

**Invisible Allies: Algorithmic Consumer Profiling and the Rise
of New Group Harms**

Andrew Miller*

An algorithm, not a person, can now determine whether we get bumped from an oversold plane, wait two hours on hold, pay a surcharge that others don't, or suffer a wide array of other consumer harms large and small. Businesses make these decisions based on an endless stream of data points about each one of us, using our behavioral quirks and vulnerabilities to make increasingly accurate predictions about our future behavior. Scholars in law and other fields have shown persuasively how such tools can reproduce and even worsen existing patterns of discrimination. They are correct to sound the alarm.

But the existing literature devotes much less attention to the worrying prospect that our list of groups in need of protection—those persistently on the losing end of marketplace decisions—may now be woefully incomplete. This Article seeks to broaden the scope of our concern, explaining how these same algorithmic tools can create new, and in some ways more invisible, patterns of discrimination and disadvantage. Even more worrying, members of the new harmed groups may not realize that such groups exist as “groups” at all, or that they are members, hobbling their ability to join forces and fight back. Yet from the business point of view, these groups are very real, be they people with low self-esteem, introverts, or late risers, among innumerable other categories known and unknown.

Drawing from scholarship on privacy law, consumer protection law, civil procedure, and social theory, this Article

*Assistant Professor of Law, Brooklyn Law School. Suggestions, questions, comments, and advice from the following people were an enormous help throughout my time working on this project: Bill Araiza, Miriam Baer, Jack Balkin, Ray Brescia, Sam Daly, Robin Effron, Jon Endean, Arianna Fabri, Lally Gartel, Nik Guggenberger, Max Helveston, Joel Hochman, Amy Kapczynski, Josh Macey, Daniel Markovits, Sari Mazzurco, Christina Mulligan, Kali Murray, Robert Post, Vijay Raghavan, Kate Redburn, Noah Rosenblum, Geoff Shaw, Jocelyn Simonson, Alicia Solow-Niederman, Adam Steinman, Olivier Sylvain, Rory Van Loo, and Rebecca Wexler. I also benefited from the opportunity to receive feedback from members of the LPE-NYC Junior Scholars Workshop and the Information Society Project at Yale Law School. I owe all of them my thanks.

argues that the invisibility of these new harmed groups has troubling implications for distributive justice and democratic deliberation. The argument proceeds in two parts. First, I take the position that as long as such algorithms function without any meaningful chance of accountability, they impose a democratic harm on the entire polity. This is because the norms we use to assess the fairness of a given distributive mechanism aren't independent of empirical reality. Rather, we update and apply our norms in response to our shared understanding of what's happening on the ground—or at least, as I argue, that is how the deliberative process ought to work. The invisibility of these algorithmic tools short-circuits the mechanisms we would otherwise use to ensure oversight of emerging technologies and their uses in the marketplace. I further explain why giving members of the new harmed groups the opportunity to communicate among themselves would be essential for any such deliberative process.

Second, I argue (perhaps counterintuitively) that litigation could play a crucial role in kickstarting this sort of democratic deliberation. Before more “classically” democratic modes of engagement can fulfill their potential, we may need the coercive power of the courts to help uncover what precisely is happening. Our legal system provides us with helpful procedural tools to shed light on these collective harms, but only if we use and refine them strategically. In particular, we must ensure our system of aggregate litigation remains able to surface and redress new patterns of harm where they occur. Only by giving the new harmed groups the chance to mobilize against the market logics that created them, empowering them to pursue collective action and articulate both moral and legal claims, can we stop market actors from dividing us all into winners and losers on whatever basis they choose.

Article Contents

Introduction	526
I. A New Generation of Consumer Profiling:	
Everything Everywhere All at Once	529
A. What Businesses Are Trying to Learn or Predict About Their Customers.....	530
B. What Sorts of Data Do Businesses Use?	532
C. Consumer Profiling in Context	533
1. “Matt” and His “Own” Data.....	533
2. Matt’s Data Combined with Data from Other Consumers	539
D. What Sorts of Individualized Decisions Result from Consumer Targeting?	543
II. Individualized Treatment, Invisible Patterns of Harm	550
A. The New Harmed Groups	551
B. Who Wins, Who Loses, and Why: Distributive Concerns	554
1. Context-Specific Norms and Harms.....	556
2. The Credit Market: Personalization is Expected, But What Kind?	559
3. The Retail Market: Different Distributive Logics and Norms	564
C. Finding Invisible Allies	573
1. Collective Problem-Solving and the Role of Groups	574
2. Top-Down and Bottom-Up Stimuli for Solidarity.....	578
3. Who’s in This With Me?	582
4. How We Conceive of the New Harmed Groups	588
III. What Litigation (and Procedure) Can Do.....	591
A. Aggregate Litigation as Part of the Democratic Ecosystem	595
B. Procedural Shortcomings.....	602
1. Class Actions	603
2. Multidistrict Litigation	606
3. Article III Standing	608
4. Additional Procedural Challenges	610
C. Green Shoots and Paths Forward.....	611

1.	Class Actions.....	612
2.	Multidistrict Litigation.....	614
3.	Article III Standing	615
Conclusion.....		616

Introduction

Algorithmic decision-making has begun to radically reshape interactions between buyers and sellers in the consumer marketplace. Attributes of consumer goods and services that were once uniform no matter who the buyer was—sticker prices, warranties, and other terms and conditions—can now be customized depending on a wide range of consumer data, from purchasing history and zip code to the identities of friends and family members. Even transactions that have long allowed for some adjustment according to the identity of the buyer (e.g., the decision to extend credit) can now take into account a much richer set of data points and predictive models than ever before. Many consumers remain unaware that these calibrations are even being made, let alone on what basis.

There may be much to like in this trend: sellers can now better match products to the needs of individual customers, make more accurate ability-to-pay determinations, and avoid falling back on old biases as they gain access to more granular data about human behavior. But I will argue that the growing use of predictive analytics in the consumer marketplace raises serious questions of distributive justice that ought to be the subject of collective, democratic decision-making. And as I will discuss, the way in which these marketplace decisions are made—opaquely and on a hyper-individualized basis—undermines our ability to act collectively to assess and combat the resulting distributive harms. Although aggregate litigation ought to be a key catalyst for precisely this sort of collective action, a number of relevant procedural tools may make for an increasingly poor fit with the algorithmic harms chronicled in this Article. After identifying these procedural shortcomings, I propose some paths forward that might nevertheless exist under current law.

This Article contributes to the emerging literature on algorithmic decision-making in a number of ways. First, there has been widespread writing about the need to monitor decision-making algorithms whose outcomes end up masking

or reinforcing pre-existing social biases.¹ But much of this work takes for granted that we know in advance who needs protecting,² or which sorts of harms we want to address. This Article highlights the inherently social and collective process of determining which sorts of groups need protection, including newly disfavored groups that emerge over time. And yet, as this Article argues, the very ways in which these algorithmic tools operate prevent that collective conversation from taking place. Similarly, while commentators have noted the importance of making algorithms more transparent and accountable,³ this Article adds to their analysis by zeroing in on a particular goal that (a) would justify the move toward greater transparency and accountability, and (b) could guide how and for whom we make patterns of harm more visible. Namely, it is the need to solve a specific sort of political economy problem—the invisibility of harms and the resulting difficulty of counter-mobilization—that militates in favor of solutions focused on revealing patterns of shared disadvantage. The potential upsides and downsides of “dynamic pricing” have long been a topic of debate in the literature,⁴ and several commentators

¹ See, e.g., Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1 (2014); Anupam Chander, *The Racist Algorithm*, 115 MICH. L. REV. 1023 (2017); CATHY O’NEILL, *WEAPONS OF MATH DESTRUCTION: HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY* (2d ed. 2017).

² An important exception is Sandra Wachter’s piece, *The Theory of Artificial Immutability: Protecting Algorithmic Groups Under Anti-Discrimination Law*, 97 TUL. L. REV. 149 (2023), which grapples explicitly with the phenomenon of “new” groups that are created and disadvantaged by algorithmic tools. As discussed *infra* Part II.C.4, however, my approach differs substantially from Wachter’s in that her normative framework emphasizes the resulting danger to individual autonomy. In contrast, I am more concerned with the aggregate-level democratic harms that flow from such invisible social sorting (and the mismatch with the procedural tools we have designed to surface collective harms). In this way, my approach has more in common with, and draws from, the broader privacy literature that focuses on relational and democratic harms. See, e.g., Salomé Viljoen, *A Relational Theory of Data Governance*, 131 YALE L.J. 573, 586 (2021); Alicia Solow-Niederman, *Information Privacy and the Information Economy*, 117 NW. U. L. REV. 357, 362 (2022).

³ See, e.g., FRANK PASQUALE, *THE BLACK BOX SOCIETY* (2015); Joshua A. Kroll et al., *Accountable Algorithms*, 165 U. PA. L. REV. 633 (2017).

⁴ See, e.g., Mark Klock, *Unconscionability and Price Discrimination*, 69 TENN. L. REV. 317 (2002) (arguing that price discrimination should be treated as a presumptively unconscionable practice); Matthew A. Edwards,

have recognized its potential (regardless of its effect on aggregate welfare outcomes) to raise serious distributive justice concerns.⁵ But this Article adds to this literature by exploring in greater detail how distributive harms might differ across specific sorts of consumer markets in terms of both material outcomes and threats to existing norms.

Finally, this Article seeks to put the literature on algorithmic decision-making into dialogue with scholarship on private enforcement and aggregate litigation. As I will argue, the recent procedural trends that have sparked trans-substantive critique among scholars may pose a special danger for the new harmed groups with which this Article is concerned. In highlighting these challenges, this Article aims to stimulate a research and advocacy agenda focused on how to make our procedural law function more effectively in the era of mass consumer profiling.

The argument proceeds below as follows. First, in Part I, I provide an overview of the algorithmic decision-making tools and strategies that are growing increasingly widespread in consumer markets. After explaining the business motivations that lead firms to make predictions about their customers, I outline the various tranches of data they can use to do so. I then survey the varieties of individualized decisions that can result from such profiling tools, including price, customer service, advertising, and credit.

Part II focuses on the distributive ramifications of algorithmic decision-making. I begin by arguing that the resulting winners and losers are best understood at the level of groups, rather than as a sea of atomized individuals. After outlining my theoretical approach to distributive justice — which centers the importance of democratic deliberation for updating and applying our distributive norms — I consider why

Price and Prejudice: The Case against Consumer Equality in the Information Age, 10 LEWIS & CLARK L. REV. 559 (2006) (arguing against an equal-pricing mandate).

⁵ See, e.g., Akiva A. Miller, *What Do We Worry About When We Worry About Price Discrimination? The Law and Ethics of Using Personal Information for Pricing*, 19 J. TECH. L. & POL'Y 41 (2014); Rory Van Loo, *Helping Buyers Beware: The Need for Supervision of Big Retail*, 163 U. PA. L. REV. 1311 (2015); Lauren Henry Scholz, *Algorithmic Contracts*, 20 STAN. TECH. L. REV. 128, 143, 163 (2017).

some of the specific consumer profiling strategies we know about are likely to raise serious concerns. Part II concludes by arguing that helping similarly situated consumers or “invisible allies” find each other will likely be crucial for ensuring that our collective conversation adequately confronts these new tools. Part III is concerned specifically with litigation as a potential (but perhaps unrealized) route toward this sort of democratic deliberation. I begin by explaining my decision to focus on litigation at all, and aggregate litigation in particular. Drawing from civil procedure literature, I call attention to recent trends in our procedural law that threaten to hobble the ability of aggregate litigation to perform this deliberative function in the context of consumer profiling claims. I then sketch out a procedural toolkit that might nevertheless exist under current law for those who advocate on behalf of the new harmed groups.

I. A New Generation of Consumer Profiling: Everything Everywhere All at Once

In order to assess the potential harms associated with algorithmic consumer profiling, it is necessary to get a basic sense of what is happening on the ground. Because the evidence is so incomplete and consumer markets themselves are so varied, a comprehensive accounting of all the uses of such tools is beyond the scope of this Article. A major motivating force behind this project is the vastness of what we don’t know and the dangers posed by such gaps. But enough information exists — from a mix of empirical academic research, news reporting, and economic theory — to develop at least a stylized picture of how some businesses are using these automated profiling tools in their interactions with consumers.

This Part begins by outlining a few ideal types of consumer profiling. This discussion develops a functionalist account of the kinds of questions that businesses are interested in asking and can now answer with a high degree of accuracy. After establishing some of the broad goals of consumer profiling from a business perspective, I outline some of the key sources of data used for such purposes. This discussion highlights both

the breadth and depth of the relevant data inputs.

I then give an overview of the different sorts of *outputs* that can result from such profiling, *i.e.*, the decisions that consumer-facing businesses can now render in an individualized and automated fashion. These include advertising, pricing, and the outcomes of customer complaints. Throughout the discussion of automated decisions, I aim to (a) explain why a business might opt to pursue such strategies and (b) to flesh out what we can learn from the concrete examples that have been uncovered. I will also try to tease out what is new about these business practices vis-a-vis the pre-machine learning era.

A. What Businesses Are Trying to Learn or Predict About Their Customers

When businesses use algorithmic tools to profile their current or potential customers, they are almost always doing some version of the same thing: making predictions about future behavior. Of course, when we zoom in, there is a great deal of variety in the consumer behaviors or attributes businesses are interested in modeling. Sometimes, businesses aim to arrive at a global or holistic assessment about a consumer. Many merchants, for example, use a global “customer lifetime value” (or CLV) score.⁶ These scores attempt to project how much profit the business can expect to derive from a given customer over the lifetime of the business relationship.

These macro-level assessments can also focus on the chance of some downside risk materializing, be it a consumer defaulting on a loan or engaging in other “costly behaviors” like complaining or returning items.⁷ Credit scores are perhaps the best-known example of these risk assessments. For

⁶ Khadeeja Safdar, *On Hold for 45 Minutes? It Might Be Your Secret Customer Score*, WALL ST. J. (Nov. 1, 2018), <https://www.wsj.com/articles/on-hold-for-45-minutes-it-might-be-your-secret-customer-score-1541084656> [https://perma.cc/N8AM-ADU8].

⁷ See generally Yonathan A. Arbel & Roy Shapira, *Theory of the Nudnik: The Future of Consumer Activism and What We Can Do to Stop It*, 73 VAND. L. REV. 929, 968 (2020); Joseph Turow et al., *Open to Exploitation: American Shoppers Online and Offline*, ANNENBERG PUB. POL’Y CTR. UNIV. PA. (June 1, 2005), <https://scispace.com/pdf/open-to-exploitation-america-s-shoppers-online-and-offline-i578cc9mbl.pdf> [https://perma.cc/8V4U-CJKM].

businesses that have a special interest in avoiding some negative outcome — like a default, a money-laundering transaction, or some major reputational harm — a risk-focused assessment tool can be very valuable.

At other times, a business will be interested in predicting a much narrower response to a specific stimulus, such as: what sort of price or special deal might induce a customer to switch from a competitor?⁸ Or, conversely, how much more can we charge our existing customer without losing them? What if we enforce rather than waive a late fee?

Businesses might also try to estimate the probability that a customer *currently* possesses some attribute. Algorithms designed to identify pregnant, Jewish, or LGBT consumers may be able to answer a yes-or-no question about a current state.⁹ But even tools like these still help project customers' *future* conduct.

For example, if a seller decides to reject some customers' retail returns and not others', the question isn't merely whether the consumer is telling the truth about the product's defect. For the seller, it is also important (perhaps more important) to know how a given customer is likely to react to a positive or negative outcome.¹⁰ These predictive assessments may bear no relation to the *prima facie* explanation the seller provides, *e.g.*, that a "fraudulent return alert" was triggered. Sellers may also want to know who precisely their consumers are now so they can determine who is likely to need, or perhaps pay more for, certain products in the near future.

In short, businesses use algorithms to estimate probabilities about their customers, making predictions about a range of possible consumer actions and their likely impact on the bottom line. Below, I will outline some of the data inputs that allow businesses to pursue these kinds of predictions.

⁸ See, *e.g.*, Miller, *supra* note 5, at 58–59.

⁹ Zeynep Tufekci, *We Need to Take Back Our Privacy*, N.Y. TIMES (May 19, 2022), <https://www.nytimes.com/2022/05/19/opinion/privacy-technology-data.html> [<https://perma.cc/T4DS-EWTW>]; Michal Kosinski, David Stillwell & Thore Graepel, *Private Traits and Attributes Are Predictable from Digital Records of Human Behavior*, 110 *Proc. Nat'l Acad. Sci.* 5802 (2013).

¹⁰ Rory Van Loo, *The Corporation as Courthouse*, 33 *YALE J. REG.* 547, 565 (2016).

B. What Sorts of Data Do Businesses Use?

In their quest for ever more accurate predictions and the profits they enable, sellers have begun integrating vast quantities of data into virtually every stage of their decision-making. But where does this data come from, and what sort of information does it contain? The answers to both these questions reveal clearly how difficult it is to assess the new generation of harms faced by consumers in the era of algorithmic profiling.

It is worth first saying a brief word about the role that information privacy law does and does not play in my overall analysis. I do not approach the puzzle of AI-era consumer protection solely as a “privacy” or “privacy law” issue (though I do draw from the privacy scholarship). Those who don’t care or think they don’t care about protecting their “privacy,” “data,” or “personal information,” *per se*, might nonetheless react with indignation if they learned they were being charged more for a particular product, denied a discount, or selectively refused leniency. Given what I will argue is the essential role of bottom-up democratic deliberation and activism in ensuring a fair marketplace, my framing of the problem attempts to capture as broad a swath of interested citizens as possible. I want the stakes to be clear, *i.e.*, that these tools help decide the new winners and losers in our economy.

Yet there are also analytical, rather than *realpolitik*, reasons to nudge the discussion away from a laser focus on privacy. As this Section will explain, much of the data that businesses use to make decisions about a given consumer *may not be generated by that consumer at all*. Instead, data from and about other people often proves even more important in modeling the likely behavior of a given consumer, irrespective of how much “personal” information he has chosen to give up.¹¹ Even when a piece of relevant data does originate in part with the individual consumer who is subject to a marketplace decision, it may have started its life-cycle in the form of seemingly innocuous or already-public information—that is, before being

¹¹ See generally Salomé Viljoen, *A Relational Theory of Data Governance*, 131 YALE L.J. 573, 586 (2021); Alicia Solow-Niederman, *Information Privacy and the Information Economy*, 117 NW. U. L. REV. 357, 362 (2022).

fed into a model that spit out a bespoke treatment. While scholars have increasingly sought to grapple with the *privacy* impacts¹² of non-private information used in this way, they have also highlighted, as does this Article, the structural and material consequences of consumer profiling. Thus, “privacy” may simply not be the best framework for apprehending the nature of this particular problem. Instead, at least when we are talking about run-of-the-mill marketplace transactions, I suggest that consumer protection and/or political economy may be at least as useful a set of lenses for understanding and ultimately tackling the attendant harms.

C. Consumer Profiling in Context

Below, I turn to the data flows that enable fine-grained consumer targeting. What sort of data is it? Just about any data that sellers can acquire in a (putatively) legal¹³ and cost-effective manner can influence their decision-making calculus.

1. “Matt” and His “Own” Data

Consider the case of Matt, a young man in a big city who shops for clothing both online and in-person at Retailer A. Below is a non-exhaustive list of the kinds of data that could flow *directly from Matt’s own conduct* or that Matt himself *has generated or made available*, and that Retailer A might use to learn more about him and tailor its decision-making appropriately. (Later on, I will discuss what can be gleaned from the data generated by people other than Matt.)

Purchase History: Unsurprisingly, Retailer A will want to know what Matt has bought from them in the past. But Matt’s purchases will also reveal other patterns, such as where, when, and how he shops for one kind of thing or another, and how much in dollar terms. Matt’s purchase data will also reveal trend lines, *e.g.*, whether his average purchase amount (*e.g.*, per visit) has been going up or down and for how long.¹⁴ Matt’s purchase data may reveal how he tends to respond to nudges

¹² See generally Viljoen, *supra* note 11; Solow-Niederman, *supra* note 11.

¹³ While this section does not focus on the substantive law governing such data practices, I note relevant legal frameworks in passing.

¹⁴ An on-background conversation between the author and a former executive in the consumer profiling industry focused on the importance of this data point.

like temporary offers. Crucially, neither the law nor technology would necessarily limit Retailer A to using data that it has itself collected. Rather, Matt may also have generated all or some of the above information in his interactions with *other retailers* in such a fashion that Retailer A can make use of it. (Third-party data aggregators play a key role in this part of the ecosystem, as discussed below.) Thus, even if Matt is relatively new to Retailer A, a much richer data set from his other shopping experiences may nonetheless allow Retailer A to model his likely behavior with great accuracy.

