financial services firms is undetermined as of the writing of this chapter.

Ideally, these momentous decisions would be aided and influenced by economic and policy research. We are concerned, though, that an inadequate amount of research will be produced that focuses on the welfare perspective of consumers of financial services. In our preparation of this manuscript, the imbalance away from such studies was evident. Because the algorithms are complex and the datasets are proprietary, much research, by necessity, is conducted in partnership with technology companies. This research is valuable, but the imbedded incentive structure will likely result in research that focuses more on the questions of interest to companies than to consumers.

We end finally on a positive note. The technological transformation of financial services is making enormous inroads into improving choices, competition, and access for millions of people. And although we indeed conclude with the surprising thought that this technological transformation may do the opposite of leveling the playing field, this outcome is not set in stone. The choices made in the next few years will influence whether discrimination is a pervasive feature in the markets for financial services or not, and whether technological innovation results in these markets becoming more fragmented or more inclusive. The goal of our chapter is to bring these fault lines forward to improve the chances of reaching the inclusive outcome.

References

- Alesina, Alberto F., Francesca Lotti, and Paolo Emilio Mistrulli, 2013. "Do Women Pay More for Credit? Evidence from Italy." *Journal of the European Economic Association*, 11(1):45-66.
- Alexandrov, Alexei and Sergei Koulayev, 2018. "No Shopping in the U.S. Mortgage Market: Direct and Strategic Effects of Providing Information." Consumer Financial Protection Bureau Office of Research Working Paper Series, 2017-01.
- Ali, Muhammad, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke, 2019. "Discrimination through Optimization: How Facebook's Ad Delivery Can Lead to Skewed Outcomes." arXiv:1904.02095 [cs.CY].
- Altenburger, Kristen M. and Daniel E. Ho, 2018. "When Algorithms Import Private Bias into Public Enforcement: The Promise and Limitations of Statistical Debiasing Solutions." *Journal of Institutional and Theoretical Economics*, Vol. 174(1).
- Arrow, Kenneth J., 1973. "The Theory of Discrimination." In Orley Ashenfelter and Albert Rees (eds.), *Discrimination in Labor Markets*, Princeton, NJ: Princeton University Press.
- Athey, Susan and Michael Luca, 2019. "Economists (and Economics) in Tech Companies." *Journal of Economic Perspectives*, 33(1):209-30.
- Avery, Robert B., Mary F. Bilinski, Audrey Clement, Tim Critchfield, Samuel Frumkin, Ian H. Keith, Ismail E. Mohamed, Forrest W. Pafenberg, Saty Patrabansh, and Jay D. Schultz, 2018. "A Profile of 2016 Mortgage Borrowers: Statistics from the National Survey of Mortgage Originations." National Mortgage Database Technical Report 6.0, April 18.
- Ayres, Ian, 1991. "Fair Driving: Gender and Race Discrimination in Retail Car Negotiations." *Harvard Law Review*, 104(4):817-872.
- Ayres, Ian, Gary Klein, and Jeffrey West, 2017. "The Rise and (Potential) Fall of Disparate Impact Lending Litigation." In Lee Anne Fennell and Benjamin J. Keys (eds.), *Evidence and Innovation in Housing Law and Policy*, Cambridge: Cambridge University Press.
- Barocas, Solon, and Andrew D. Selbst, 2016. "Big Data's Disparate Impact." 104 *California Law Review*, 104:671-732.
- Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace, 2019. "Consumer Lending Discrimination in the FinTech Era." Working paper.
- Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace, 2020. "Algorithmic Accountability: A Legal and Economic Framework." Working paper.
- Bayer, Patrick, Fernando Ferreira, and Stephen L. Ross, 2017. "What Drives Racial and Ethnic Differences in High-Cost Mortgages? The Role of High-Risk Lenders." *Review of Financial Studies*, 31(1):175-205.
- Becker, Gary, 1957. *The Economics of Discrimination*. Chicago: University of Chicago Press.