Geolocation: Retailer A can learn about Matt's physical movements at virtually every level, from the hyper-local to the global. Starting narrow, Retailer A can use surreptitious "Bluetooth beacons" to track Matt's movements within its brick-and-mortar stores down to the level of feet or even inches.¹⁵ Retailer A could potentially figure out how long Matt spends in a particular aisle, the fitting room, or anywhere else in the store.¹⁶ Matt's broader use of Wi-Fi, mobile data, and GPS can supplement the picture of his overall movements, revealing what other stores he visits and when. And as Justice Sotomayor acknowledged in *United States v. Jones*, the trail of geolocation data Matt leaves behind may reveal much more, adding up to a startlingly detailed record of his life.¹⁷ Unlike the government actors in *Jones*, private actors face few meaningful constraints on their ability to collect, share, and draw conclusions from Matt's geolocation data.¹⁸

Online Social Activity: A good deal of Matt's social activity will also be legible in the trail of data he leaves behind online. Matt may be an active participant in social networks that reveal his friendships, political opinions, media consumption, hobbies, and anything else he shares through his posts, affiliate groups, or "likes." (In fact, his social media activity probably

¹⁵ Michael Kwet, *In Stores, Secret Surveillance Tracks Your Every Move*, N.Y. TIMES (June 14, 2019), <https://www.nytimes.com/interactive/2019/06/14/opinion/bluetooth-wireless-tracking-privacy.html> [https://perma.cc/2T4R-KHGE].

¹⁶ *Id.*

¹⁷ 565 U.S. 400, 413 (2012) (Sotomayor, J., concurring).

¹⁸ Often, the only legal barrier (if any) is a thin, easy-to-circumvent consent requirement. See generally Daniel J. Solove, *Privacy Self-Management and the Consent Dilemma*, 126 HARV. L. REV. 1880, 1890 (2013).

reveals far *more* than he believes he's choosing to share, thanks to the data furnished by others. *See infra* Part I.C.2.) Matt may also belong to and/or post in online forums that reveal even more intimate information, e.g., about sexual habits or medical problems. And if he thinks he's participating anonymously or in an untraceable fashion, odds are he isn't.¹⁹

Complaints and Customer Service History: If Retailer A has any history at all with Matt, it will likely have kept a record of his past complaints, returns, and other interactions with customer service. As with his purchase history, Matt's track record with customer service at other retailers may also be available, at least in the aggregate or through composite profiles sold by third parties. The existence of such records may not itself be new. But the scope of information available in an instant, the sophistication of the resulting models, and their easy integration with other automated systems (e.g., at the point of sale) now allow for even brick-and-mortar retailers to tailor their customer service interactions on an individualized basis. When customers interact with a retailer by phone or online, both the data they generate and the opportunities for hyper-individualized service — which may or may not be to their benefit, as discussed below — increase dramatically.

Financial Situation and Habits: Credit reporting is one of the oldest forms of consumer profiling and long predates the current era of AI-driven individualization²⁰. Arguably, the credit reporting system that emerged in the late 20th century is the direct ancestor of the “everything reporting” system that has taken root in the early 21st century. Partly because of its long pedigree, credit reporting is also one of the most-regulated forms of consumer data harvesting, thanks to the Fair Credit Reporting Act (FCRA).²¹ For Retailer A, following the strictures of the FCRA for every consumer would probably not make good business sense even if legally viable. But it can still acquire other financial information about Matt that wouldn't fall under the ambit of FCRA, from public records of liens or

¹⁹ *See, e.g.,* Carissa Véliz, *Self-Presentation and Privacy Online*, 9 J. PRAC. ETHICS 2, 39 (2022).

²⁰ Citron & Pasquale, *supra* note 1, at 16–18.

²¹ 15 U.S.C. § 1681 et seq. The statute limits the kinds of information that can be included on a consumer's credit report and the situations in which the report can be accessed by third parties, among other provisions.

judgments to transaction histories Matt reveals through third-party loyalty programs.

Browsing History and App Use: If Matt uses Retailer A's app or website, Retailer A will already have access to a valuable cache of information about his searches, browsing habits, and other such information. But Retailer A can also acquire (via third-party aggregators) a treasure trove of information about Matt's *global* browsing history and app use. An entire industry exists to collect this information more-or-less invisibly and then sell it to Matt's once and future counterparties.²² Essentially, any time Matt visits a website or uses an app that in turn sends data to one of these third-party aggregators,²³ information is added to a three-dimensional profile of his online activity. While Matt may not conceive of such online activity as any kind of knowing disclosure of information, the resulting data is nonetheless data that his own conduct generates and that the law, in many instances, views as fair game for others to analyze and monetize. From life events like births or deaths to political interests, media consumption, anxieties, dreams, attentiveness to core tasks, and much more, Matt's browsing and app history could reveal a remarkably detailed picture of who he is and how he uses his time.

Demographics: Basic demographic information like age, gender, race, and sexual orientation can inform all sorts of decisions about how to market to and transact with consumers. (While some such decisions might violate public accommodations laws, businesses often enjoy a surprising amount of latitude to adjust their treatment of a given consumer based on his demographic categories.)²⁴ Matt may provide some of this information directly to Retailer A if he signs up for a loyalty program or some other company-specific service. More likely, however, he has provided this information

²² See, e.g., Complaint for Permanent Injunction and Other Relief at 10, *FTC v. Kochava Inc.*, No. 2:22-cv-377 (D. Idaho May 4, 2023) [hereinafter *Kochava Complaint*].

²³ By integrating these third-party software development kits or SDKs, app or website designers get reciprocal benefits in terms of the ability to track and analyze data about their own app or website's users.

²⁴ While refusing to contract may be unlawful, a wide range of differential treatment has been held not to implicate such statutes. See Suja A. Thomas, *The Customer Caste: Lawful Discrimination by Public Businesses*, 109 CALIF. L. REV. 1 (2021).

in other contexts such that it has been aggregated and made available for sale. Even absent explicit disclosures, some of Matt's group memberships (e.g. a mention of Outlaws or a Black Students Association on LinkedIn) might work just as well. Businesses that want to tailor their decision-making are likely interested in a much more detailed demographic profile of Matt than the above categories suggest—and they can probably get one thanks to comparisons with others and their data, as discussed in the following Subpart.

Language Use: If Matt has ever called or emailed Retailer A in any capacity, his use of language may provide a rich trove of data about his needs and wants, his emotional states, and his education level, among many other insights. Even if Matt is a brand-new customer, he has likely left behind a trail of language-related data elsewhere.

Biometric Data: The category of biometric data can include all sorts of immutable information specific to a person's biological makeup or body, including their facial geometry, fingerprint, gait, and voice. Consider the most widely used variety of biometric data in commercial settings — facial recognition data — as an example of what is now possible. Simply by walking around in public, Matt has exposed himself to thousands of cameras that may now be able to cheaply and efficiently compute an individualized, mathematical representation of his face. This data can in turn be cross-referenced with any identifiable picture of Matt that appears anywhere online, allowing for near-instant identification. This is the general model that underlies the Clearview AI point-and-shoot facial recognition app, its rough analog PimEyes, and other similar products.²⁵ While some U.S. jurisdictions would prohibit Retailer A from using this sort of technology on Matt without his informed consent,²⁶ most would not as of this writing. At a basic level, in-store cameras and associated facial recognition technology can quickly establish that Matt is present and associate him with all the other kinds of data

²⁵ Kashmir Hill, *A Face Search Engine Anyone Can Use Is Alarmingly Accurate*, N.Y. TIMES (May 26, 2022), <https://www.nytimes.com/2022/05/26/technology/pimeyes-facial-recognition-search.html> [<https://perma.cc/M2WH-BB76>].

²⁶ See, e.g., Illinois's Biometric Information Privacy Act ("BIPA"), 740 ILL. COMP. STAT. 14/15(b).

Retailer A possesses about him (does he buy more when interacting with female shop assistants?), including some categories not yet mentioned (*e.g.*, its assessment of his past suspicious behavior, likelihood of shoplifting, or actual arrest records).²⁷

Technical Information: When Matt shops online, Retailer A may be able to learn his (a) computer or phone model; (b) operating system; (c) browser; (d) mode of connecting to the internet, and (e) IP address, among other data points. This data may be useful to identify Matt and tie him to other Matt-related information. Other, more granular data points — such as the battery level of Matt’s phone²⁸ or the fact that he is using an unfamiliar machine — might point to an emergent situation that is relevant to the seller’s decision-making in the moment.

Below, I will explain how the data generated by *other* people can be combined with data directly traceable to Matt in order to yield actionable insights for Retailer A. But before broadening the scope of the inquiry, it is worth pausing to consider a few insights that emerge from the illustrative list above.

First, none of the above categories exists in isolation. Retailer A or third-party aggregators it hires will instead combine information from multiple tranches of data in order to build out a fuller picture of who Matt is. Thanks to Matt’s search history, it can become clear not just *that* Matt visits a multi-practice medical office building once every week — ascertainable from his geolocation data — but *why*. (“Siri, help me find the best downtown psychiatrist with a specialty in depression.”) Once we know about Matt’s job as a judge or

²⁷ Jeffrey Dastin, *Rite Aid Deployed Facial Recognition in Hundreds of U.S. Stores*, REUTERS INVESTIGATION (July 28, 2020), <https://www.reuters.com/investigates/special-report/usa-riteaid-software/> [https://perma.cc/K2DS-L9YC]; Flores v. Motorola Sols., Inc., No. 1:20-CV-01128, 2021 WL 232627, at *1 (N.D. Ill. Jan. 8, 2021) (describing a business whose service cross-references arrest records with biometrics).

²⁸ Amit Chowdhry, *Uber: Users Are More Likely to Pay Surge Pricing if Their Phone Battery Is Low*, FORBES (May 25, 2016), <https://www.forbes.com/sites/amitchowdhry/2016/05/25/uber-low-battery/> [https://perma.cc/9V35-9LKT]; VICE Staff, *Uber Accused of Charging People More If Their Battery Is Low*, VICE NEWS (Apr. 11, 2023), <https://www.vice.com/en/article/uber-surge-pricing-phone-battery/> [https://perma.cc/849Z-WRDX].

doctor or EMT, along with his putative working hours, perhaps his three-hour-a-day Candy Crush habit starts to appear in a different light. One result of this cross-fertilization, as legal scholar Alicia Solow-Niederman and others²⁹ have recognized, is that it can be difficult *ex ante* to figure out what kind of data is or should be considered “sensitive.”

Second, Matt has almost no ways of knowing which of these types of data will be available or useful to Retailer A at any given time. In some jurisdictions, Matt may have some limited rights to learn about the sources of data used to render a decision about him under a narrow set of circumstances. Even in those states, many consumer transactions would fall outside the purview of the relevant laws. And where the data isn’t “Matt’s” at all, the situation to which we will turn below, many of those narrow rights more or less evaporate.

Third, it may be irrelevant whether or not Retailer A itself collects or gains access to the raw data Matt has left behind. Instead, third-party data aggregators or processors,³⁰ who are often repeat players, can save Retailer A time and money by providing composite analyses in an actionable form (*e.g.*, Matt spends big on Fridays). Such strategies become supercharged when these third parties integrate the data of millions of others.

2. Matt’s Data Combined with Data from Other Consumers

The stylized examples above adopted something of a fiction: that Matt’s data is *ever* truly considered or processed without being combined with the data of others. In a technical sense, virtually every AI model relevant to the retail context would have to be built and tested using data from a substantial number of consumers. If Retailer A purchased a model from a software company that claimed to predict Matt’s demand for shoes at a given moment using data about Matt’s past purchases, the model capable of doing that almost certainly emerged thanks to the data of many non-Matt consumers.

Thus, as privacy scholars and others have recognized,³¹ if

²⁹ Solow-Niederman, *supra* note 11, at 22–24.

³⁰ Solow-Niederman, *supra* note 11, at 54–55 (discussing information processors).

³¹ See, *e.g.*, Viljoen, *supra* note 11, at 607–08.

Retailer A wants to make automated predictions about Matt's likely shoe-buying behavior, it will have many tools in its arsenal even if Matt jealously guards much of his "own" data. Thanks to the correlations that have been discovered about the features and behaviors of *others* and their likelihood of purchasing shoes (or some brand of shoes, or shoes at some dollar amount), much about Matt's likely behavior can still be modeled. The upshot is that models of consumer protection or privacy law that depend on Matt's own data choices may leave him out to dry.³² For this reason, it is worthwhile as an analytical matter to consider how Retailer A might use data *not* generated or revealed by Matt to personalize its treatment of him. Failing to grapple with this dynamic has resulted in the largely toothless, consent-based legal protections consumers enjoy today.³³

Below, I have labeled the following categories with a "-plus" to emphasize the special importance of other people's data to predictions that can be made about Matt.

Geolocation-Plus:

Even if Matt has reverted to a "dumb phone" and thinks he has escaped most of the geolocation dragnet, it wouldn't take much to yank him back in. By otherwise obtaining information about his address (public records, IP addresses, credit card records), data aggregators can cross-reference information from *other* people's much richer set of geolocation data. As a result, models can now do much more than rough correlations between zip codes and demographic features or income levels. If Matt's partner Monica plays fast and loose with her data, Retailer A may be able to learn not just *that* they live together, but also that they travel together, and even what they talk about. (Indeed, many examples of phones appearing to "listen" to us likely result from such techniques.)³⁴

While these examples might seem anodyne if vaguely creepy, few barriers would prevent them from taking a dark turn. Leaving our shoes behind for a moment, the same sort of

³² *Id.*

³³ Solove, *supra* note 18.

³⁴ See, e.g., Coco Khan, *Is My Phone Listening to Me? We Ask the Expert*, THE GUARDIAN (Oct. 29, 2021), <https://www.theguardian.com/lifeandstyle/2021/oct/29/is-my-phone-listening-to-me-we-ask-the-expert> [https://perma.cc/QJA3-54U6].

geolocation information could reveal that Matt lives with, say, someone going through a personal crisis, a difficult illness, or a financial crunch. Such knowledge might be hugely relevant to how Retailer A treats Matt and what kind of behavior it expects from him. Matt's other neighbors could yield even more useful insights. Perhaps his apartment building has lost power or flooded, or his block has experienced a sudden uptick in crime or COVID. It is easy to see why certain sellers might find such information valuable. Perhaps a seller like Retailer A might itself not understand *why* its model is flashing an “upsell” or “deny the return” light in the brick-and-mortar store—or why its website or app is displaying a particular price. But if the seller or the designer of its predictive model is interested in how much Matt will pay for a good or service right now, the exigencies described above could prove quite relevant.

Biometric Data-Plus:

Facial biometric data has use cases far beyond recognizing who someone is. Using measurements of people's facial movements, retailers can now undertake “sentiment analysis” that theoretically reveals how a person is feeling at a given moment. Walgreens, for instance, used cameras embedded in its beverage refrigerators to assess the moods of customers who walked by, with the aim of adjusting the displayed ads accordingly.³⁵ Companies have also claimed to be able to measure qualities like tenaciousness or capacity for teamwork based on individuals' facial movements.³⁶ Other companies now use analogous sentiment analysis tools to measure the tone of voice a consumer uses on its customer service phone lines and adjust their treatment of the consumer accordingly.³⁷

³⁵ Joseph Pisani, *Cameras That Guess Your Age, Sex and Mood Pop Up in Walgreens, Other Retailers*, NBC N.Y. (Apr. 23, 2019), <https://www.nbcnewyork.com/news/local/cameras-store-shelves-consumer-behavior-walgreens-kroger-supermarket/1770749/> [<https://perma.cc/K66Y-7BXM>].

³⁶ There is debate about whether such claims have a scientific basis. Angela Chen & Karen Hao, *Emotion AI Researchers Say Overblown Claims Give Their Work a Bad Name*, MIT TECH. REV. (Feb. 14, 2020), <https://www.technologyreview.com/2020/02/14/844765/ai-emotion-recognition-affective-computing-hirevue-regulation-ethics> [<https://perma.cc/FPS4-59TB>].

³⁷ See, e.g., Pradeep Vasudev, *What Is Speech Analytics and How Does It Work?*, SPRINKLR CXM GUIDE (Jan. 31, 2024),

In both the above examples, so long as the seller or the designer of the model had access to the faces or voices of lots of *other* people and verified information about how they were feeling, the same correlations can (they claim) be applied to Matt even when he shows up to a store or places a phone call for the very first time. Matt's choices about safeguarding his "own" data in the past simply don't matter to Retailer A's ability to use these tools at that moment.

Language Use-Plus:

Many of the same tools that exploit correlations tied to people's facial movements or tone of voice can draw similar conclusions based on someone's use of language. One company, for instance, discovered that the frequency with which a consumer used the word "please" on social media *in general* was associated with a higher risk of loan non-repayment (perhaps because of a history of making requests for help).³⁸

Technical Information-Plus:

Even if Matt is a particularly assiduous user of VPNs or other masking devices, basic information about the device he uses to shop on Retailer A's website or app will likely still be ascertainable. The same may also be true about the device he carries with him as he enters Retailer A's brick-and-mortar store, partly thanks to the Bluetooth beacons discussed above. Sellers long ago recognized correlations between obvious data points like operating systems and more difficult-to-measure ones like wealth or willingness to spend. Such correlations likely explain why Apple users sometimes get steered toward higher-priced products,³⁹ and may influence the price Matt gets charged.

This Subpart has catalogued the sources of data that sellers can collect, analyze, and ultimately mine for useful insights about their consumers. The next Subpart considers what a

<https://www.sprinklr.com/cxm/speech-analytics> [https://perma.cc/FQV2-B58D].

³⁸ Yanhao Wei et al., *Credit Scoring with Social Network Data*, 35 MKTG. SCI. 234 (2016).

³⁹ *Report: Orbitz Steers Mac Users to Pricier Hotels*, CBS News (June 25, 2012, 9:19 PM), <https://www.cbsnews.com/news/report-orbitz-steers-mac-users-to-pricier-hotels/> [https://perma.cc/G8B3-6XFX].

business might do with such insights.

D. What Sorts of Individualized Decisions Result from Consumer Targeting?

What, if any, are the consumer-facing decisions that might differ based on a seller's data-driven personalization strategies? In order to determine who the potential winners and losers might be (a focus of the inquiry in Part II) and assess possible avenues of redress (Part III), we must canvass what sorts of things can be won or lost.

Price:

The seemingly straightforward term “price” masks a good deal of complexity. Businesses not only display different sticker prices to different customers, but can also use customized search results to “steer” customers to higher- or lower-priced items,⁴⁰ tweak final prices using unexpected add-on fees, or offer individualized discounts.

Even when the inquiry is limited to disparities in the displayed sticker price, variations can take many forms. Broadly, the term “dynamic pricing” refers to prices that change over time, while “personalized” or “individualized” pricing refers to prices that differ from consumer to consumer (at any given time).⁴¹ These pricing strategies sometimes go together; a business can charge a given person prices that change over time and differ from the prices shown to others.⁴²

To begin, though, let us consider the simplest case of price discrimination: when individual consumers are “shown inconsistent [sticker] prices for the same product” at the same time.⁴³ E-commerce helps make this strategy possible. Along with reams of data about consumers past and present, businesses on the internet have the advantage of interacting with shoppers on a one-on-one basis. As one study put it, when

⁴⁰ Aniko Hannak et al., *Measuring Price Discrimination and Steering on E-commerce Web Sites*, in PROCEEDINGS OF THE 2014 INTERNET MEASUREMENT CONFERENCE 305, 305 (2014), <https://personalization.ccs.neu.edu/static/pdf/imc151-hannak.pdf> [<https://perma.cc/BQM6-B98Y>].

⁴¹ See generally Miller, *supra* note 5.

⁴² *Id.*

⁴³ Hannak et al., *supra* note 40, at 307.

we shop online, “we shop alone.”⁴⁴ We usually don’t know what our neighbors pay for the same flight, hotel, or bag of dog food.

From the perspective of a profit-maximizing seller, customized pricing offers significant upside potential. When products are offered at a single uniform price, some consumers will value the product in excess of what they are charged, representing a welfare gain for the consumer and a missed opportunity for the seller. This difference aggregated across all consumers is the “consumer surplus.”⁴⁵ At the other end of the spectrum of pricing practices, “first-degree” or “perfect” price discrimination would allow sellers to “extrac[t] the entire consumer surplus” by ensuring that prices are calibrated to the idiosyncratic preferences of each consumer.⁴⁶

It has long been assumed that first-degree price discrimination would be impossible. But for high-volume, repeat-player businesses that aggregate and analyze data on millions of consumers at low cost — or sellers who buy data from such businesses — it is now possible to get closer to first-degree price discrimination than ever before. This data helps businesses predict with a high degree of confidence who consumers are, what they need, and what influences their behavior.⁴⁷

In Part II, I will discuss in greater depth the potential distributive outcomes that might flow from this sort of individualized pricing. A key point to which I will return is that there is likely no one-size-fits-all answer about who the winners and losers are in the consumer marketplace. While lower-income shoppers may be offered lower prices some of the time for certain products, such an outcome is far from a guarantee.

The above discussion explains why some businesses can, and may want to, use individualized pricing in pursuit of

⁴⁴ Aditya Karan et al., *Your Browsing History May Cost You: A Framework for Discovering Differential Pricing in Non-Transparent Markets*, in PROCEEDINGS OF THE 2023 ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 717, 717 (2023).

⁴⁵ Klock, *supra* note 4, at 327–28.

⁴⁶ Edwards, *supra* note 4, at 566.

⁴⁷ Cf. Vian Bakir, *Psychological Operations in Digital Political Campaigns: Assessing Cambridge Analytica’s Psychographic Profiling and Targeting*, FRONTIERS IN COMM’N, Sept. 2020, at 1.

greater profits. But even assuming data quality, model accuracy, and cost effectiveness, individualized pricing techniques carry some serious potential risks for sellers. Consumers might, for instance, judge a business harshly for individualized pricing or find ways to claw back some of the surplus through strategic behavioral responses. While public awareness as of today is under-developed enough that these risks likely do not outweigh the potential benefits for most businesses, even the specter of public outcry or strategic consumer behavior puts a thumb on the scale in favor of opacity. Thus, whatever the evidence suggests about variability in sticker prices, our concern must extend to more under-the-radar strategies that affect the prices consumers *actually* pay, from steering techniques to explicitly price-modifying practices like targeted discounts, temporary offers, or hidden surcharges.

Given the means and the motivation to individualize their prices, do businesses actually do so? Unfortunately, the reliable empirical data is thin and in some cases outdated. Developing a reliable picture of a company's pricing strategies without access to its internal algorithms is also extremely challenging. The strong incentives favoring secrecy and subterfuge make this task even harder. That said, between the few rigorous studies that exist and a range of additional "anecdotal," we can say with a reasonable degree of confidence that individualized pricing likely happens often enough to matter both to particular consumers and to overall economic outcomes.

In a landmark 2013 study, Mikians et al. were already able to show that a range of different online retailers were systematically engaging in personalized pricing.⁴⁸ Using a crowdsourcing tool they designed, Mikians et al. collected prices from over 600 online retailers, eventually identifying 21 that appeared to display different sticker prices for the same good at the same time. The research team then conducted a much more in-depth study of that subset of retailers, collecting

⁴⁸ Jakub Mikians et al., *Crowd-Assisted Search for Price Discrimination in E-Commerce: First Results*, in PROCEEDINGS OF THE NINTH ACM CONFERENCE ON EMERGING NETWORKING EXPERIMENTS & TECHNOLOGIES 1, 1 (2013).

a total of 188,000 prices from up to 100 products per retailer.⁴⁹ What they discovered confirmed the suggestive data that they'd earlier gleaned from the crowd. Of the subset of 21 retailers, 17 exhibited “near complete (100%)” price variation, *i.e.*, displayed non-uniform sticker prices for almost all sampled products. For some of the retailers, the researchers even found “100% coverage, pointing to the fact that price variations are a persistent and repeatable phenomenon.”⁵⁰ The group of frequent price customizers spanned many kinds of websites. Amazon.com, for example, used variable pricing on over half of its products, while Home Depot did so for more than a fifth. The size of the price differentials was also meaningful, with most of the retailers varying their prices between an average of 10 to 30% across their product lines.⁵¹ In all instances, variable prices were captured at the same time, meaning the differences were not attributable to generalized shifts in demand.

Although this research is more than a decade old, it helps establish something of an empirical baseline. Since then, researchers have continued to uncover evidence of customized pricing online. More recent empirical research supports two overall conclusions that dovetail with the results that Mikians et al. found: (a) a minority of online retailers seem to be responsible for most personalized pricing, but (b) when consumer-to-consumer price differentials show up, they are big enough to matter, especially in the aggregate.

A much more recent study of airline ticket and hotel prices (Karan et al. 2023) found, for example, that most sellers displayed equal prices most of the time, but the ones that customized their prices showed meaningfully different prices to the various behavioral “profiles” that the researchers had constructed. A random consumer would, on average, expect to see prices \$3.00 higher than would the most favored profile when shopping for hotels, and \$6.00 higher than would the most favored profile when shopping for flights.⁵² As the authors note, “[i]n non-transparent markets with millions of

⁴⁹ *Id.*

⁵⁰ *Id.*

⁵¹ *Id.*

⁵² Karan et al., *supra* note 44, at 717–18.

transactions, this is a significant impact.”⁵³ And while the overall incidence of individualized pricing remained low, it was higher than in some of the 2010s studies, with 4.5% of flight queries and almost 2% of hotel queries showing evidence of individualized pricing.⁵⁴

Perhaps the most dramatic demonstration of how widespread such pricing practices have become was a blockbuster 2025 investigation of pricing on the grocery-shopping app Instacart by Consumer Reports and Groundwork Collaborative.⁵⁵ The investigation revealed that Instacart was using AI tools to individualize prices for consumers shopping at a wide range of popular grocery chains, including Albertsons, Costco, Kroger, and Safeway.⁵⁶ Every survey participant with valid data was shown individualized prices for at least some products, with the app increasing the price for certain products “by as much as 23 percent” depending on the consumer.⁵⁷ And the aggregate effects were profound: “The price of the same basket of food at a Seattle-area Safeway on the Instacart platform, for example, ranged from \$114.34 to \$123.93—roughly a \$10 difference,” a difference that “could translate into a cost swing of about \$1,200 per year” for a family of four.⁵⁸

Even more recently, the same Consumer Reports team enlisted more than 170 volunteers to test the prices quoted by Uber and Lyft for the very same rides at almost exactly the same time.⁵⁹ Their findings revealed that Uber and Lyft tended to separate individual consumers into higher and lower “price

⁵³ *Id.* at 718.

⁵⁴ *Id.* at 724.

⁵⁵ Derek Kravitz, *Instacart’s AI-Enabled Pricing Experiments May Be Inflating Your Grocery Bill*, CR and Groundwork Collaborative Investigation Finds, CONSUMER REPORTS (Dec. 22, 2025), <https://www.consumerreports.org/money/questionable-business-practices/instacart-ai-pricing-experiment-inflating-grocery-bills-a1142182490/> [<https://perma.cc/QW8J-Q64L>].

⁵⁶ *Id.*

⁵⁷ *Id.*

⁵⁸ *Id.*

⁵⁹ Derek Kravitz, *Different Prices for the Same Ride: How Uber and Lyft Use AI to Get More Money Out of You*, CONSUMER REPORTS (June 16, 2026), <https://www.consumerreports.org/money/questionable-business-practices/uber-lyft-different-prices-for-same-ride-and-fake-discounts-a1093538909/> [<https://perma.cc/ZUP8-ZKMD>].

clusters” for the exact same ride, with a median difference of 42.4 percent. Of the 30 routes the researchers tested, *every single one* had “at least two price clusters separated by at least 5 percent.”⁶⁰ Taken together, these two Consumer Reports studies illustrate that even when buying essential goods like food or transportation, consumers must now be on alert for secretly individualized prices.

Adjudications:

In recent years, increasing attention has been paid to what scholar Rory Van Loo refers to as the “Corporation as Courthouse” in his 2016 piece.⁶¹ Whether a consumer seeks some compensation for a complaint, attempts to return an item, or challenges a late fee or some other penalty, Van Loo explains how consumer-facing companies often make decisions much like “adjudications” in the traditional sense. Moreover, companies invest substantial resources in their adjudicative apparatuses, both staffed and automated. Virtually the entire spectrum of “adjudications” covered by Van Loo could potentially be customized based on the types of consumer data discussed above. And substantial evidence exists that this is already happening, *e.g.*, through automated systems that decide whether to accept or reject retail returns.⁶²

Other Customer Service and Qualitative Treatment:

Data-driven profiling can also power a much broader range of decisions about how to treat a given consumer. Sometimes, the issue is how to apportion scarce customer service resources—put differently, who gets disproportionately more or better customer service. Bank of America, for instance, was found to route phone calls from customers with perceived higher “value” to a more “elite” line that provided better care.⁶³

Certain businesses, however, may actually target particular consumers for purposefully poor service. Prof. Suja Thomas has documented the creation of these “consumer castes” in a variety of contexts (and the failure of anti-discrimination law

⁶⁰ *Id.*

⁶¹ Van Loo, *supra* note 10.

⁶² *Id.* at 565.

⁶³ *Id.* at 565–66.

to grapple with qualitatively poor service).⁶⁴ More recently, Prof. Hila Keren has argued for private law to recognize “market humiliation” as a breach of the duty of good faith.⁶⁵

But both these scholars focus on examples where, at least partly, animus or hostility toward established identity categories was at play. Importantly, sellers can also use subpar customer service to drive away a set of customers they view as undesirable for more “economically rational” reasons. Some retailers, for example, have been found to use “intentionally poor service” to repel customers identified as “‘bottom feeders’—consumers who visit them mostly for bargains and return products too often.”⁶⁶ While such strategies may only make sense in some contexts, the fact that they *can* be market-rational may make them a particularly worthy object of analysis (as they can show up even in the absence of animus).

Advertising:

Targeted advertising has played a key role in driving the rise of such profiling tools. While advertisers and sellers have used broadly defined population segments for decades, it is now possible to (1) predict what sorts of people are most likely to respond to an ad, and (2) ensure that those people are the ones that see it. Moreover, the content of the ads themselves can sometimes be personalized.

At first blush, targeted ads may seem different in kind from the more clearly win-or-lose decisions covered in this Article, such as prices or the quality of customer service. But the line between targeted advertising and other forms of consumer-facing decisions is fuzzy. As mentioned above, sellers can “steer” buyers toward higher- or lower-priced items using customized search results. Targeted advertising works similarly. Even if a product or service is technically available to consumer X and consumer Y, a seller can make it much more likely that it is consumer X who would hear about it. This is one way in which targeted advertisement is bound up in other sorts of consumer-facing decisions. Furthermore, thanks to behavioral profiling, advertisements can also be used to match

⁶⁴ Thomas, *supra* note 24.

⁶⁵ Hila Keren, *Beyond Discrimination: Market Humiliation and Private Law*, 95 U. COLO. L. REV. 1, 87 (2024).

⁶⁶ Turow et al., *supra* note 7, at 32 n.22.

exploitative schemes like usurious loans or expiring free trials with the customers most likely to fall for them (in general or right now).⁶⁷ Finally, as discussed below, advertising can reveal how sellers *themselves* understand groups of consumers that they wish to treat differently, giving us some purchase on how the slicing and dicing of consumers plays out in practice.

Credit, Insurance, and Whether to Contract at All:

Credit decisions encompass (at least) two genres: the decision whether to extend credit to a particular would-be borrower, and on what terms. Both decisions have long been individualized based on, *inter alia*, the information in an individual's credit report. Similarly, insurance companies must also decide both whether to offer a policy at all and, if so, with what premiums.⁶⁸ Other than credit and insurance, there are other contracting decisions that can also take a yes-or-no form, such as whether a request for an Uber ride, a restaurant reservation, or an AirBnB rental is granted. Interestingly, the refusal to contract with a consumer at all may expose some businesses (though not others) to liability under the federal civil rights laws, even if other differential-treatment decisions might fall outside the statutes' purviews.⁶⁹

II. Individualized Treatment, Invisible Patterns of Harm

Having outlined the range of available data and decision-making strategies, we can now turn to our assessment of who the potential winners and losers might be. This Part begins by explaining the group-based framework that animates this Article's normative account and its prescriptive suggestions. I aim to demonstrate why a focus on patterns of shared, collective harms — “the new harmed groups” — makes sense from both an analytical and strategic perspective.

I then attempt to identify who some of these groups might be in particular consumer contexts. Again, the paucity of

⁶⁷ See generally Van Loo, *supra* note 5 (noting categories of consumers who are vulnerable to such techniques).

⁶⁸ See Max N. Helveston, *Reining in Commercial Exploitation of Consumer Data*, 123 PENN. ST. L. REV. 667, 676 (2019) (discussing consumer profiling in insurance decisions).

⁶⁹ See, e.g., Thomas, *supra* note 24, at 12.

reliable data and the heterogeneity of market contexts makes it impossible to issue any sweeping verdicts. Instead, through a combination of illustrative examples and theoretical observations, my hope is to build a persuasive case that these tools are a cause for serious concern at least some of the time.

Importantly, not every pattern of advantage or disadvantage might bother us if fully known. Others might cause immediate outcry. As the examples will illustrate, there is also a vast, gray muddle in between, where benefits and burdens are distributed according to logics that our society hasn't yet confronted head-on. It is precisely for this reason that members of harmed groups must have the opportunity to come together and articulate moral or legal arguments in opposition to such treatment. Drawing from scholarship on privacy law, social theory, and democratic theory, I explain why harmed groups must play an essential role in the deliberative process we use to update both (a) our empirical knowledge about the state of play in the marketplace and (b) our shared notions of proper and improper marketplace conduct. Without such deliberation, the conceptual categories and legal tools we use to police the consumer marketplace will ossify. If we fail to meet the moment, the invisible, large-scale social sorting that marketplace actors have already unleashed will continue on autopilot.

The final Subpart will then argue for a mechanism of action that any resultant legal and policy reforms ought to focus on as a first step: stimulating relationships among similarly harmed parties. While the literature on data privacy and consumer protection rightly draws attention to the hidden and specialized knowledge possessed by data-driven businesses, I argue that there are also bottom-up insights, *e.g.*, about how consumer harms play out and why they may be wrongful, that may be missing from the conversation. Collective knowledge-building and action by these new harmed groups may be the only feasible alternative to the invisible and unaccountable patterns of harm that would otherwise remain the default.

A. The New Harmed Groups

At the outset, it is reasonable to ask whether it makes sense to think about affected individuals as members of "groups" at

all. Some proponents of data-based decision-making, such as Lior Strahilevitz, celebrate the move toward more individualized treatment as a major improvement on invidious group-based heuristics (e.g. racism and sexism).⁷⁰ Indeed, anti-discrimination doctrine frequently frames the project of anti-discrimination law as an attempt to *promote* individualized treatment.⁷¹

Even granting that individualization has its upsides vs. the status quo ante, it certainly does not mean that aggregate patterns of advantage and disadvantage in the marketplace are about to disappear. On the contrary: these new tools threaten both to (a) perpetuate already-existing patterns of harm by less observable means, and (b) generate new groups of harmed individuals who won't recognize what is in fact a shared experience of material disadvantage. While I do not discount the urgency (perhaps even primacy) of the former concern, it is this latter phenomenon — much less discussed in the literature — that I focus on in this Article. The rapid emergence, heterogeneity, opacity, and complexity of the new harmed groups make for a deeply challenging puzzle for all those concerned with distributive justice, marketplace discrimination, and democratic decision-making.

There are several sound reasons for focusing our analysis on shared patterns of harm that play out, at least partly, in the form of advantaged and disadvantaged groups. First, sellers themselves often conceive of their own use of consumer profiling in this way. According to a recent complaint against the data broker Kochava, the company offered its business customers access to tranches of consumers with labels like “cancer,” “bereavement,” “eldercare,” “special needs kids,” and “vaccines.”⁷² This is not an isolated story. Almost a decade ago, for-profit colleges were targeting groups with labels such as “low self-esteem” and “few people in their lives who care about them.”⁷³ These examples demonstrate how these algorithmic tools can generate new groups of individuals

⁷⁰ Lior Strahilevitz, *Information and Exclusion* 141 (2011).

⁷¹ Owen Fiss, *Groups and the Equal Protection Clause*, 5 PHIL. & PUB. AFFS. 107, 126–27 (1976) (noting but taking issue with this conception of anti-discrimination law).

⁷² Kochava Complaint, *supra* note 22, at 20.

⁷³ O'NEILL, *supra* note 1, at 70–73.

harmed on the basis of a perceived shared feature. Such groups often exist in a measurable and verifiable way whether their members have any clue or not. And sometimes — perhaps with appropriate discovery tools, *see infra*, Part III — these groups even come ready-labeled.

Second, as is widely acknowledged in the literature on consumer data harvesting,⁷⁴ the key marketplace actors that compile and sell access to these “groups” of consumers are often big, powerful repeat players. As a result, rather than a vast universe of idiosyncratic patterns of harm that might somehow cancel each other out, we instead face the threat of a few methods of social sorting rising to dominance. If these patterns of apportioning benefits and burdens metastasize throughout the consumer economy and start popping up in seemingly unrelated contexts, the aggregate results could be profound. We already have a precedent for this sort of development. Scholars have written about the “mission creep” that led credit reporting and scoring from a narrow focus on the likelihood of defaulting on a loan to a much wider range of use cases, including housing and employment decisions.⁷⁵ For many landlords and employers, the cheapness and ubiquity of the tool was worth it. Similar stories may already be playing out elsewhere in the market.

Third, given the millions of datapoints that can now inform virtually any consumer-facing decision, an analytic focus on groups rather than individuals can help make the resulting patterns more intelligible (and, relatedly, more suitable as a subject for democratic deliberation and political action). The contrary approach, which accepts the notion that every consumer is an island treated according to some incomparable and bespoke logic, threatens to entrench a world of atomized consumers who have no choice but to accept their lot.

While a framework focused on collective data impacts is relatively new in the literature, a new generation of scholars has pursued a version of this approach. Salomé Viljoen, for

⁷⁴ See, e.g., Solow-Niederman, *supra* note 11; Sari Mazzurco, *Privacy Law's Role in an Information Economy*, 46 CARDOZO L. REV. 123 (2024); Viljoen, *supra* note 11.

⁷⁵ Luke Herrine, *Credit Reporting's Vicious Cycles*, 40 N.Y.U. REV. L. & SOC. CHANGE 305, 321 (2016).

example, has advanced an influential “relational theory of data governance” — discussed at greater length below — that stresses how the collection and use of data is only meaningful at the “population level.”⁷⁶ Neither the knowledge that can be gleaned from a piece of data nor the consequences of its use can be understood only through the lens of an individual “data subject.”⁷⁷ Sari Mazzurco, too, argues that decisions about data use ought to be made on the level of a recognized social group or identity category, analogizing the operative decision-making structures to a sort of union. As discussed below, though, my approach differs from Mazzurco inasmuch as I don’t start from the presumption that our relevant or salient “social roles” are fully knowable without reference to the algorithmic tools themselves.

Sandra Wachter, as I do, explicitly contemplates how profiling tools might create new “algorithmic groups” that fall outside conventionally understood identity categories.⁷⁸ But unlike the argument in this paper, Wachter’s framework (a) views “immutability” as the crucial inquiry for triggering protections in the substantive law and (b) emphasizes the resulting harm as primarily an injury to individual autonomy, rather than to democratic ordering.⁷⁹ More broadly, this Article’s focus on distributive justice norms (in this Part) and litigation as a trigger for democratic deliberation (Part III) mark somewhat of a departure from the core inquiries of these pieces.

The differences in these approaches will be fleshed out more fully in the succeeding sections. But zooming out, this scholarship on collective decision-making about data flows marks an exciting new conversation in the literature to which I hope this Article can contribute.

B. Who Wins, Who Loses, and Why: Distributive Concerns

Algorithmic tools have the potential to radically reshape

⁷⁶ Viljoen, *supra* note 11, at 578.

⁷⁷ *Id.*

⁷⁸ Sandra Wachter, *The Theory of Artificial Immutability: Protecting Algorithmic Groups Under Anti-Discrimination Law*, 97 TUL. L. REV. 149 (2023).

⁷⁹ *Id.* at 157.

both distributive *outcomes* and distributive *logics*. Both moves represent a potential break from the (assumed) status quo, yet one that remains largely invisible to democratic scrutiny and shielded from countermeasures.