- Berg, Tobias, Valentin Burg, Ana Gombovic, and Manju Puri, forthcoming. "On the Rise of FinTechs - Credit Scoring using Digital Footprints." *Review of Financial Studies*.
- Bhutta, Neil, Andreas Fuster, and Aurel Hizmo, 2019. "Paying Too Much? Price Dispersion in the US Mortgage Market." Working paper.
- Bhutta, Neil, and Aurel Hizmo, 2020. "Do Minorities Pay More for Mortgages?" Finance and Economics Discussion Series 2020-007. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2020.007>.
- Björkegren, Daniel, and Darrell Grissen, 2019. "Behavior Revealed in Mobile Phone Usage Predicts Credit Repayment." *World Bank Economic Review*.
<https://doi.org/10.1093/wber/lhz006>
- Blodgett, Su Lin, and Brendan O'Connor, 2017. "Racial Disparity in Natural Language Processing: A Case Study of Social Media African-American English." arXiv:1707.00061v1 [cs.CY] 30 Jun 2017.
- Bohren, J. Aislinn, Kareem Haggag, Alex Imas, and Devin G. Pope, 2019. "Inaccurate Statistical Discrimination," Working paper.
- Buolomwini, Joy, and Timnit Gebru, 2018. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." *Proceedings of Machine Learning Research*, 81:1-15.
- Brevoort, Kenneth, Philipp Grimm, and Michelle Kambara, 2015. "Data Point: Credit Invisibles." Consumer Financial Protection Bureau Office of Research, May.
- Caliskan, Aylin, Joanna Bryson, and Arvind Narayanan, 2017. "Semantics Derived Automatically from Language Corpora Contain Human-like Biases." *Science*, 356(6334):183-186.
- Charles, Kerwin Kofi, and Erik Hurst, 2002. "The Transition to Home Ownership and the Black-White Wealth Gap." *Review of Economics and Statistics*, 84(2):281-297.
- Cowgill, Bo, and Catherine Tucker, forthcoming. "Economics, Fairness, and Algorithmic Bias." *Journal of Economic Perspectives*.
- Dastin, Jeffrey, 2018. "Amazon Scraps Secret AI Recruiting Tool that Showed Bias against Women." Reuters, October 9.
- Deku, Solomon Y., Alper Kara, and Philip Molyneux, 2016. "Access to Consumer Credit in the UK." *European Journal of Finance*, 22(10):941-964.
- Detting, Lisa J., Sarena Goodman, and Jonathan Smith (2018). "Every Little Bit Counts: The Impact of High-Speed Internet on the Transition to College." *The Review of Economics and Statistics*, 100(2): 260-273.
- Dietrich, Jason, Feng Liu, Akaki Skhirtladze, Misha Davies, Young Jo, and Corinne Candilis, 2019. "Data Point: 2018 Mortgage Market Activity and Trends," Consumer Financial Protection Bureau, October.

- DiMaggio, Marco and Vincent Yao, 2019. “Fintech Borrowers: Lax-Screening or Cream-Skimming?” Working paper.
- Dobbie, Will, Andres Liberman, Daniel Paravisini, and Vikram Pathania, 2019. “Measuring Bias in Consumer Lending.” Working paper.
- Doleac, Jennifer L., and Luke C.D. Stein, 2013. “The Visible Hand: Race and Online Market Outcomes.” *Economic Journal*, 123(572):F469-F492.
- Donnelly, Robert, Francisco R. Ruiz, David Blei, and Susan Athey, 2019. “Counterfactual Inference for Consumer Choice Across Many Product Categories,” arXiv:1906.02635 [cs.LG].
- Edelman, Benjamin, and Michael Luca, 2014. “Digital Discrimination: The Case of Airbnb.com.” Harvard Business School Working Paper 14-054.
- Edelman, Benjamin, Michael Luca, and Dan Svirsky, 2017. “Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment.” *American Economic Journal: Applied Economics*, 9(2):1-22.
- Edelman, Benjamin, and Abby Stemler, 2018. “From the Digital to the Physical: Federal Limitations on Regulating Online Marketplaces.” *Harvard Journal on Legislation*, 56(1):141-198.
- Evans, Carol, 2017. “Keeping Fintech Fair: Thinking about Fair Lending and UDAP Risks.” *Consumer Compliance Outlook*, 2:1-9.
- Evans, Carol and Westra Miller, 2019. “From Catalogs to Clicks: The Fair Lending Implications of Targeted, Internet Marketing.” *Consumer Compliance Outlook*, 3:1-9.
- Federal Trade Commission, 2016. “Big Data: A Tool for Inclusion or Exclusion? Understanding the Issues.” Federal Trade Commission Report.
- Ficklin, Patrice and Paul Watkins, 2019. “An Update on Credit Access and the Bureau’s First No-Action Letter.” Consumer Finance Protection Bureau Blog, August 6.
- Financial Conduct Authority, 2019. “The Impact and Effectiveness of Innovate.” April.
- Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther, 2018. “Predictably Unequal? The Effects of Machine Learning on Credit Markets.” Working paper.
- Fuster, Andreas, Matthew Plosser, Phillipp Schnabl, and James Vickery, 2019. “The Role of Technology in Mortgage Lending.” *Review of Financial Studies*, 32(5):1854-1899.
- Goodman, Bryce, and Seth Flaxman, 2017. “European Union Regulations on Algorithmic Decision-making and a `Right to Explanation,” *AI Magazine*, 38(3):50-57. arXiv:1606.08813 [stat.ML].
- Grundl, Serafin, and You Suk Kim, 2019. “Consumer Mistakes and Advertising: The Case of Mortgage Refinancing.” *Quantitative Marketing and Economics*, 17: 161-213.