Before attempting to outline the potential distributive impacts, it is worth clarifying a few definitions. By “distributive outcomes,” I mean, crudely, who gets what—who wins and who loses, when, by how much, and under what circumstances. By “distributive logics,” I mean the mechanisms and criteria by which the distribution of particular goods or services is determined in a specific context. I then use the term “distributive norm” to capture the notion of an *idealized* normative framework surrounding a particular social good, or the standard against which society judges a distribution to be just or unjust. A real-world “distributive logic” may or may not violate a “distributive norm.”⁸⁰

An example may help flesh out what I mean. Consider a bank with a scarce number of live human assistants on its call line. The bank could apportion these slots in many different ways. Perhaps, to stylize a well-documented example,⁸¹ callers with higher account balances or even wealthier family members might get the right to jump the queue. Alternatively, the bank could shorten wait times for callers with higher social media followings. The bank could also go with a first-come-first-serve approach. In this example, the term “distributive outcomes” would refer to the outputs of this decision-making calculus, *i.e.*, the wait time (or lack thereof) assigned to each caller and, in aggregate, the patterns that such outcomes would reveal if known. The “distributive logic” refers to the decision-making criteria the bank itself uses, such as account balance, social media following, or any other weighting or combination of variables used to sort or rank the callers.

The term “distributive norm,” in turn, refers to citizens’

⁸⁰ Others have used “distributive logics” to refer to society’s *values* surrounding a distribution, *i.e.*, what I call a “distributive norm.” See, e.g., MICHAEL WALZER, SPHERES OF JUSTICE: A DEFENSE OF PLURALISM AND EQUALITY 86 (1984).

⁸¹ See Van Loo, *supra* note 10.

own beliefs⁸² about what constitutes a fair or appropriate distributive logic (or outcome)⁸³ within a given market. In the bank example, the relevant distributive norm would be what society considers a fair or appropriate way to determine wait times. There might be a widely held norm that first-come-first-serve is the only fair approach, or perhaps that deviations from that standard are permissible if advertised clearly in advance (*e.g.*, for an elite membership tier). Or maybe this is the sort of market decision where society is not prepared to rule any particular distributive logic or outcome out of bounds, leaving it instead to the full discretion of the firm.

I do not mean to suggest that distributive norms are either monolithic or static. On the contrary, they are essential subjects for public debate and democratic deliberation. A diversity of views is likely in most cases. What matters for purposes of the present analysis is that in a democratic society, citizens must have the opportunity to (a) learn enough about emergent market practices to assess their compatibility with distributive norms and (b) to discuss and refine the norms themselves in response to changing technologies, social practices, and market conditions. When we don't even know which distributive norms are implicated by a market practice or how, we allow the market actors to order our society in a way that cannot be collectively assessed or challenged. Central to this Article is the claim that where this happens, it is a serious harm that — if left unremedied — is incompatible with democratic society.

1. Context-Specific Norms and Harms

Even well-established distributive norms are not one-size-fits all. No single criterion could ever govern how a society allocates all of its scarce goods and services. Rather, as Michael Walzer has persuasively argued, different goods are governed by different social understandings of a just distribution. “[P]rinciples of justice are themselves pluralistic in form [and] different social goods ought to be distributed for different

⁸² As Viljoen notes, deciding which “public” counts for purposes of assessing uses of data is a major challenge in trying to democratize such decisions. Viljoen, *supra* note 11, at 638.

⁸³ As with anti-discrimination law’s distinction between disparate treatment and impact, it is conceivable that some distributive *outcome* would violate our norms even if the underlying decision-making logic didn’t.

reasons, in accordance with different procedures, by different agents . . .”⁸⁴ To illustrate, Walzer explains how Marx’s criterion of need “doesn’t help at all with regard to the distribution of political power, honor and fame, sailboats, rare books, beautiful objects of every sort.”⁸⁵ All goods, which for Walzer include not just consumer goods and services but also symbolic goods, are “social goods,” in that their value depends on socially and historically contingent understandings of what it means to possess, give, or receive them. Thus, “all distributions are just or unjust relative to the social meanings of the goods at stake.”⁸⁶ The same would be true even if one adopted a narrower conception of goods than Walzer and included only scarce goods or services that are exchanged to some extent in the market. (Distributions of airline seats, broccoli, therapy, and Stradivarius violins likely could not be judged by a single, overarching criterion of justice.)

Accordingly, as we assess the implications of consumer profiling tools, we should remain attentive to the distributive norms that have governed the markets for particular goods and services. Even where the gap between market practices and distributive norms is uncertain or hypothetical, there is a strong case for collective knowledge-building and democratic deliberation. As scholars of anti-discrimination law have long noted, most instances of “small-d” discrimination in the consumer marketplace pit the values of instrumental market rationality against other sorts of values, with law deciding when and where the former must give way to the latter (even if it does not acknowledge doing so).⁸⁷ This process cannot occur unless the nature of such tradeoffs is understood.

Of course, the market *in general* operates by “rules of the game” that differ from those that once predominated in other fields of human endeavor.⁸⁸ A variety of legal and social critics have raised the alarm about the imposition of this “market

⁸⁴ MICHAEL WALZER, SPHERES OF JUSTICE: A DEFENSE OF PLURALISM AND EQUALITY 6 (1984).

⁸⁵ *Id.* at 25.

⁸⁶ *Id.* at 9.

⁸⁷ *See, e.g.*, ROBERT POST ET AL., PREJUDICIAL APPEARANCES: THE LOGIC OF AMERICAN ANTIDISCRIMINATION LAW 40–43 (2001).

⁸⁸ PIERRE BOURDIEU, THE FIELD OF CULTURAL PRODUCTION 183 (1994).

logic” on more and more spheres of human life.⁸⁹ But the important question of whether market transactions of any sort are an appropriate mode of distribution for a given good or service is beyond the scope of this Article.

The key point here is that even *within* consumer markets, where some degree of market exchange is taken as a given, there is no single distributive norm that can answer the questions about fine-grained discrimination in price, service, and terms with which this Article is concerned. Depending on the good or service being purchased, we operate under different assumptions and expectations about how buyer/seller interactions ought to proceed. While status-quo distributive norms differ along all sorts of dimensions, two are of particular relevance here.

First, there is the question of whether formal equality among consumers is widely assumed to be an inherent feature of the market in question. Do I expect to receive the same price in exchange for the same product under the same terms as every other consumer? For many goods bought in a brick-and-mortar retail context (e.g. groceries), the usual answer is probably yes.⁹⁰ The discovery that ten different people paid ten different prices at the checkout counter would likely be received by many as illegitimate. But there are other markets in which consumers do not expect to be treated as equals. Here, the paradigmatic example is the market for consumer credit, discussed below.

In practice, what I have presented as an essentially binary question—do we expect “customization” of prices, products, or terms, or not—may be more like a gradient or spectrum. Airline tickets are an example of an “in-between” market, in that some degree of customization in price (more expensive last-minute tickets) and product (diamond, gold, sapphire status, etc.) is expected, even if a fully individualized calibration might not be. It is in these “in-between” markets where the questions raised by profiling tools may prove especially vexing.

⁸⁹ See generally Jedediah Purdy & David Singh Grewal, *Introduction: Law and Neoliberalism*, 77 L. & CONTEMP. PROBS. 1 (2015).

⁹⁰ Targeted sales or discounts (e.g. rewards programs) are high-salience deviations from this baseline.

The internal distributive norms of specific consumer markets differ in a second important respect: even if consumers do expect that sellers are customizing prices, products, and terms for each individual consumer, *on what basis* are sellers understood to make such adjustments?

A shopper buying an electronics product may accept that his or her purchase history would affect the price on offer. But that does not mean that she would accept her neighborhood affecting the price she paid. A borrower would likely understand that his terms of credit will be customized to his income level and payment history, but probably would not expect his circle of friends to have any impact. At a granular level, this question of the expected or accepted logic of distribution within a particular market becomes an empirical question, and without data about prevailing beliefs among consumers, it can be difficult to establish an authoritative picture of what would or would not offend status-quo beliefs.⁹¹ But the examples below will show how the use of consumer profiling tools could unleash distributive logics within a given market that are at least somewhat in tension with identifiable distributive norms.

2. The Credit Market: Personalization is Expected, But What Kind?

The modern consumer credit-scoring industry already takes as a given that the availability, pricing, and terms of credit are tailored to the risk that a lender associates with the individual in question. Broadly, credit reports use an algorithm to transform data about the consumer's past payments and debts into a single, quantified measure of the consumer's risk of default or another form of "negative financial event."⁹² By allowing lenders to evaluate consumers based on quantifiable, individualized criteria, the advent of credit scoring may in fact

⁹¹ Some data suggests that consumers view a list or reference price as a default to which they are entitled. See Daniel Kahneman et al., *Fairness as a Constraint on Profit Seeking: Entitlements in the Market*, 76 AM. ECON. REV. 728, 732, 734 (1986); see also William Fisher III, *When Should We Permit Differential Pricing of Information*, 55 UCLA L. REV. 1, 11–12 (2007).

⁹² Mikella Hurley & Julius Adebayo, *Credit Scoring in the Era of Big Data*, 18 YALE J.L. & TECH. 148, 153–54 (2016).

have reduced the lending industry's reliance on discriminatory proxies and/or on pre-existing social networks.⁹³ In this way, credit scoring is a specific example of what theorists like Strahilevitz celebrate: better data allowing for the abandonment of flawed and discriminatory proxies. As in Owen Fiss's discussion of labor market discrimination, where our status-quo distributive norms privilege notions of merit or desert,⁹⁴ individualized decision-making may sometimes strike us as not only consistent with, but a prerequisite for, fair treatment.

The decades-old credit reporting system in the United States (a sort of baseline against which we might measure the new profiling tools) is certainly not without flaws. The industry is plagued by widespread complaints about inaccuracy, reliance on overly narrow payment data, and more broadly, the unavoidable fact that using past economic outcomes as evidence of future risk likely perpetuates already-existing patterns of inequality.⁹⁵ Motivated partly by the desire to capture more and better data, and to give marginalized consumers more "credit" for traits or behaviors that apparently correlate with lower risk, a wide variety of alternative credit models has emerged in recent years. Some of these models simply include additional types of payments that the earlier models ignored, such as utility bills.⁹⁶ But others now aim to provide lenders with a more three-dimensional picture of who potential borrowers are, factoring in behaviors, social connections, and personality traits that had previously been invisible to lenders.

One particular variant of "alternative credit scoring" has

⁹³ Robinson & Yu, *Knowing the Score: New Data, Underwriting, and Marketing in the Consumer Credit Marketplace: A Guide for Financial Inclusion Stakeholders*, UPTURN 26–27 (2014), https://www.upturn.org/static/files/Knowing_the_Score_Oct_2014_v1_1.pdf [https://perma.cc/85BG-KRYV].

⁹⁴ Owen Fiss, *Groups and the Equal Protection Clause*, 5 PHIL. & PUB. AFFS. 107, 127 (1976).

⁹⁵ Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 10–16 (2014).

⁹⁶ See, e.g., Sara Burr & Virginia Carlson, *Utility Payments as Alternative Credit Data: A Reality Check*, ASSET BUILDERS OF AM. INC. (2007), https://www.assetbuilders.org/CaseforPAID_final.pdf [https://perma.cc/F6N8-3X4B].

used information culled from an applicant's social media activity. The US company LendUp, for example, rated consumers with more friends and higher volumes of social media interactions as more creditworthy, because "the company views an active social media life as an indicator of stability."⁹⁷ Models such as these have led to fears of discrimination "against applicants who essentially appear socially undesirable."⁹⁸ Similar models purport to assess how responsible or trustworthy someone's friends are.⁹⁹ As discussed above, other lenders have discovered correlations between the language people use in their social media messages and their likelihood to make repayments.¹⁰⁰ In some of these examples, it may be harder to parse precisely who the new winners or losers from such a model are likely to be. But other distributive outcomes might be more predictable. Consumers who use better grammar in their writing have also been found to be more creditworthy, creating potential barriers for otherwise responsible consumers with lower levels of educational attainment or English fluency.¹⁰¹

In all the alternative credit models discussed above, lenders are "discriminating" amongst consumers—distributing credit opportunities unequally—along lines that depart meaningfully from the conventional criteria. And these are just the models we know a bit about in a notoriously opaque industry.¹⁰² Individual traits or behaviors that had previously been thought irrelevant to the credit market will now mark out new

⁹⁷ Erika Eichelberger, *Your Deadbeat Facebook Friends Could Cost You a Loan*, MOTHER JONES (Sept. 18, 2013), <https://www.motherjones.com/politics/2013/09/lenders-vet-borrowers-social-media-facebook/> [<https://perma.cc/3BXQ-TKBR>].

⁹⁸ *Id.*

⁹⁹ *The Surprising Ways that Social Media Can Be Used for Credit Scoring*, KNOWLEDGE AT WHARTON (Nov. 5, 2014), <http://knowledge.wharton.upenn.edu/article/using-social-media-for-credit-scoring/> [<https://perma.cc/LS4A-DMXQ>].

¹⁰⁰ Wei et al., *supra* note 38.

¹⁰¹ Wolfie Christl, *Corporate Surveillance in Everyday Life: How Companies Collect, Combine, Analyze, Trade, and Use Personal Data on Billions*, CRACKED LABS (June 2017), <https://crackedlabs.org/en/corporate-surveillance> [<https://perma.cc/HW5E-QBKX>].

¹⁰² See generally Janine S. Hiller & Lindsay Sain Jones, *Who's Keeping Score? Oversight of Changing Consumer Credit Infrastructure*, 59 AM. BUS. L.J. 1, 118–19 (2021).

individuals and, in the aggregate, new groups, for worse treatment (*e.g.*, introverts).

Taken together, these new models raise the specter of at least two different kinds of potential normative concerns. First, and most obviously, the newly disfavored behaviors or traits will visit real material harms on those who exhibit them, often compounding existing patterns of inequality. This is particularly true for the models that measure the creditworthiness of a consumer's friend group. Second, the departure from existing norms might itself be experienced or understood as an injustice by those subject to these new forms of surveillance, sorting, and evaluation, if they were aware of them. Even those other than the "losers" in such a model might be offended by, and favor limits on, distributive logics that depended on social surveillance or that they found unjust for other reasons.

As I see it, the key normative inquiry is not whether these newly considered variables are "relevant" to the provision of credit in the sense that they have predictive value or are otherwise instrumentally useful for lenders. Rather, the issue is whether the people being sorted or the broader polity would view this new distributive logic as just or unjust, permissible or impermissible (if they knew about it and had a chance to make their views known).

The objection may be raised that social norms are neither static nor exogenous to these technological changes. Indeed, the emergence of a particular market practice may influence the development of social norms about its legitimacy. But the directionality of that relationship is open-ended. If a perceived social ill becomes more widespread, society might respond by doubling down on its existing social norms or by updating those norms. Yet even that process is short-circuited when a democratic society fails to learn that its distributive norms are out of step with the distributive logics at play in the economy.

Zooming out from the credit example, the issue of how and on what basis businesses in a particular market can discriminate amongst consumers—which sorts of groups it can legitimately single out for favor or disfavor—is quintessentially a question for collective, democratic decision-making. Only in this manner can we divine (a) what the status quo norms are,

(b) which innovations or deviations society can tolerate, and (c) which departures are a bridge too far. Admittedly, the complexity and opacity of credit scoring algorithms would likely make such a collective conversation extremely difficult. But this is precisely why forward-looking solutions should focus on ways to make these patterns of favor and disfavor more visible and susceptible to democratic deliberation.

The credit market is simply one example of a broader cluster of consumer markets where individualized pricing and terms based on a consumer's perceived risk are already assumed. But these other, similar markets have developed their own norms about the means of controlling for risk that are just or unjust, irrespective of how rational or actuarially sound they might be. The Affordable Care Act's prohibition on exclusions for preexisting conditions is a recent example of how (a) a collective conversation can identify a market-rational form of differential treatment based on risk that nonetheless violates norms of distributive justice, and (b) democratic decision-making can specify where, when, and how those countervailing values (egalitarianism, discomfort with distributions based on moral luck, etc.) must take precedence.

Notably, there is nothing natural or essential about treating all people with "pre-existing conditions" as a unified group with unified interests and political goals.¹⁰³ Breast cancer, HIV, emphysema, and Marfan syndrome are very different sorts of "pre-existing conditions" that carry different kinds of health risks to the individual, actuarial risks to the insurer, incidences among the population, and perceived dependence on an individual's own choices. By one theory of distributive justice, it might be desirable to sort out which pre-existing conditions were more likely to have resulted from voluntary choices that turned out poorly ("option luck") as opposed to those that were merely random or genetically hard-wired, the result of "brute luck."¹⁰⁴ But that is not the choice the American polity made. Instead, an unjustly disadvantaged group called "people with pre-existing conditions" emerged from our collective

¹⁰³ See 42 U.S.C. § 300gg-3.

¹⁰⁴ See, e.g., Eric Cavallero, *Health, Luck and Moral Fallacies of the Second Best*, 15 J. ETHICS 387 (2011) (critiquing while acknowledging the prevalence of such distinctions).

conversation as one with social recognition and political salience. And democratic deliberation led to the decision to remedy what was understood to be a shared, equivalent injustice for all those who belonged to the group. These are the sorts of social and political processes that can occur when the distributive logics of particular markets are made visible and, as a result, accountable to our context-specific norms of distributive justice. But these conversations cannot occur when both the mechanisms of differential treatment and their effects remain hidden.

3. The Retail Market: Different Distributive Logics and Norms

The market for many retail goods and services is one in which prices and terms are assumed to be more or less standard. Consumers in such markets appear to judge discriminatory pricing as unfair when it is brought to light.¹⁰⁵ Although special discounts for new shoppers or rewards programs certainly exist, these are high-salience strategies for sellers precisely because the status quo in such markets is the expectation that under normal circumstances, the price is the price and the product is the product.

The belief that prices, products, and terms in a given market are fixed and uniform does not make it so. On the contrary, as illustrated in Part I, sellers in the retail market increasingly make use of algorithmic decision-making tools to calibrate the prices, products, and terms offered to particular consumers. As with new credit scoring models, some of these practices have not only generated new groups of winners and losers, but also reshaped the distributive logic of the retail market.

Targeting Willingness to Pay:

Even though many economists argue that individualized price discrimination would increase overall efficiency and social welfare,¹⁰⁶ they also acknowledge that the disappearance of all consumer surplus would represent a welfare transfer from consumers to sellers as a group.¹⁰⁷ As discussed above, it has

¹⁰⁵ Fisher, *supra* note 91, at 11–12.

¹⁰⁶ See, e.g., *id.*; Hal Varian, *Price Discrimination and Social Welfare*, 75 AM. ECON. REV. 870, 870–71 (1985).

¹⁰⁷ Edwards, *supra* note 4, at 566.

long been assumed that first-degree price discrimination would be impossible.¹⁰⁸ Yet predictive analytics tools now allow sellers to get very close. Given the strong incentives for sellers to do so, determining the potential distributive impacts of near-first-degree price discrimination has grown more important than ever before.¹⁰⁹

Sometimes, a consumer may be willing to pay more for a given product simply because of greater spending power overall (higher income, wealth, etc.) Where this is the case—in markets where price sensitivity predictably declines as income rises—sellers have an incentive to charge higher prices to users they estimate to be better-resourced. (This is the assumed logic behind the Apple users being steered to higher-priced items, as discussed *supra*.) The more closely a consumer's willingness to pay for a given thing tracks his overall ability to pay, the more we would expect price discrimination to have a downwardly redistributive effect on net.

But crucially, it is *not* the case that the willingness to pay for a particular product will always predictably scale upward with financial resources. First, there are instances in which individuals exhibit low price sensitivity because of localized market failures that constrain their set of available choices. Here, even if the market as a whole is competitive enough to put downward pressure on prices, the market available to a particular consumer is narrower or less competitive. This creates an opportunity that some companies have learned to exploit. In a well-known example, “Staples’s algorithm for online prices looks at the zip code to determine whether the consumer has other options nearby, which essentially enables the company to better pinpoint potential market failures relating to insufficient choice.”¹¹⁰ (Other online retailers have

¹⁰⁸ See, e.g., Miller, *supra* note 5; Klock, *supra* note 4, at 327–28.

¹⁰⁹ See, e.g., Andrew Odlyzko, *Privacy, Economics, and Price Discrimination on the Internet*, in PROCEEDINGS OF THE 5TH INTERNATIONAL CONFERENCE ON ELECTRONIC COMMERCE 355, 357–58 (2003); Klock, *supra* note 4, at 328.

¹¹⁰ Van Loo, *supra* note 5 (citing Jennifer Valentino-Devries et al., *Websites Vary Prices, Deals Based on Users’ Information*, WALL ST. J. (Dec. 24, 2012), <https://www.wsj.com/articles/SB10001424127887323777204578189391813881534> [<https://perma.cc/P9W2-TFK8>]).

used similar zip-code-based strategies.)¹¹¹ While the geographic and socioeconomic incidence of market failures probably depends on the product in question, such tools threaten to raise prices precisely on communities that already face the effects of non-competitive local markets. This worry may be especially acute in minority communities.¹¹² Conversely, those with a wealth of options in a given market may now receive what are essentially targeted discounts.

Consumers may also value a product more, and therefore be willing to pay more, in certain exigent circumstances. This is another situation in which willingness to pay would not necessarily correlate positively with financial resources.¹¹³ And it is also the sort of situation that is becoming increasingly visible to sellers, who now have the tools to track or predict major life events, from pregnancies to the purchase of a new pet.¹¹⁴ Algorithmic tools can also help sellers detect when someone's life has taken a turn for the worse, due to a recent divorce or death in the family. In all these situations, sellers have an incentive to raise prices at precisely the moment a consumer has a pressing need for a particular good or service.

Certainly, price gouging in general is nothing new, and the law and economics literature is replete with defenses (and critiques) of such practices. I will not attempt here to resolve that debate. Rather, I want to call attention what *is* relatively new: the ability of sellers of all sorts to calibrate prices from moment to moment to match the ups and downs of our lives at extraordinary levels of detail. (To capture something of the character of this change, Uber knows that users will pay more

¹¹¹ Neil Howe, *A Special Price Just for You*, Forbes (Nov. 17, 2017), <https://www.forbes.com/sites/neilhowe/2017/11/17/a-special-price-just-for-you/#> [<https://perma.cc/3VGT-6SM3>].

¹¹² Executive Office of the President, *Big Data and Differential Pricing* 17 (2015), https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/docs/Big_Data_Report_Nonembargo_v2.pdf [<https://perma.cc/66V6-Q8LV>].

¹¹³ Of course, the probability of exigent circumstances would not always predictably correlate with lower incomes, either.

¹¹⁴ See, e.g., Michal S. Gal & Niva Elkin-Koren, *Algorithmic Consumers*, 30 HARV. J.L. & TECH. 309, 317 (2017).

for a ride when their phone batteries are lower.)¹¹⁵ In essence, we must now grapple with the potential for a new distributive logic, one in which people in moments of particular exigency can be singled out and forced to pay more for what they need most.

Admittedly, there are normative considerations tugging on both sides of the issue of need-based pricing. Moreover, the distribution of material outcomes may turn out to be something of a wash over the long term, since most people aren't in exigent circumstances all the time. But a world in which new dog owners pay more for rug cleaner or new parents pay more for earplugs is very different than the one most consumers have grown accustomed to. Such practices might very well offend distributive norms (more in some markets than in others) sufficiently to merit regulatory attention.

Perhaps more important, there is one sense in which "exigency-based" pricing could generate new disadvantaged groups over the long term and/or result in more regressive distributive outcomes on net. For some people, a heightened need for a particular product is not the effect of a transitory life moment, but a more permanent feature of their lives. The example that comes most readily to mind is the physically disabled, though categories of "heightened need" might flow from other pre-existing disadvantages. (Consumers might, for instance, have a heightened need for in-home educational products depending on the quality of education or after-school programs in their areas.) Here, the importance of making visible the newly disadvantaged groups—and for society to assess whether it is just or unjust for sellers to treat them differently—is especially acute.

To the extent those with permanent, structural exigencies cannot afford the higher prices at all, predictive analytics tools might themselves reveal such constraints on their ability to pay and lower their prices accordingly. But for a single business providing the good or service for which someone is known to have inelastic demand—essentially a necessity for this one

¹¹⁵ Amit Chowdhry, *Uber: Users Are More Likely to Pay Surge Pricing if Their Phone Battery Is Low*, FORBES (May 25, 2016), <https://www.forbes.com/sites/amitchowdhry/2016/05/25/uber-low-battery/> [https://perma.cc/9V35-9LKT].

person but not for another—there is a strong incentive to ratchet the price as high as the person is able and willing to go, *i.e.*, just before it starts to impinge on other non-negotiable consumption necessities.

Are we comfortable with a world in which consumers who have special material needs can be selectively pushed to forgo other, less-essential consumption through higher prices for those high-priority goods? We might be, we might not be, or we might arrive at some in-between answer, where there is a ceiling on how much ratcheting-up we are willing to tolerate. But these are difficult normative questions that require sustained collective conversation and democratic deliberation. In order to have a fighting chance at articulating moral claims, and convincing the unaffected among us that they are suffering an impermissible harm, those who are targeted on the basis of idiosyncratic material needs must know—both about themselves and those in a similar position.

The discussion above has focused on scenarios in which willingness to pay might correlate upward with the incidence of particular life circumstances (other than wealth or income). Increasingly, sellers in the consumer marketplace have also been able to discover and track specific *behaviors* that correlate with a higher or lower price sensitivity. In one sense, much like price gouging, behavioral pricing at a broad level is nothing new. Retailers have long known, for example, that consumers are willing to tolerate higher prices when they end in “9” rather than “0,” or that “anchoring” a consumer’s expectations with an artificially high base price can make a putative discount more attractive.¹¹⁶

But the breadth of the data now available about the aggregate mass of consumers—combined with new data-mining techniques—allows for much more fine-grained assessments of an individual consumer’s price sensitivity. In one sense, this might not raise serious concerns if it could be understood as rewarding or encouraging bargain hunting. Yet retailers now invest significant sums in obfuscating the true price of their products (and impeding apples-to-apples

¹¹⁶ Van Loo, *supra* note 5, at 1328.

comparisons) through a wide variety of tactics.¹¹⁷ Some retailers, for instance, create multiple, functionally identical versions of the same product, all with different prices.¹¹⁸ The ability of sellers to sow confusion in this way has grown more powerful due to the ability to run “randomized controlled trials that easily provide statistical significance across thousands of stores and millions of transactions.”¹¹⁹ And just as they develop these new weapons in their arsenal, sellers are also developing the ability to divine *ex ante* which consumers have the weakest defenses. Some analysts have written about this possibility as the sorting of the sophisticated from the unsophisticated,¹²⁰ while others are concerned with the ability to predict who will fall prey to specific kinds of cognitive biases or mistakes.¹²¹

However we conceive of this filtering mechanism, the offering of higher prices to those less able to notice them will certainly generate new groups of winners and losers. Two factors suggest that consumers already facing resource constraints are likely to be especially vulnerable to these tactics. First, conditions of economic scarcity can themselves dampen the quality of one’s financial decision-making due to the cognitive effect of repeated financial stresses.¹²² Second, any tools aimed at combatting such practices would likely require an increasing investment of intellectual and material resources in order to keep up. The very consumers who might be most sensitive to price differentials if they were cognizant of them might be most ill-positioned to detect or avoid any particular small-dollar upcharge.¹²³

More broadly, all sorts of behaviors or personality traits

¹¹⁷ *Id.* at 1336; *see also* Oren Bar-Gill, *The Behavioral Economics of Consumer Contracts*, 92 MINN. L. REV. 749, 769–76 (2008).

¹¹⁸ Van Loo, *supra* note 5, at 1337–38.

¹¹⁹ *Id.* at 1322.

¹²⁰ Miller, *supra* note 5, at 60.

¹²¹ Van Loo, *supra* note 5, at 1316 (referring to these as market failures).

¹²² *Id.* at 1359 n.219 (citing Anandi Mani et al., *Poverty Impedes Cognitive Function*, 341 SCI. 976 (2013)); *see generally* Cara Feinberg, *The Science of Scarcity*, HARV. MAG. (May–June 2015), <https://www.harvardmagazine.com/social-sciences/the-science-of-scarcity> [<https://perma.cc/8QVU-A46B>].

¹²³ The curve might turn out to have an upside-down U-shaped, in which it is middle-income consumers who have the combination of the resources (more than the low-income) and the income incentives (more than the high-income) to invest in such tools.

might correlate with price sensitivity for a given product or service, including the times of day consumers are awake and transacting¹²⁴ or their openness to forming new relationships with acquaintances.¹²⁵ Many of these targeting techniques might strike us as unproblematic if we knew about them. But perhaps others would bother us, particularly if the harmed groups were able to make a forceful claim that they should not be targeted for differential treatment on such a basis.

Targeting “High-Value” Consumers:

The preceding discussion has taken as its assumption that, all things being equal, sellers will attempt to maximize the surplus they can extract from consumers by charging as close as possible to the maximum price each consumer is willing to pay for a given transaction. This idea has strong support in economic theory and is the primary mechanism scholars of price discrimination have in mind when they discuss personalized pricing.¹²⁶ But sometimes, companies take a longer view beyond the transaction at hand, instead focusing on which sorts of customers they want to retain over the longer term. There is evidence that companies will use differential pricing—or differential service provision—to selectively retain “high-value customers.”

But who are the “high-value customers” that sellers are interested in retaining? This question turns out to be a complex one. Sometimes, the answer is nearly literal: stores want customers who can and will spend more money overall. In the past, sellers would do this by pushing goods or services that their preferred customers wanted. Supermarkets, for instance, would “sometimes offer special discounts on higher-end luxury

¹²⁴ Patrick Witschi et al., *Understanding the Global Price Sensitive Consumer*, BOS. CONSULTING GRP. (2021), <https://www.bcg.com/publications/2021/consumer-price-sensitivity> [https://perma.cc/EQG3-QRL9] (finding that price sensitivity for coffee varied by time of day).

¹²⁵ One study suggested “friendship-oriented behaviors” by staff to increase gym members’ “psychological switching costs” in an attempt to dilute price sensitivity. Lei Ouyang, *An Empirical Study of the Determinants of Consumer Price Sensitivity for the Health and Fitness Club Industry* (Sept. 2020) (Ph.D. dissertation, University of Northern Colorado), <https://digscholarship.unco.edu/cgi/viewcontent.cgi?article=1708&context=dissertations> [https://perma.cc/YG3Y-YM39].

¹²⁶ See, e.g., Edwards, *supra* note 4; Klock, *supra* note 4.

products in order to garner the loyalty of high-value customers, while at the same time raising prices on staple items that are preferred by low-value customers,” a practice that “has been criticized as a cross-subsidy of the wealthy by those of lesser means—essentially ‘paying the wealthy for being wealthy.’”¹²⁷

Now, though, sellers have the tools to identify well-resourced shoppers individually and to target them appropriately for special benefits—often without them knowing why, and without the non-beneficiaries knowing it is happening at all. Bloomingdale’s, for instance, offered special point-of-sale discounts for customers its algorithms deemed to be actual or likely big spenders, as well as “special services that increase the value of that person’s purchases compared to other customers.”¹²⁸ Bank of America has developed technology that takes into account not only the wealth of particular customers, but also their family members, in adjudicating disputes.¹²⁹

Other times, however, the most valuable customer may *not* be the most well-resourced. Instead, sellers may be looking specifically to attract customers who are relatively price insensitive (their “strong market,” according to one formulation) as opposed to customers who are relatively price sensitive (their “weak market”).¹³⁰ And customers with an especially low price sensitivity may be reflecting market failures specific to the product in question, personal exigencies of special interest to the seller, or cognitive biases the seller has uncovered, rather than merely greater resources. Crucially, the seller is targeting these customers precisely so that it can extract greater surplus from them in the future. Customers tagged as actual or likely bargain-hunters—*i.e.*, more price-sensitive—may be singled out for *disfavor*. As discussed above, some retailers use poor service to repel customers they’d prefer not to have.

Businesses may assess customers as “high value” or “low value” for any number of additional reasons. Sometimes these criteria might be idiosyncratic, such as a nightclub seeking out

¹²⁷ Miller, *supra* note 5.

¹²⁸ *Id.* at 94–95.

¹²⁹ Van Loo, *supra* note 10, at 565–66.

¹³⁰ Miller, *supra* note 5, at 58–59.

more physically attractive clientele. But patterns could emerge in which certain consumers consistently get labeled as “high value” or “low value” across many of their market transactions. (People with larger or wealthier social networks, for instance, may prove to be more valuable customers in multiple lines of business.) These consumers will enjoy what Akiva Miller aptly calls “price advantages,” in that their dollars will buy them more than the dollars of other consumers.¹³¹

From the perspective of distributive justice, there may be reasons to be especially concerned specifically about the differential provision of customer service to perceived high-value customers (as opposed to differential sticker prices or targeted discounts). First, as Joseph Turow and others have pointed out, the quality of customer service attached to a given purchase affects its value.¹³² Charging two people the same price for products of different value is economically equivalent to charging different prices for products of the same value: both are forms of price discrimination.¹³³ As difficult as apples-to-apple price comparisons can be, it will almost always be *more* difficult—often impossible—to predict what kind of personalized downstream treatment you will receive. From wait times to the likelihood of a return being accepted, downstream customer-service differentials are something most consumers cannot even putatively bargain over.¹³⁴ Given that commentators have long been concerned about low-salience contract terms that extract costs from the vast majority of consumers who do not notice them,¹³⁵ “no-salience” contract terms should cause even more concern.

A second reason for heightened scrutiny is that such techniques may represent a particularly stark departure from

¹³¹ *Id.* at 91–92. Miller argues that such advantages should be viewed as a commodity.

¹³² Turow, *supra* note 7.

¹³³ Klock, *supra* note 4, at 370.

¹³⁴ *But see* Daniel Markovits, *Good Faith Bargaining in the Shadow of a Form*, in DEFENSES IN CONTRACT (Andrew Dyson, James Goudkamp & Frederick Wilmot-Smith eds., 2017) (conceptualizing the handling of customer complaints as a site of displaced bargaining over low-salience contract terms).

¹³⁵ *See, e.g.*, Russell Korobkin, *Bounded Rationality, Standard Form Contracts, and Unconscionability*, 70 U. CHI. L. REV. 1205 (2003); Van Loo, *supra* note 5, at 1318.

status-quo norms. Particularly in situations where consumers need some kind of dispute adjudicated, such as returning a product or challenging a charge, research has shown they “value having an unbiased and factual decision-making process in which the rules are applied in a consistent manner.”¹³⁶ Where a new distributive logic violates a pre-existing norm of procedural justice in such a manner, it may merit special attention.

The strategies outlined above are likely to generate particular groups of winners and losers in different marketplace contexts. Whatever the net effect on social welfare—and notwithstanding the widespread view that price discrimination is likely to be welfare-increasing overall — the issue of who gets the increased welfare at whose expense, and under what circumstances, implicates a wide variety of distributive justice concerns that merit sustained democratic deliberation.

C. Finding Invisible Allies

It is hard enough to determine if any one person is paying higher prices than others or otherwise getting the short end of the stick. In the unlikely event a consumer can acquire such information, going a step further and figuring out *why* her individualized outcome looks a certain way is often impossible. In fact, the difficulty runs deeper. Lacking even this basic knowledge about what is happening to each one of us as individuals, we are deprived of the chance to communicate with each other about the rapidly changing marketplaces we inhabit. In this Subpart, I examine in more detail the potential consequences of this deliberative deficit.

I begin this Subpart by framing my theoretical approach to understanding the structural features and functional consequences of democratic deliberation. With this framework in mind, I then explain why the new harmed groups previewed above have a crucial role to play in helping society decide how to confront the new AI tools in the consumer marketplace. I argue that a core challenge is ensuring that citizens can find

¹³⁶ Rebecca Hollander-Blumoff & Tom R. Tyler, *Procedural Justice in Negotiation: Procedural Fairness, Outcome Acceptance, and Integrative Potential*, 33 L. & SOC. INQUIRY 473, 492 (2008).

their “invisible allies,” *i.e.*, those who have been harmed in similar ways, have a basis for believing that they might be so harmed in the future, or otherwise have some special reason for understanding or empathizing with each other’s predicaments. To make this concept more concrete, I sketch out how some such coalitions might form. This discussion forms the basis for my thinking in the following Part, where I consider how litigation might catalyze these latent but unrealized connections.

1. Collective Problem-Solving and the Role of Groups

My analysis focuses on three modes of collective problem-solving without which it is difficult to establish effective and legitimate democratic governance. First, there is the basic issue of “detection,” by which I mean the process of apprehending that some real-world phenomenon (material, cultural, environmental, etc.) is afoot in society. At this stage, the primary task is for citizens to be alerted *that* there is some change to be grasped at all. Then, there is the more complex process of “diagnosis,” the admixture of storytelling, analysis, and debate by which groups of citizens develop some shared understanding of the phenomenon. When a critical mass of people arrives at the conclusion that this new phenomenon needs to be acted upon — halted, curtailed, reversed, or even accelerated — they can join together in a third stage, “mobilization,” in which they use tools of formal and informal power to try to accomplish that common goal.

All these processes, as I conceive of them, have a few key features that are worth highlighting. First, all of them are inherently communicative. Unlike the “aggregative model” of collective decision-making, which envisions politics as a system for collecting pre-existing citizen preferences and filtering them through the policymaking process, I adopt a more deliberative conception of democracy.¹³⁷ In line with those who endorse this deliberative model,¹³⁸ I take the position that the act of deliberation itself helps give rise to the viewpoints that emerge.

¹³⁷ IRIS MARION YOUNG, INCLUSION AND DEMOCRACY 17–21 (2002) (describing and critiquing the aggregative model).

¹³⁸ *Id.* at 21–26; *see also* JÜRGEN HABERMAS, THE THEORY OF COMMUNICATIVE ACTION (1981).

Put differently, through citizens' communicative encounters with each other, their understandings of empirical reality, their norms, and their political priorities take shape.

In contrast, the aggregative model of democracy is especially ill suited to tackling emerging social issues of uncertain dimension. To use the vocabulary developed at the outset of this Subpart, it is unclear what role exists for "detection" or "diagnosis" if one's conception of democracy treats citizens' views and preferences as a given, as if they could have come into being independent of their interactions with others. Even if the aggregative model might make sense for figuring out how a democratic polity should work through longstanding social issues on which citizens already have fairly hardened positions (and I doubt it does even then), it cannot provide a model for how citizens can come to an understanding they currently lack. If, instead, we aspire for democracy to function not only as a vehicle for gathering and acting upon citizen preferences but also for *forming* them, the deliberative model is a much better fit.

None of this is to suggest that people are free from material interests or ideological priors that influence what they bring to (or take from) these encounters with their fellow citizens. The key point is that such priors are neither static nor exogenously given. People's broader conceptions of their own material interests and beliefs may bear the imprint of the conversations in which they have participated. The "priors" they carry into next month's debates may not be the same as the "priors" that influenced them this month.

And yet, there can be no question that citizens' perceptions of their own material situations, group memberships, affinities, and alliances can have a profound impact on the trajectory of democratic deliberation. Just as individual citizens do not spring into democratic life with their knowledge and opinions fully formed and frozen, it would be equally implausible to expect the undifferentiated mass of all citizens to learn from or influence each other all together, all the time, to an identical extent. A useful framework for understanding democratic deliberation in a pluralistic society must leave room for the material, cultural, affective, and epistemological bonds that situate groups of citizens in different kinds of relationships with

each other. At any given time, a citizen is more likely to know, hear from, or care about some people than others.

So far, this account of group membership as a structuring, but not deterministic, influence on democratic deliberation is intended to be mostly descriptive rather than normative. But I also align myself with theorists who make a strong normative case for acknowledging and harnessing social difference in service of democracy. To simplify a longstanding divide among theorists of democracy, some (*e.g.*, Jürgen Habermas) are more inclined to see social difference as a barrier to overcome in service of reaching consensus. For Habermas and his allies, the public giving of reasons and testing of arguments is the route by which society is to reach this goal. Others, most notably Chantal Mouffe, have advanced a more “agonistic” model of democracy, which conceives of conflict as an inescapable and even desirable feature of democratic life. In this framework, a quest for consensus may at times be anti-democratic, especially inasmuch as it ignores or flattens actually existing conflicts amongst groups.¹³⁹ Mouffe and other proponents of an agonistic approach argue that democratic life is best understood as an arena for managing and domesticating those conflicts, rather than eliminating them through reasoned debate.

Drawing from the work of political theorist Iris Marion Young, I favor a pluralistic framework that holds out hope for the potential of problem-solving communication (as in Habermas) while *also* insisting (as in Mouffe) on the importance of making room for differently situated groups to develop and articulate their points of view. Young describes this “third way” as a “process of public discussion and decision-making which includes and affirms the particular social group positions relevant to issues . . . in order to draw on the situated knowledge of the people located in different group positions as resources for enlarging the understanding of everyone and moving them beyond their own parochial interests.”¹⁴⁰

Young’s conception of the proper role of differentially situated groups in democratic dialogue includes both an

¹³⁹ Chantal Mouffe, *Deliberative Democracy or Agonistic Pluralism?*, 66 SOC. RSCH. 745 (1999).

¹⁴⁰ YOUNG, *supra* note 137, at 109.

outward-facing and an inward-facing component. For members of a given group (especially a disadvantaged one), engaging with outsiders is an opportunity to try to win support for more favorable social and political conditions over which the broader society has influence. But as the above quote illustrates, it is crucial to Young's framework that *everyone else*, *i.e.*, the polity as a whole, also derives benefits from the engagement of groups with particular knowledge or perspectives that might otherwise be missing. Conversely, when democratic processes exclude or obscure such voices, we end up with democratic outcomes that are both less legitimate and less responsive to the facts on the ground.