- Gurun, Umit, Gregor Matvos, and Amit Seru, 2016. “Advertising Expensive Mortgages.” *Journal of Finance*, 71(5): 2371-2416.
- Hale, Brenda, 2018. “Equality and Human Rights.” Oxford Equality Lecture 2018. October 29.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan, 2018. “Human Decisions and Machine Predictions,” *Quarterly Journal of Economics*, 133(1):237-93.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Cass R. Sunstein, 2018. “Discrimination in the Age of Algorithms.” *Journal of Legal Analysis*, 10:113-174.
- Lambrecht, Anja, and Catherine Tucker, 2019. “Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads.” *Management Science*, 65(7): 2966-81.
- Morse, Adair, 2015. “Peer-to-Peer Crowdfunding: Information and the Potential for Disruption in Consumer Lending?” *Annual Review of Financial Economics*, 7:463-482.
- New York Department of Financial Services, 2019. “Governor Cuomo Calls on DFS to Investigate Claims that Advertisers Use Facebook Platform to Engage in Discrimination.” Press release, July 1.
- Perry, Vanessa G., and Carol M. Motley, 2009. “Reading the Fine Print: Advertising and the Subprime Mortgage Crisis.” *California Management Review*, 52(1):1-16.
- Pew Research Center, 2019. “Mobile Technology and Home Broadband 2019.” June 13.
- Phelps, Edmund, 1972. "The Statistical Theory of Racism and Sexism." *The American Economic Review*, 62(4): 659–661.
- Pop, Valentina, 2013. “German Data Chief Attacks Credit-Profile Firms.” EU Observer, April 25. <https://euobserver.com/economic/119930>.
- Pope, Devin G., and Justin R. Sydnor, 2011a. “Implementing Anti-Discrimination Policies in Statistical Profiling Models.” *American Economic Journal: Economic Policy* 3: 206–231.
- Pope, Devin G., and Justin R. Sydnor, 2011b. “What's in a Picture? Evidence of Discrimination from Prosper.com.” *Journal of Human Resources*, 46(1):53-92.
- Ravina, Enrichetta, 2019. “Love and Loans: The Effect of Beauty and Personal Characteristics in Credit Markets.” Working paper.
- Rothstein, Richard, 2017. *The Color of Law*. New York: Liveright Publishing Corporation.
- Safronova, Valeriya, 2019. “Gender Stereotypes Banned in British Advertising.” *New York Times*, June 14. <https://www.nytimes.com/2019/06/14/style/uk-gender-stereotype-ads-ban.html>

Scott Morton, Fiona, Florian Zettelmeyer, and Jorge Silva-Russo, 2003. "Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?" *Quantitative Marketing and Economics*, 1(1):65-92.

Stowe, Stacey, 2019. "New York City to Ban Discrimination Based on Hair." *New York Times*, February 18. <https://www.nytimes.com/2019/02/18/style/hair-discrimination-new-york-city.html>

Woodward, Susan E. and Robert E. Hall, 2012. "Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence." *American Economic Review*, 102(7): 3249-76.