Young also recognizes that robust *intra-group* communication is a precondition for the sorts of broader political engagement envisioned above. Similarly situated people may only be able to apprise themselves of their shared situations when they confront social reality together. As Young puts it, an "account of someone's life circumstances contains many strands of difficulty or difference from others that, taken one by one, can appear to be the result of decision, preferences, or accidents. When considered together, however, and compared with life story of others, they reveal a net of restricting and reinforcing relationships."¹⁴¹ Relatedly, this discovery of commonality can itself give rise to the evolution of politically consequential groups that previously did not exist as such, "encourag[ing] solidarity among those with a group affinity, and a sense of political agency in making justice claims to the wide society."¹⁴²

For democracy to function in this way, therefore, there must be a means for citizens to identify those who are "similarly positioned," who share "similar constraints or enablement, particular modes of expression or affinity," such that these bonds can form the basis for collective knowledge-building, argumentation, and action.¹⁴³ When these bonds instead remain hidden, those stories may never get told and

¹⁴¹ *Id.* at 93.

¹⁴² *Id.* at 103.

¹⁴³ *Id.* at 100. Young's theory has something in common with Bourdieu's theory of the social field, in which social and political action is conditioned by a given agent's "habitus" (modes of perception and expression) and "capital" (including both economic and social capital).

those justice claims may never get made.

2. Top-Down and Bottom-Up Stimuli for Solidarity

There is no single process by which disadvantaged groups develop the collective consciousness necessary to mobilize politically. As sociologists Francesca Polletta and James Jasper explain in their review of the literature, scholars of social movements have noted a variety of different political contexts and social practices that have caused “structural inequality [to get] translated into subjective discontent.”¹⁴⁴ Sometimes, the state itself can shape and give content to the rise of politically salient solidarities through its own legal categories or discriminatory actions. In this way, “the legal institutionalization of racial privilege in the United States and South Africa generated severe and pervasive inequality and eventually provided the basis for demands by blacks for legal equality. By contrast, the absence of legalized racial categories, agencies, and statistics in Brazil impeded black mobilization.”¹⁴⁵

The state can also affirmatively recognize and thereby lend its backing to certain conceptions of shared harm. As Owen Fiss has noted, even though federal antidiscrimination law doctrine proclaims itself as a vindication of the right of the individual to proceed unfettered in the marketplace, its functions can be better explained as a collective identification of groups of people who suffer persistent disadvantage, and whose needs must sometimes take precedence over market norms.¹⁴⁶ Different polities would be expected to arrive at different conclusions about who such groups are. (California’s Unruh Civil Rights Act, for example, recognizes far more protected classes than do the equivalent federal public

¹⁴⁴ Francesca Polletta & James M. Jasper, *Collective Identity and Social Movements*, 27 ANN. REV. SOCIO. 283, 287 (2001) (quoting Verta Taylor & Nancy E. Whittier, *Collective Identity in Social Movement Communities: Lesbian Feminist Mobilization*, in FRONTIERS IN SOCIAL MOVEMENT THEORY 104, 104 (Aldon D. Morris & Carol McClurg Mueller eds., 1992)).

¹⁴⁵ Polletta & Jasper, *supra* note 144, at 283, 287 (citing ANTHONY MARX, MAKING RACE AND NATION: A COMPARISON OF SOUTH AFRICA, THE UNITED STATES, AND BRAZIL (1998)).

¹⁴⁶ Owen Fiss, *Groups and the Equal Protection Clause*, 5 PHIL. & PUB. AFFS. 107, 126–27 (1976).

accommodations laws.)¹⁴⁷

Of course, long before these conversations rise to the level of legislative debates, they happen on the level of friends and neighbors, congregants and colleagues. People tell stories about where they can and cannot get served, who refuses to hire them, and why. In this way, conceptions of shared harm can also arise in a bottom-up fashion out of “prior bonds with others that make solidaristic behavior a reasonable expectation” and “[motivate] participation through norms of obligation and reciprocity.”¹⁴⁸ Although not always necessary, a helpful ingredient to transform such social bonds into political mobilization networks is some kind of institutional “free space in which people can develop counterhegemonic ideas and oppositional identities,”¹⁴⁹ such as Black churches or gay bars. Organizers can also use narrative and framing devices to “make a compelling case for the ‘injustice’ of the condition and the likely effectiveness of collective ‘agency’ in changing that condition.”¹⁵⁰

All these routes involve individuals establishing (1) mutual recognition of a characteristic they have in common and (2) a shared understanding of the place they occupy in the social and political context on the basis of that characteristic. To the extent that the new mechanisms of algorithmic profiling get in the way of either of these two essential ingredients, they impede the ability of similarly situated individuals to engage in the processes of detection, diagnosis, or mobilization.

To illustrate how this might work in practice, even if three people with debilitating back problems all know each other to suffer from that condition—or could at least find each other via social media—that would merely satisfy some degree of criterion (1), mutual recognition. But what do they recognize each other as? In contrast, imagine the counterfactual in which, prior to meeting, each of the three received a message from Uber saying, “having determined based on your web activity

¹⁴⁷ CAL. CIV. CODE § 51.

¹⁴⁸ Polletta & Jasper, *supra* note 144, at 290.

¹⁴⁹ While I do not necessarily think “identity” is the best framework for conceptualizing the new harmed groups, *see infra* Part II.C.4, scholars of collective identity formation have much to contribute to our understanding of how mutual recognition comes about.

¹⁵⁰ Polletta & Jasper, *supra* note 144, at 291.

that you are incapable of walking to the subway and are not in possession of a car, we will be charging you 25% more per ride.”

In essence, the establishment of a new basis for political solidarity represents a shift, or an expansion, in our sense of those commonalities or differences among people that *matter*. The social theorist Charles Taylor sees the act of social and political self-definition as the discovery of the distinctions that carry social weight: “When we come to understand what it is to define ourselves, to determine in what our originality consists, we see that we have to take as background some sense of what is significant. Defining myself means finding what is significant in my difference from others. I may be the only person with exactly 3,732 hairs on my head, or be exactly the same height as some tree on the Siberian plain, but so what?”¹⁵¹ The anonymity, opacity, and hyper-personalized nature of algorithmic discrimination work to mask the true significance of a variety of social categories and thereby rob them of their political salience. And when they take shape, new fault lines of potential political significance become harder to uncover. For that reason, members of newly disfavored groups may barely recognize themselves as such.

Before turning to some illustrative examples that will help flesh out who these new groups might be, I want to be clear about what I am *not* arguing. In laying out my theoretical approach, I have been careful to avoid speaking of new social or political “identities,” instead favoring terms like “affinities,” “commonalities,” or even just “groups.”¹⁵² I avoid the identity framework for a few reasons. First, like Young herself, I am concerned that casting every social or political grouping as an “identity” threatens to yield to “essentializing tendencies” that “freeze[] the experienced fluidity of social relations by setting up rigid inside-outside relations among groups.”¹⁵³ Wendy Brown, similarly, points out that identity categories have a tendency to ossify over time, inhibiting the more organic,

¹⁵¹ CHARLES TAYLOR, *ETHICS OF AUTHENTICITY* 36 (2018).

¹⁵² Some of the theorists I rely on are more comfortable with the identity frame, however. *See, e.g., id.* at 100–10.

¹⁵³ YOUNG, *supra* note 137, at 87–88.

bottom-up bonds of solidarity that might otherwise form.¹⁵⁴ Especially in an era of new and shifting patterns of harm, it seems unwise for advocates of collective action to take the position that every shared experience of harm ought to give rise to a new identity that the bearer will carry everywhere.

A distinct problem with an identity framework — of special relevance here — is that it overlooks the importance of changes over time. While there are exceptions, most identity categories are understood to persist throughout a person's life. Even identity categories that don't quite fit this mold, like age or professional affiliation, still have a predictable permanence over some span of time. In contrast, the categories that matter for consumer profilers are often of an inherently ephemeral character, be they emotional states, exhibited behaviors, or particular life events (*e.g.*, a recent death in the family).¹⁵⁵ Other times, a personal characteristic might be assessed as relevant by a consumer profiling tool in some corners of the market and not in others, meaning the trait in question might not rise to the level of an overarching "identity." And yet it still might matter some of the time, enough to cause annoyance, empathy, and perhaps action.

As discussed below, the possibilities for solidarity and collective action also grow as people are able to see themselves or their loved ones as *potential* or *future* sufferers of an analogous harm, or to recall relevant circumstances from their own pasts. In this way, I borrow from the literature on disability rights, which has long used terms such as "not yet disabled" to express the notion that everyone will one day be disabled to at least some extent.¹⁵⁶ Many of the new harmed groups involve similarly unstable categories for which the chance of eventual membership extends to much larger numbers of people than might count themselves as present-day members. Given the right conditions, such a broad conception of potential allies could turn a seemingly narrow interest group into something much more powerful.

¹⁵⁴ WENDY BROWN, *STATES OF INJURY: POWER AND FREEDOM IN LATE MODERNITY* 55 (1995).

¹⁵⁵ Kochava Complaint, *supra* note 22, at 20 (alleging that "bereavement" is a profiling category used by the defendant company).

¹⁵⁶ Elizabeth Emens, *Mindful Debiasing: Meditation as a Tool to Address Disability Discrimination*, 53 CONN. L. REV. 835, 864 n.145 (2022).

3. Who's in This With Me?

So who might comprise these coalitions of latent, but heretofore invisible, allies? I focus here on three illustrative examples drawn from a combination of documented and hypothetical algorithmic profiling strategies. Even the hypothetical strategies, however, are grounded in what is known about the technological capabilities of automated profiling tools and the business incentives that drive them.

Cognitive Differences and Mental States:

Businesses have the ability to measure and exploit a wide range of cognitive shortfalls, from inattentiveness to impulsivity.¹⁵⁷ Relatedly, businesses can and do measure the emotional states of their customers in real time. Reporting has revealed that Facebook, for instance, uses analyses of photos and other posts to assess when users are feeling “stressed, defeated, overwhelmed, anxious, nervous, stupid, silly, useless, and a failure.”¹⁵⁸ Other digital profiling tools have focused on categories such as “low self-esteem” and “bereavement.”¹⁵⁹ Even some brick-and-mortar businesses use facial recognition technology¹⁶⁰ and other data to make similar assessments.

Little attention has been paid to how such profiling might generate organized counter-movements. In a way, this makes sense: people may not be fully aware of their own cognitive biases or vulnerable emotional states, and even if they were, they might not wish to build relationships on that basis (perhaps because they aspire to a different situation). Add to that the problem of impermanence — just because I am grieving today doesn't make me a “griever”— and the difficulties grow further.

Yet there are reasons to be optimistic that bonds of solidarity could form on such a basis. Perhaps most important, while several of these categories or profiles might not look like a social group with any kind of self-awareness as such, many of

¹⁵⁷ See *supra* pp. 542–43.

¹⁵⁸ Sam Levin, *Facebook Told Advertisers It Can Identify Teens Feeling 'Insecure' and 'Worthless,'* THE GUARDIAN (May 1, 2017), <https://www.theguardian.com/technology/2017/may/01/facebook-advertising-data-insecure-teens> [<https://perma.cc/63JV-F2AP>].

¹⁵⁹ See *supra* pp. 542–43.

¹⁶⁰ Pisani, *supra* note 35.

them contain or overlap with groups that do. (And the existence of these groups, in turn, helps ameliorate a bit the concern that people would shy away from establishing relationships on the basis of cognitive or mental traits.) In recent years, for instance, reporting has shown that young people diagnosed with ADHD and autism spectrum disorder have been increasingly vocal and visible on social media platforms such as TikTok, making friends and building community through their shared stories.¹⁶¹ While this community-building doesn't always (or doesn't yet) have an explicitly political cast or aim for specific policy goals, it can certainly take that form. In Brazil, for instance, autism awareness is arguably more ingrained in public life, with a legally recognized symbol for autism and other "invisible disabilities" recognized at a federal level.¹⁶² Interestingly, self-perceived social categories in this domain have already begun to stretch across formal diagnostic categories. Along with "invisible disabilities," concepts like "neurodivergence" and "neurodiversity" have emerged in recent years as widespread umbrella terms that capture a range of minority cognitive traits.¹⁶³

At first blush, the above examples might not immediately cry out for new forms of mutual recognition or collective action. Some of these groups may, for instance, already fall under the rubric of the Americans with Disabilities Act (ADA). But importantly, sellers charging higher prices or targeting exploitative offers based on some cognitive or

¹⁶¹ See, e.g., Ellen Simpson et al., "Hey, Can You Add Captions?": *The Critical Infrastructuring Practices of Neurodiverse People on TikTok*, PROCS. ACM ON HUM.-COMPUT. INTERACTION, Apr. 2023, at 1.

¹⁶² *Brazil Passes Law About the Sunflower*, Hidden Disabilities Sunflower, <https://hdsunflower.com/us/insights/post/brazil-passes-law> [<https://perma.cc/TJN3-AZXH>] (last visited July 6, 2026); see also Brazilian Chamber of Deputies Bulletin, *Lei Prevê uso de Fita com Desenho de Girassóis para Identificar Pessoa com Deficiência Não Aparente*, CÂMARA DOS DEPUTADOS (July 18, 2023), <https://www.camara.leg.br/noticias/980154-lei-preve-uso-de-fita-com-desenho-de-girassois-para-identificar-pessoa-com-deficiencia-nao-aparente/> [<https://perma.cc/SD6A-RBT8>].

¹⁶³ Caitlin H. Cochran & Jessica Edonick, *The Socialization of Neurodivergent Graduate and Professional Students at Big State University 18* (Aug. 2022) (Ph.D. capstone, Vanderbilt University) (noting the "emerging community" of neurodiverse students).

psychological profile likely don't care if the consumers in question match up with the diagnostic criteria, let alone whether they have been formally diagnosed. The operative question is, instead, whether the consumers in question are willing to pay the higher price, likely to fall for the expiring free trial or "drip pricing" trap, or otherwise transact in some desired and predictable manner that sets them apart from others. Our established categories of disability — even the new, more bottom-up categories driven by activism — almost certainly fail to capture a wide swath of affected consumers in such situations.

There is of course a difference between (a) people who haven't been formally diagnosed, say, as depressive, and so don't self-conceive in this way, despite nonetheless matching the formal criteria for someone with depression, and (b) those who do not in fact suffer from depression, but nonetheless exhibit relevant traits or behaviors at a given moment ("sadness" or "hopelessness") that look the same from the perspective of a seller. Those in group (b), though, might still be upset to learn they were profiled in such a way, even for a fleeting moment or day or week, especially if such profiling was tied to measurably worse treatment. The resulting outcry, among both affected individuals and their friends or relatives, might be even more profound if their transient emotional states resulted from difficult life moments such as the loss of a family member (*e.g.*, the "bereavement" category above). If it were made known, even from a single company's actions, the social category "targeted based on sadness" could give rise to collective conversation, storytelling, and eventually justice claims that spread beyond the transactions in question. This is the sort of democratic deliberation that isn't currently happening, leaving atomized individuals at the mercy of marketplace actors that stand to benefit from exploiting their emotional states.

Similar logic applies to consumer targeting based on cognitive traits such as inattentiveness or impulsivity. For some people, these measurable qualities may flow from an established disability such as ADHD. For others, however, the situation might be quite different. Perhaps the behaviors they exhibit result from a more isolated personality trait that falls

short of a diagnosable disorder. Even more people might display tell-tale signs of inattentiveness or impulsivity at some times or not others, or during particular stretches. Granted, certain sellers might care about these distinctions only inasmuch as they can make global predictions about a given consumer. But other sellers (potentially quite a few) may be more interested in whether someone is impulsive or inattentive *today*, or *recently*. Moreover, people don't necessarily exhibit such behaviors randomly. Rather, conditions of acute financial stress, such as the periods right before payday for those living paycheck-to-paycheck, cause people to perform worse on measures of cognitive function.¹⁶⁴ Thus, for a seller interested in exploiting distracted or impulsive consumers, it is likely that some consumers would regularly and predictably show up as prime targets more than others. Here, again, affected individuals might be inclined to recognize each other as similarly situated victims.

The full list of analogous potential groups is essentially limitless and, on the basis of current information, unknowable. But if we are to update our distributive norms to account for the new winners and losers in the consumer marketplace and adjust our legal and policy responses accordingly, we must find ways to alert affected populations and to give them an opportunity to communicate amongst each other (and then with the rest of us).

Social Media Abstainers:

Businesses may have all sorts of reasons for taking an interest in the quality and quantity of their customers' social media activities. Even one person's social media activity could yield a dizzying range of measurable data, so for purposes of this discussion, let us consider the stylized question of whether a person is "more" or "less" active on social media. Why might this matter to a business, and how might this affect the decisions they render as to a particular consumer?

Some businesses are especially concerned with those who are likely to complain about poor customer service. While it is sometimes possible to identify serial complainers in advance,

¹⁶⁴ Anandi Mani et al., *Scarcity and Cognitive Function Around Payday: A Conceptual and Empirical Analysis*, 5 J. ASS'N CONSUMER RSCH. 365, 365–76 (2020).

businesses might also use a customer's social media reach as a coarser gauge of the risks associated with dissatisfaction. As scholars of consumer protection have noted, businesses have an incentive to respond by defanging the potential complainers in advance, giving them better-than-average service and leaving everyone else without the benefit of what might have been helpful feedback.¹⁶⁵ Interestingly, a business more interested in affirmatively spreading the word about *positive* outcomes would often be expected to favor more or less the same strategy: showering benefits on those with social media reach. In both cases, consumers with smaller networks or fewer followers will likely get the short end of the stick. As discussed above, there are other, non-PR-related reasons (e.g., “stability” or “trustworthiness”) that businesses might treat consumers differentially based on their perceived social connections.

As with the examples of “sad” or “distracted” consumers, those targeted for worse service or other adverse outcomes based on their absence from social media (or their small number of visible connections) do not currently constitute a social group whose members conceive of themselves as such. Whether or not they would recognize each other as similar situated victims of a shared harm, even if all the patterns of differential treatment were visible, is unknowable. But under present circumstances, the affected individuals don't even have a chance to contemplate this option.

Lest this potential harmed group seem too trite or low-stakes to spark any sort of meaningful collective action, it is worth considering who might make up its membership. Once again, we would expect some predictable overlap with groups whose existence and special vulnerability we *do* already recognize, such as the elderly or the very poor (on the wrong side of the digital divide). Those who have been or fear becoming targets of harassment, such as victims of domestic violence, may have special reason to stay off social media. And then there is the great mass of people who simply choose to

¹⁶⁵ Yonathan A. Arbel & Roy Shapira, *Theory of the Nudnik: The Future of Consumer Activism and What We Can Do to Stop It*, 73 VAND. L. REV. 929, 968 (2020). Other strategies exist: some businesses might also try to avoid transacting with such consumers at all. *Id.*

stay away from social media based on personal preference, such as a general inclination toward privacy. We as a society might decide — or the affected individuals at least might try to persuade us — that such choices should not be punishable by higher prices or longer wait times. Furthermore, these various “subgroups” have little reason to view each other as fellow travelers in any sort of social or political endeavor. Together, their claims might take a form (and pack a punch) that they would not be able to if each subgroup proceeded separately.

Dependencies on Particular Products or Services (Dating Apps):

The two potential coalitions mapped out above involved modes of consumer profiling that could follow consumers across a variety of consumer transactions. But there are other kinds of profiles that might correspond to dependencies on a given product or service. Such dependencies have the potential to reveal consumers to be a “strong market,” a relatively price-insensitive segment of the customer base that might be willing to pay more than average.¹⁶⁶ Of course, not every such price-sensitivity determination would make a natural rallying cry for a new set of justice claims. In some markets, though, affected groups might have some serious arguments to make.

Putting aside more obvious examples, such as insulin or other products or services essential to physical well-being, consider an example that may at first seem trivial: the cost of a dating app. This hypothetical is inspired by the real-life case of *Candelore v. Tinder*, in which a plaintiff challenged Tinder’s practice of charging users over 30 a higher price for using the dating app’s premium service.¹⁶⁷ A common interpretation of this dual-pricing regime attributed the difference to the notion that singles start to grow “more desperate” as they age.¹⁶⁸ But Tinder had an entirely different explanation: the lower price

¹⁶⁶ See *supra* Part II.B.

¹⁶⁷ *Candelore v. Tinder, Inc.*, 228 Cal. Rptr. 3d 336, 340 (Cal. Ct. App. 2018).

¹⁶⁸ See, e.g., Alison Griswold, *California Told Tinder to Stop Discriminating Against 30-Somethings*, QUARTZ (Jan. 30, 2018), <https://qz.com/1193515/california-told-tinder-to-stop-discriminating-against-30-somethings> [<https://perma.cc/U2PQ-4M3Z>]; Amanda Schupak, *Over 30? Tinder Plus Will Cost You More*, CBS NEWS (Mar. 3, 2015), <https://www.cbsnews.com/news/over-30-tinder-plus-will-cost-you-more/> [<https://perma.cc/2SEK-CF8H>].

for younger users reflected their lower “relative income and willingness to pay for the service, as a group, as compared to the older cohort.”¹⁶⁹

For reasons that have to do with California’s unique public accommodations law, the Unruh Act, the California Court of Appeals held that Tinder could not use the protected category of age as a proxy at all, regardless of the “true” motive.¹⁷⁰ In contrast, the Court held that Tinder *could* permissibly charge different prices based on an “economic distinction,” specifically citing willingness-to-pay as a legitimate reason for doing so. The Court’s distinction between “economic distinctions” and “personal characteristics” is logically unsatisfying for reasons that require a more extended discussion than this Article permits. But putting that puzzle aside, we might still ask: who is it, ultimately, that might be willing to pay more for a dating app?

For a business like Tinder, all sorts of qualities that might be highly correlated with a customer’s willingness to pay, such as the length of time a person has been single, loneliness, introversion, or how far they live from nightlife venues. The list could go on. I mention these only to show that one *could* imagine a story — including one ascribed to Tinder’s actual conduct, the targeting of desperate singles — in which willingness-to-pay no longer looks so uncontroversial.

4. How We Conceive of the New Harmed Groups

Having reviewed some examples of who the new winners and losers might be, it would be reasonable to ask: should we care? By what criterion should we separate out the groups worth protecting, or the forms of small-d “discrimination” we view as out of bounds? These are thorny, theoretically challenging questions that resist easy answers. I do not attempt here to make a single overarching argument from first principles about where society ought to draw these difficult lines. Rather, I take the position that the very complexity and open-endedness of these questions (and the heterogeneity of the contexts in which we’re now asking them) *requires* sustained democratic deliberation before any legitimate

¹⁶⁹ *Candelore*, 228 CAL. RPTR. 3d at 336.

¹⁷⁰ *Id.* at 340.

framework could be even tentatively developed. To use the language of Chantal Mouffe, when a pluralistic democracy is confronted with “traces of power and exclusion” that potentially challenge the foundational ideals of the polity, the core task is to “make them visible so that they can enter the terrain of contestation.”¹⁷¹ To state my position a bit more plainly: not all of the new harmed groups would succeed in persuading their fellow citizens that they deserve protection from marketplace exploitation. But they deserve the chance to try.

To show why such contestation would be necessary, it is helpful to consider just a small sample of the competing conceptions of unjust or invidious discrimination that already command a following within U.S. law and policy. My aim here is to explain by way of illustration how these conceptions are not only at odds with each other, but can also lead to strikingly different answers when applied to the vexing questions of algorithmic consumer profiling raised in this Article.

The debates within U.S. equal protection doctrine offer one view of how varied and unsettled our notions of improper discrimination are. There is, for example, the famous notion in “Footnote Four” of *United States v. Caroline Products* that “discrete and insular minorities” may need special protection from the law to account for their inability to advocate for themselves through the usual political channels.¹⁷² But scholars such as Bruce Ackerman have pointed out that it can actually be “anonymous and diffuse” groups, rather than the “discrete and insular,” that face special challenges in winning support for their favored political positions.¹⁷³ Just as these two frameworks would lead to different conclusions as to whether to institute protections for the poor, or for sexual minorities (groups likely to be “anonymous and diffuse” but *not* “discrete and insular”) they would also counsel very different approaches to dealing with the (likely anonymous and diffuse)

¹⁷¹ CHANTAL MOUFFE, *THE DEMOCRATIC PARADOX* 33–34 (2009).

¹⁷² *United States v. Carolene Prods. Co.*, 304 U.S. 144, 152 n.4 (1938).

¹⁷³ See, e.g., Bruce Ackerman, *Beyond Carolene Products*, 98 HARV. L. REV. 713 (1985). For an overview of the literature assessing the origins of political powerlessness, see generally Nicholas O. Stephanopoulos, *Political Powerlessness*, 90 N.Y.U. L. Rev. 1527, 1532 (2015) (citing Ackerman among others as participants in this long-running debate).

groups formed by algorithmic profiling tools.

Indeed, as Nicholas Stephanopoulos has noted, being “discrete and insular” or “anonymous and diffuse” are only two among many reasons a group might experience the “political powerlessness” that U.S. law has recognized as a justification for special protection.¹⁷⁴ To make matters even more complicated, “political powerlessness” is far from the only justification recognized by Equal Protection doctrine for carving out a suspect class. U.S. courts have recognized other “indicia of suspect status,” including “(1) whether a group has experienced a history of discrimination; (2) whether a group is defined by an immutable characteristic; and (3) whether a group is defined by a trait that typically bears no relation to the group’s ability to contribute to society,” without making it at all clear how those conceptions of at-risk groups relate to one another.¹⁷⁵ And those criteria, in turn, represent only those used within a single, fairly circumscribed area of U.S. law: equal protection doctrine.¹⁷⁶

Each of these criteria, if adopted as the go-to approach for regulating algorithmic consumer profiling, would lead us in very different directions. And, relatedly, each of these criteria would also fail to address some of the potentially troubling practices chronicled in this Article. Sandra Wachter, for example, argues that it is indeed the “immutability” of certain algorithmically created groups that should make them objects of particular concern.¹⁷⁷ But as discussed above, some people’s “sadness” or “distractibility” might be more immutable than others’. Should the targeting of transiently-depressed

¹⁷⁴ Stephanopoulos, *supra* note 173, at 1558–59 (noting how different conceptions of political powerlessness would lead to different doctrinal answers in an Equal Protection analysis).

¹⁷⁵ *Id.* at 1563.

¹⁷⁶ The above discussion has also elided the more macro-level debate — ongoing and unsettled — about whether equal protection doctrine ought to be more concerned with dubious classifications as such or rather with persistent patterns of subordination. *See, e.g.*, Laura Portuondo, *The Myth of Equal Protection*, 16 U.C. Irvine L. Rev. (forthcoming 2026) (manuscript at 13–14) (explaining the “competing” conceptions of anti-discrimination law).

¹⁷⁷ *See generally* Wachter, *supra* note 78 (arguing that immutability is a helpful framework for understanding those algorithmic groupings that rob individuals of their autonomy and agency).

consumers be unconstrained due to their failure to satisfy the “immutability” criterion? Other theorists of discrimination, such as Kasper Lippert-Rasmussen, take the position that “social salience” is the appropriate criterion for separating groups that merit anti-discrimination protection from those that don’t.¹⁷⁸ (These arguments roughly track the “history of discrimination” criterion in equal protection law.) On his view, isolated or “idiosyncratic” instances of differential treatment fail to even qualify as discrimination if they do not flow from some more widespread, already-existing pattern of animus or disadvantage.¹⁷⁹ This view, however, would leave untouched the algorithmic tools that single people out for better or worse treatment based on *newly* relevant categories or traits.

I do not mean to suggest by this discussion that an anti-discrimination principle is only useful or normatively defensible if it can capture the full range of consumer profiling practices that might bother us. No such principle exists. Rather, I aim to demonstrate that there are tensions and tradeoffs between the different approaches that would make any sort of universal or *a priori* solution impossible. All of these anti-discrimination principles are perhaps better viewed as resources upon which we can draw as we decide where the line between permissible and impermissible marketplace practices ought to be. But without a better sense of empirical reality, we can’t even begin to understand who we’d be leaving unprotected — the stakes, essentially — of adopting any one criterion. It is only through sustained democratic deliberation that we can work through the shared consequences of rejecting (or endorsing) any particular mechanism of consumer profiling. It is not only likely, but in my view normatively desirable, that such deliberation would result in very different answers for different sorts of consumer transactions.

III. What Litigation (and Procedure) Can Do

Even the most basic facts about AI-driven profiling remain

¹⁷⁸ Karl Lippert-Rasmussen, *Discrimination: An Intriguing but Underexplored Issue in Ethics and Political Philosophy*, 2 MORAL PHIL. & POL. 207, 207–17 (2015), <https://doi.org/10.1515/mopp-2015-0026> [<https://perma.cc/839Q-E9EJ>].

¹⁷⁹ *Id.*

shrouded in mystery, protected by a combination of corporate secrecy, complex technology, and the atomized experiences of individual consumers. The previous section explored how bonds of mutual recognition among consumers could, at least in theory, help level the playing field. While such relationships could develop along a number of dimensions, one of the most promising ways forward would be assembling groups of similarly harmed individuals, the coalitions of (currently) invisible allies.

How can these groups of invisible allies be made visible, either to each other or to the broader universe of legal and political actors? While it had been largely absent from the policy conversation, the importance of collective action has recently emerged as a point of agreement among scholars concerned with the prospect of forcing individuals to navigate the new algorithmic marketplace alone.¹⁸⁰ These scholars have done essential work by making the case that we need to develop a more democratic, less individualistic approach to governing the use of our data. But the harder question is *how*. On this front, for obvious reasons, the ideas are in their infancy.

Viljoen, for instance, does not go into much detail as to the concrete reforms that might make “data as a democratic medium” a reality. But her work appears to envision a situation in which state mandates might force patterns of data use out of their balkanized corporate niches and into some kind of collective institution(s) subject to democratic control. This approach has much to recommend it, especially inasmuch as it insists on creating publicly accountable entities that can assess “population-level effects” of mass surveillance.¹⁸¹ Yet in the admirable move away from an individualistic paradigm, Viljoen’s framework threatens to swing instead toward a totalizing model in which all parties having an “interest in” some form of data collection or use get a right to participate in the relevant decision-making.¹⁸² Leaving the diffuse mass of affected parties to somehow manage all their complex and conflicting data interests through an as-yet-undetermined process of collective governance carries, in my view, a risk of

¹⁸⁰ See, e.g., Viljoen, *supra* note 11; Mazzurco, *supra* note 74.

¹⁸¹ Viljoen, *supra* note 11, at 583.

¹⁸² *Id.* at 578.

over-centralization or even elite capture.¹⁸³ Relatedly, Viljoen's model does not reckon explicitly with the theoretical or practical implications of pluralism. In a large and diverse country like the United States, locating the nexus of decision-making at *too* high a level or in too centralized an institution would have some important potential downsides. What we might gain in terms of a more complete, bird's-eye-view of the costs and benefits of data use, we could lose in terms of the smaller, divergent, variegated points of view that often struggle for attention in majoritarian systems.

One scholar's approach to democratic control over algorithmic decision-making does attempt to take on these issues more directly. Embracing an explicitly pluralistic vision, Sari Mazzurco argues for union-like "collective bargaining associations" of Internet users, which could negotiate together over the collection and use of their members' data.¹⁸⁴ Mazzurco's framework envisions that such cooperative associations could form on the basis of not only shared employment, but also membership in many different kinds of "social roles," including "mother," "tenant," or "parishioner." This model nicely captures both the importance of collective action as a counterweight to corporate power over mass data collection and use, and the intuitive truth that we have different views about our data within our separate social spheres.

Of course, there is nothing mutually exclusive about the two visions sketched out by Viljoen and Mazzurco. One could envision a regime in which a central data-governance institution or equivalent citizens' committee sat on top of, and helped to police, a range of decentralized data collectives (perhaps like the National Labor Relations Board functions vis-a-vis unions or workplaces). Still, this sort of future sounds

¹⁸³ In fairness to Viljoen, she explicitly lists as an open question "how to identify the relevant 'public' or institutional level of civic life at which to coalesce and govern such collective interests." Viljoen, *supra* note 11, at 638. But given Viljoen's arguments about the interconnectedness of all of our data decisions and the heterogeneity of the competing interests at play, her analysis suggests to me that the new institutions she envisions would have to be big enough to encompass a staggeringly wide range of social and political actors. Perhaps this is a good thing. My point here is merely that making these decisions at such a scale comes at a cost, both from the perspective of pluralism and bottom-up democratic participation.

¹⁸⁴ Mazzurco, *supra* note 74.

a long way off. Both in Mazzurco's model and in Viljoen's, there is relatively little discussion about how to get from point A to point B, *i.e.*, from our status quo of near-total opacity and atomization to a future where citizens participate meaningfully in deciding how their data can be used in the marketplace.

In this Part, I explain why litigation, and especially aggregate litigation, could be a crucial (if counterintuitive) bridge to that future. The present era is not the first in which we have had to devise ways to tackle widely dispersed harms that are difficult both to detect and deter. Drawing from civil procedure scholarship, I argue that, with the right tools in place, aggregate litigation has the potential to assist us with all the key functions of collective political communication identified above: detection, diagnosis, and mobilization. The upshot is that there are sound reasons to see litigation as a force multiplier for democratic deliberation — indeed, as a *part* of it — rather than an alternative to political action. This function of litigation is especially important in the context of rapid social and technological changes like those discussed above.

But just because litigation can perform this democracy-enhancing role doesn't mean it will. In the following Subpart, I will take a closer look at some of the key procedural tools, rules, and doctrines, as they actually exist today. The picture that emerges from the following discussion, ranging from class actions and multidistrict litigation to standing doctrine and discovery, is of a system that may not be a good fit for the new marketplace realities unleashed by ubiquitous AI. This overview aims to lay the groundwork for a targeted procedural reform agenda that would empower private plaintiffs to play a stepped-up role in algorithmic governance.

The news is not all bad, however. Within each area of procedure covered in the below discussion, even current law reveals some possible paths forward. Where I see these green shoots, I explain how litigators might be able to use them to force more of the new algorithmic harms into the public arena. We need not wait on big procedural reforms (let alone a wholesale revolution in data governance) to start bringing more accountability to AI-driven consumer markets. Private plaintiffs, with the help of strategic litigators and thoughtful judges, can help kickstart that process now.

A. Aggregate Litigation as Part of the Democratic Ecosystem

It is first worth clarifying why an argument about democratic oversight should spend so much time on litigation. After all, courts are an unelected branch of government, access to which is controlled by a maze of impenetrable rules and a guild of expensive gatekeepers. Appropriately, scholars have criticized court-led, top-down theories of change that have downplayed the importance of bottom-up political action.¹⁸⁵

But if the goal is to engender meaningful public deliberation about marketplace AI, the moment calls for a “both/and” approach. We should be looking for ways to harness the potential of courts to work *in service of* broader democratic deliberation, political action, and ultimately new legislation and regulation. As other scholars have pointed out in a variety of contexts, litigation in general and aggregate litigation in particular can indeed play this role in American political life.¹⁸⁶ Developing an account of the potentially symbiotic relationship between private lawsuits and public deliberation can help orient our more granular discussions about the strengths and weaknesses of particular procedural tools.

So what is it about litigation that could make it not only a proper, but perhaps indispensable, part of our much-needed democratic reckoning with the new algorithmic business practices?

In his classic work *Adversarial Legalism: The American Way of Law*, Robert Kagan catalogues a number of features that make America’s legal system unique when viewed from a comparative perspective. His normative assessment is deeply ambivalent, replete with competing lists of pros (e.g., the “tradition of entrepreneurial legal advocacy” and the “creativity and boldness of its . . . judiciary”) and cons (e.g., “costliness,” and “legal uncertainty”). But on one metric relevant here, America’s system really stands out. Kagan

¹⁸⁵ See, e.g., Amna Akbar, *Non-Reformist Reforms and Struggles over Life, Death, and Democracy*, 132 YALE L.J. 2497 (2023).

¹⁸⁶ See, e.g., Elizabeth Chamblee Burch & Abbe R. Gluck, *Plaintiffs’ Process: Civil Procedure, MDL, and a Day in Court*, 42 REV. LITIG. 225, 231 (2023).

identifies the U.S. as having “the most politically and socially responsive court system in the world.”¹⁸⁷ As Kagan describes it, “the judiciary and lawyers [are] more fully part of the governing process and more fully democratic in character” than in other systems.¹⁸⁸ What he means by this is that the U.S. provides routes for litigants, lawyers, and judges to take a leading role in tackling big public problems—lawsuit by lawsuit, ruling by ruling—that would be handled more comprehensively by expert bureaucrats and/or elaborately articulated rules in other systems.¹⁸⁹

Kagan does not see it as an unalloyed positive that private litigation and political decision-making are so closely intertwined in the U.S. While our system enjoys an uncommon degree of openness to new and innovative advocates, he also critiques its unpredictability, inefficiency, and vulnerability to capture and misuse, among other ills.¹⁹⁰ Like Kagan, I do not attempt here to render any overarching verdict on whether the upsides are worth the downsides. Instead, I invoke his work because it makes an important analytical point about the structural relationship between litigation and democracy in the U.S.

For better or worse, identifying problems on the ground, filling gaps in the substantive law where possible, and spurring broader advocacy where social reality has outstripped our existing legal frameworks are all *functions that our political system assumes litigation will perform*. As long as this basic structural fact remains true, even the more classically democratic modes of action — organizing, protesting, lobbying, voting, legislating — will be impoverished if litigation isn’t functioning well. (To put the argument a bit differently: the various undesirable effects of the politics-litigation nexus

¹⁸⁷ ROBERT A. KAGAN, ADVERSARIAL LEGALISM: THE AMERICAN WAY OF LAW 16 (2001).

¹⁸⁸ *Id.*

¹⁸⁹ Because of his comparative approach, Kagan’s sketch of our system is of course an ideal type meant to throw the unique features of adversarial legalism into sharp relief. Nothing in his (or my) analysis discounts the fact that the U.S. has a robust regulatory state, a hefty body of laws, and plenty of non-court avenues for advocacy and policymaking.

¹⁹⁰ KAGAN, *supra* note 187.

identified by Kagan and other scholars¹⁹¹ would more clearly swamp the positive features if we allow those positives, especially robust bottom-up participation, to be neutered.) We have rigged up a system that counts on private lawsuits to uncover hidden information from secretive private actors, to educate us about new harms in need of redress, and to test out the boundaries of existing legal protections and prohibitions. The narratives and outcomes of specific lawsuits also help everyday citizens better understand the stakes of emerging legal and political debates.

As this analysis suggests, some sorts of social and political issues, more than others, may be especially dependent on litigation to help improve the odds that they get a serious public airing. Key factors include the opacity of the conduct in question; the difficulty of pursuing collective action when the harms are widely dispersed; the added difficulty when those dispersed harms don't fall neatly along established group lines; and the unsettled nature of shared norms and narratives.¹⁹² All of these facets of algorithmic consumer profiling make it particularly important that litigation is living up to its democratic potential in this context.

To make the matter even more pressing, we also face a serious lack of clarity about whether our existing substantive laws even cover the various AI-driven strategies discussed above. For reasons discussed above, many of the new categories of algorithmic discrimination diverge from or cut across our established protected classes, meaning the traditional tools of anti-discrimination law (including public accommodations laws) may not apply in a straightforward manner. Scholars have also pointed out that many of the businesses engaged in consumer profiling would likely fall outside the purview of statutes like the FCRA.¹⁹³ Instead, many of the most promising substantive provisions may be state unfair and deceptive acts and practices (UDAP) laws, whose

¹⁹¹ Stephen B. Burbank et al., *Private Enforcement*, 17 LEWIS & CLARK L. REV. 637, 661–79 (2013).

¹⁹² Wachter, *supra* note 78 (identifying a similar set of features that mark off algorithmic groups as “immutable”). While I do not adopt the immutability frame, Wachter and I are concerned with similar real-world dynamics at play.

¹⁹³ Helveston, *supra* note 68, at 684–89.

capacious and open-ended language relies especially heavily on the iterative process of litigation and judicial interpretation to take on meaning through common law-style reasoning.¹⁹⁴ Even where the resulting holdings make clear that the forms of algorithmic discrimination discussed in this Article do *not* qualify as an unfair (or anticompetitive) practice, such an outcome would at least clarify for advocates and policymakers which gaps in existing legislation need to be filled in. But if our existing procedural mechanisms make it difficult to test the limits of the law, leaving these open-ended statutes to sit inert on the shelves, the path forward for reform of the substantive law becomes muddled.

Litigation's key role in the U.S. democratic ecosystem is only possible because of some unique procedural features. First, the very fact that so many substantive legal regimes provide for private enforcement at all sets the U.S. apart. These private enforcement mechanisms can "take advantage of private information to detect violations," rather than relying entirely on the investigative capacity (or commitment) of the state.¹⁹⁵ Private enforcement regimes also "encourage legal and policy innovation," "facilitate participatory and democratic governance," and make the threat of deterrence more credible, among other virtues.¹⁹⁶ Essentially, private enforcement provides citizens with more points of access to the state apparatus and to public discourse.

Those salutary social and political effects of private enforcement can only come to pass, however, if private citizens have the information, resources, and incentives to actually bring a case. In many instances, this simply isn't feasible. That

¹⁹⁴ See, e.g., Charles W. Tyler, *Common Law Statutes*, 99 NOTRE DAME L. REV. 669 (2023). While broadly worded bans on unfair and deceptive practices would probably be the most natural fit in most U.S. jurisdictions, a small handful of states has recently begun to adopt explicit prohibitions on certain forms of individualized pricing based on personal data. See, e.g., Conn.: S.B. 4, 2026 Gen. Assemb., Reg. Sess. (Conn. 2026) (enacted); Md.: H.B. 895, 2026 Leg., Reg. Sess. (Md. 2026) (enacted). But both of these statutes rely exclusively on enforcement by the attorney general, eliminating the potential for private litigation to do any work in uncovering new patterns of harm or in testing how far the new proscriptions actually sweep.

¹⁹⁵ Burbank et al., *supra* note 191, at 662.

¹⁹⁶ *Id.*

is why a second procedural feature of the U.S. system is so essential: individual citizen enforcers can often band together. These tools of aggregate litigation — including both representative suits like Rule 23 class actions and pretrial consolidation procedures like multidistrict litigation (MDL) — supercharge what private enforcement can accomplish. This unique combination of private enforcement and aggregation is particularly important for the sorts of opaque and highly dispersed harms highlighted in this Article.

Not every private enforcement regime necessarily depends on aggregation to accomplish its goals. To take one example, many employment discrimination cases arise from the high-salience event of a single employee's termination. Even before discovery, the employee in question often has at least a reservoir of information about the workplace and some resulting suspicions about what happened to her. That is often a good enough start to file a good faith complaint. The damages associated may also be enough to incentivize a lawsuit, particularly for repeat-player attorneys who can therefore streamline their research processes and other startup costs. Once the lawsuit is up and running, it can also spark broader conversations about workplace norms, encourage more careful compliance (especially if there is a finding of liability), or point out the shortcomings of existing law.

Now consider a very different kind of harm: a consumer charged \$10 more for the exact same laptop charging cable, on the same day, as someone else across town buying from the same website.¹⁹⁷ For all sorts of reasons, the consumer would probably never find this out. Even if she did, she would likely have no idea how this came about in her specific case. Even if the consumer overcame these hurdles, the stakes don't make it worthwhile for her to even look for a lawyer, nor for a lawyer to pursue the case.

Without aggregation, that would probably be the end of the story. The infeasibility of low-damages individual suits may

¹⁹⁷ This example is not at all far-fetched based on how online retailers have been operating for over a decade. Jennifer Valentino-Devries et al., *Websites Vary Prices, Deals Based on Users' Information*, WALL ST. J. (Dec. 24, 2012), <https://www.wsj.com/articles/SB10001424127887323777204578189391813881534> [<https://perma.cc/FA2D-X5TY>].

seem an obvious point, but it is worth pausing to consider the consequences of letting such harms fly under the radar. Whether one is concerned with remedying the direct distributive effects of the harm, the broader need to deter and detect future such conduct, or deliberation about the nature of the harm itself, aggregation would likely prove crucial.

First, there is the question of compensation (and the distributive implications if it never materializes). Whether or not the consumer or society would care about that one \$10 overcharge, the seller's marketplace likely extends beyond that transaction. If the affected individual is a repeat customer, that seller's pricing algorithm could have generated hundreds or perhaps thousands of dollars in excess charges over the years.¹⁹⁸ Even if she *isn't* a repeat customer, a \$10 overcharge multiplied across thousands or millions of transactions with other buyers could represent a meaningful shift of would-be consumer surplus to that seller. In either case, the harm starts to look big enough to be worth compensating, or at least big enough that it is worth assessing whether compensation is appropriate. Aggregation makes this possible.

As Judge Posner has written, “[t]he smaller the stakes to each victim of unlawful conduct, the greater the economies of class action treatment and the likelier that the class members will receive some money rather than (without a class action) probably nothing, given the difficulty of interesting a lawyer in handling a suit”¹⁹⁹ Aggregate litigation thus has the potential to operate as a sort of “redistributive tool by society's have-nots,”²⁰⁰ helping uncover and undo what might otherwise be large-scale material exploitation by powerful and well-capitalized marketplace actors. There are, of course, trenchant critiques in the literature about the agency issues and other problems that threaten the efficacy of class actions as a

¹⁹⁸ A consumer might be “overcharged” relative to other buyers for some products and “undercharged” for others, resulting in a negligible or unclear net effect as to her. But even if the net effect might be a wash for *some* consumers, others might find themselves on the winning or losing end again and again. We just don't know how often either scenario occurs.

¹⁹⁹ *Hughes v. Kore of Ind. Enters., Inc.*, 731 F.3d 672, 675 (7th Cir. 2013).

²⁰⁰ Maya Steinitz, *Whose Claim Is This Anyway? Third-Party Litigation Funding*, 95 MINN. L. REV. 1268, 1274 (2011). While Steinitz makes this point specifically about what third-party-funded litigation can accomplish, it can also apply to aggregate litigation more broadly. *Id.*

compensation device.²⁰¹ But in many circumstances, the alternative to aggregate litigation may well be no enforcement at all.

In addition to compensation, perhaps an even more important goal of the private enforcement system is the hoped-for deterrent effect.²⁰² Here, too, aggregation is key. Where private wrongdoers stand to garner huge sums by imposing small harms on large numbers of people, the incentive to move forward with such a scheme becomes even greater when the amounts in question *wouldn't* be enough for any one consumer to seek redress. Allowing the affected individuals to bring suit together is one of the key ways the U.S. litigation system aims to rebalance the risk calculus, making what would otherwise be a tempting business strategy a bit less attractive. In this way, aggregate litigation has the potential to operate as a bottom-up engine of market oversight and an adjunct to state-led investigations and enforcement actions.

The stakes of ensuring a meaningful deterrence effect could end up being particularly important in the context of algorithmic consumer profiling. In the higher-priced charging cable example, the seller's business strategy probably doesn't exist in isolation. Rather, the overcharge often flows from a specific third-party algorithm, software, or pricing strategy that is used more widely by other sellers. Even a lawsuit against just that one retailer might therefore have rapid follow-on effects throughout the market. Conversely, the failure to hold any one seller accountable for (potentially) unlawful algorithmic decision-making strategies may impose social costs far beyond that seller's own transactions.

Finally, as discussed above, litigation can also play a broader role in the U.S. political system beyond providing compensation for, or deterring, some specified conduct. Through downstream effects on grassroots activism, electoral politics, and policymaking, private lawsuits help orient the democratic agenda. As Alexandra Lahav has written,

²⁰¹ See, e.g., John C. Coffee, Jr., *Class Action Accountability: Reconciling Exit, Voice, and Loyalty in Representative Litigation*, 100 COLUM. L. REV. 370 (2000).

²⁰² *Kore*, 731 F.3d at 681; Myriam Gilles & Gary B. Friedman, *Exploding the Class Action Agency Costs Myth: The Social Utility of Entrepreneurial Lawyers*, 155 U. PA. L. REV. 103 (2006).

aggregate litigation can itself be an “associational” activity that stimulates the sorts of intra-group communication contemplated in the previous Part.²⁰³ Elizabeth Burch, too, has written extensively about the potential for groups of plaintiffs to develop a sense of community through the litigation process.²⁰⁴ Both these scholars, among others,²⁰⁵ also emphasize how the storytelling and information-sharing functions of litigation can help develop a narrative of shared harm that didn’t previously exist, precisely the sort of deliberative activity contemplated in the previous Subpart.

But there are reasons to worry about whether, at least as regards the new patterns of harms chronicled in this Article, these crucial procedural tools are fit for purpose. The next Subpart explains why a range of procedural hurdles currently plaguing aggregate litigation and private enforcement may combine to create a perfect storm for would-be plaintiffs in such suits, threatening to keep patterns of algorithmic harm in the dark. Those hurdles notwithstanding, I will explain in the following Subpart how certain procedural strategies and arguments may still offer hope for enterprising plaintiffs’ attorneys in such cases.

B. Procedural Shortcomings

No one procedural hurdle on its own is responsible for the somewhat bleak situation facing would-be plaintiffs in suits arising out of algorithmic consumer profiling. Nor would any single procedural reform be enough to turn things around. Instead, this Subpart explains how an accretion of multiple changes in our procedural law operates to make litigation particularly ill-suited for shedding light on these new patterns of harm, and at precisely the moment when it may be most needed. Drawing from both case law and procedure scholarship, I focus on issues within the class action device,

²⁰³ Alexandra D. Lahav, *The Political Justification for Group Litigation*, 81 FORDHAM L. REV. 3193, 3210 (2013).

²⁰⁴ Elizabeth Chamblee Burch, *Litigating Groups*, 61 ALA. L. REV. 1 (2009) [hereinafter *Litigating Groups*]; Elizabeth Chamblee Burch, *Litigating Together: Social, Moral, and Legal Obligations*, 91 B.U. L. REV. 87 (2011) [hereinafter *Litigating Together*].

²⁰⁵ See, e.g., Brooke D. Coleman, *The Vanishing Plaintiff*, 42 SETON HALL L. REV. 501, 506 n.20 (2012).

multidistrict litigation, and Article III standing doctrine, while also flagging additional areas where our procedural rules may be falling short.

1. Class Actions

Many of the relevant problems with the class action device cluster around class certification. The new harmed groups are likely to be extremely heterogeneous along a number of dimensions. Making matters even more difficult, the number of data points that go into an algorithmic decision as to any given consumer means that, at least on one view, aggregating even any two consumers (let alone a vast, newly emergent group) poses serious certification challenges. In short, the ways in which the relevant plaintiffs are different jump off the page, while the ways in which they are the same are often hidden, counterintuitive, and dependent on litigation itself to uncover. In this way, procedural law exacerbates the isolation and atomization that these tools already engender among affected consumers.

One of the biggest challenges to certifying a new harmed group as a Rule 23 class is the requirement for commonality under Rule 23(a)(2). Although the standard in the Rule's text, that "there are questions of law or fact common to the class," doesn't sound especially rigorous, that changed with the Supreme Court's decision in *Wal-Mart Stores, Inc. v. Dukes*.²⁰⁶ The Court in *Dukes* clarified that to satisfy the commonality criterion, class members must have "suffered the same injury."²⁰⁷ While there can be variations in the precise cause of action or legal theory within the class, all members of the class "must depend upon a common contention," and that contention "must be of such a nature that it is capable of classwide resolution—which means that determination of its truth or falsity will resolve an issue that is central to the validity of each one of the claims in one stroke."²⁰⁸

The plaintiffs in *Dukes* — female employees of Wal-Mart who alleged that the very policy of entrusting pay and promotion decisions to the discretion of management at each

²⁰⁶ 564 U.S. 338 (2011).

²⁰⁷ *Id.* at 350 (citation omitted).

²⁰⁸ *Id.*

of the stores resulted in unlawful discriminatory outcomes — did not pass the Court’s new test. The Court found that the Plaintiffs’ aggregate methods of “statistical proof,” especially in combination with the decentralized and variable nature of the relevant decisions, could not establish any contention that would apply uniformly to all members of the class.²⁰⁹ Thus, the class was improperly certified. The Court summarized the problem by quoting with approval the dissenting opinion below, which noted that class members “held a multitude of jobs, at different levels of Wal-Mart’s hierarchy, for variable lengths of time, in 3,400 stores, sprinkled across 50 states, with a kaleidoscope of supervisors (male and female), subject to a variety of regional policies that all differed They have little in common but their sex and this lawsuit.”²¹⁰

Scholars have identified the *Dukes* decision as a major component of the Supreme Court’s effort to “curb[] class actions”²¹¹ in general. But it is easy to see how this decision could pose a particularly acute problem for class actions brought on behalf of the new harmed groups. The objections about the heterogeneous class in *Dukes* sound almost as if they are describing the groups of similarly situated consumers described in Part II, people who vary along all sorts of dimensions but who are all subject to adverse decisions *at least in part* because of a feature they do share. Furthermore, many if not all cases brought by algorithmically harmed consumers would depend on aggregate statistical measures of proof to demonstrate both the existence of legally relevant patterns and their own place within those patterns. At least on one view, the bespoke, individualized nature of algorithmic decisions in the consumer marketplace could make it extremely challenging to arrive at conclusive proof that would apply to all group members in “one stroke,” at least as the Court explains the concept in *Dukes*.²¹²

²⁰⁹ *Id.* at 356–57.

²¹⁰ *Id.*

²¹¹ Judith Resnik, “Vital” *State Interests: From Representative Actions for Fair Labor Standards to Pooled Trusts, Class Actions, and MDLs in the Federal Courts*, 165 U. PA. L. REV. 1765, 1776 (2017).

²¹² Julie Cohen has identified a similar interplay between these procedural requirements and (the difficulty of bringing) suits against large data-driven

A second roadblock to certification in algorithmic consumer profiling cases is what some circuits have termed an “ascertainability” requirement, *i.e.*, that proponents of certification demonstrate to the court’s satisfaction at the time of certification that individual class members can be identified via an “objective and administratively feasible” procedure.²¹³ As Julie Cohen has pointed out, the “technologically opaque nature of data-driven algorithmic processes” makes it possible for defendants who use such techniques to argue that “the methods proposed for ascertaining the group of consumers affected by any particular practice are just too imprecise and that including all users of a particular service would threaten crippling liability”²¹⁴ While it is true that there is a circuit split on the issue of whether administrative feasibility and objectivity are independent requirements for certification under Rule 23,²¹⁵ even circuits that do *not* recognize such a requirement may nonetheless conduct a similar analysis as part of the Rule 23(b)(3) inquiry.²¹⁶

Third, putative class actions may now not even get far enough to see a certification motion denied. That is because of the string of recent Supreme Court cases interpreting the Federal Arbitration Act to require enforceability of class action and class arbitration waivers in (*inter alia*) consumer contracts. Much ink has been spilled on the effects of these decisions, a full discussion of which is beyond the scope of this Article. But just as scholars have pointed out that these decisions are particularly harmful for the viability of “low-value, consumer-type claims” in general, especially where “a typical unsophisticated claimant may not even know she has a claim,”²¹⁷ such concerns apply *a fortiori* for cases of algorithmic consumer profiling. It is this special danger I wish to emphasize, *i.e.*, yet another procedural trend that threatens to keep already-invisible patterns of harm hidden away.

platforms. JULIE E. COHEN, *BETWEEN TRUTH AND POWER: THE LEGAL CONSTRUCTIONS OF INFORMATIONAL CAPITALISM* 159–60 (2019).

²¹³ Geoffrey C. Shaw, *Class Ascertainability*, 124 YALE L.J. 2354 (2015).

²¹⁴ COHEN, *supra* note 212, at 161.

²¹⁵ *Cherry v. Dometic Corp.*, 986 F.3d 1296, 1302 (11th Cir. 2021).

²¹⁶ *Id.*

²¹⁷ J. Maria Glover, *The Supreme Court’s “Non-Transsubstantive Class Action”*, 165 U. PA. L. REV. 1625, 1652–53 (2017).

The examples above are meant to illustrate how shifts in class action doctrine may make it more likely that the new patterns of algorithmic harm will fall through the cracks. This concern ought to be heightened even further by the pile-on of other procedural barriers.

2. Multidistrict Litigation

Class actions are, of course, not the only vehicle for aggregating the claims of multiple plaintiffs. In fact, evidence suggests that they may no longer even be the preeminent such procedural route. The rise of multidistrict litigation (or MDL) in recent years has been meteoric, with MDLs representing between 22% and 50% of all filed civil cases.²¹⁸ And certain opportunities available through the MDL mechanism, such as the ability to proceed with joint discovery *without* satisfying the commonality or predominance requirements of Rule 23, ought to be a good match for the sorts of shared harms identified earlier in this Article. But as has recently been noted by scholars of MDLs, procedural protections — especially for plaintiffs — have not caught up.²¹⁹ Two features of MDLs as they have developed in recent years make them particularly ill-suited for helping shed light on algorithmic consumer profiling: *Lone Pine* orders and the dearth of substantive adjudications.

Scholars have referred to MDLs as the “Wild West” of procedure, in which individual judges enjoy an unparalleled amount of leeway to craft ad hoc procedural rules to manage the joint pretrial proceedings.²²⁰ Left to fend largely for themselves without detailed rulemaking by the Judicial Panel on Multidistrict Litigation (JPML), judges also turn to each other for inspiration, resulting in “contagious” procedural innovations that spread far beyond their original contexts.²²¹ One particular innovation that has spread in such a fashion is the so-called *Lone Pine* order (named after its case of origin), in which each individual plaintiff inside an MDL must “supply prima facie evidence of injury, exposure, and causation—all by

²¹⁸ Burch & Gluck, *supra* note 186, at 227.

²¹⁹ *Id.*

²²⁰ *Id.*

²²¹ Nora Freeman Engstrom, *The Lessons of Lone Pine*, 129 YALE L.J. 2, 15 (2019).

a set date, under penalty of dismissal.”²²² And crucially, these orders frequently require plaintiffs to come forward with such information before discovery.²²³

The rise of *Lone Pine* orders and similar case-management techniques at the pre-discovery stage spells serious trouble for would-be algorithmic profiling plaintiffs. Any given plaintiff would almost certainly lack access to the requisite information without discovery into the underlying algorithm. To the extent some plaintiffs have enough information to survive dismissal and others don’t, *e.g.*, because some subtypes of discrimination are clearer or better-established *ex ante* than others, the perhaps unexpected commonalities and bonds of solidarity envisioned by this Article might never get the chance to form. (Put differently, *Lone Pine* orders hinder the discovery of new patterns of harm in favor of defaulting to patterns that are obvious at the point of filing.) As a result, plaintiffs whose harms may only be fully comprehensible *after* the “associational” activity of litigation takes place are deprived of the opportunity to participate.

The other feature of MDLs that nips their democracy-enhancing potential in the bud is the almost complete lack of substantive adjudications on the record. As Gluck and Burch note, although MDLs are in theory meant to aggregate cases only for pretrial proceedings before remanding back to their courts of origin, that is not how things play out in practice. Instead, judges in recent decades have “resolved 99% of cases themselves rather than sending them home as the statute envisions.”²²⁴ Those resolutions “almost always” take the form of a settlement.²²⁵ The predictable consequence is that MDLs are not able to perform the “public-regarding” functions of litigation, including “democratic participation and development of legal doctrine.”²²⁶

As discussed in Part II, communication among members of the newly harmed groups isn’t merely desirable for their own benefit. It can also help “draw on the situated knowledge of the

²²² *Id.* at 5, 17 n.58 (discussing prevalence).

²²³ *Id.* at 40–41 (noting disagreement among courts on this issue).

²²⁴ Burch & Gluck, *supra* note 186, at 228.

²²⁵ *Id.*

²²⁶ *Id.* (citing Owen Fiss, *Against Settlement*, 93 YALE L.J. 1073, 1085 (1984)).

people located in different group positions as resources for enlarging the understanding of *everyone* . . . “²²⁷ Latent groups of invisible allies exist not just among the potential plaintiffs in a given MDL, but also among the people who would hear about it, learn something about how a related technology was affecting their own lives, or update their own conceptions of fairness to account for what the lawsuit uncovered. Given the system as it currently operates, the insights gained over the course of an MDL would often fail to benefit anyone outside the four corners of the case.

3. Article III Standing

The procedural moves discussed above impede the effectiveness of aggregation as a tool of democratic deliberation. But another procedural shift has weakened the very foundation of private enforcement itself. As has been widely noted, the Supreme Court’s recent jurisprudence on Article III standing, and in particular the injury-in-fact prong, has made it increasingly difficult for plaintiffs to get past the starting gate when a statute (or potentially even the common law)²²⁸ would otherwise grant them enforcement rights. Two aspects of standing jurisprudence are particularly tricky for would-be algorithmic profiling plaintiffs: the treatment of risk and the procedural posture at which the issue is decided.

In the wake of the Supreme Court’s decisions in *Clapper v. Amnesty International USA*²²⁹ (2013) and *Spokeo v. Robins* (2016),²³⁰ scholars observed that it was no longer clear whether or when the injury-in-fact requirement for Article III standing could be satisfied by a demonstrated risk of future harm.²³¹ In particular, Solove and Citron argued that cases involving data breaches occupied precisely the unsettled area where courts had split on such questions.²³² In her own assessment of the case

²²⁷ YOUNG, *supra* note 137, at 109 (emphasis added).

²²⁸ See *infra* notes 236–37 and accompanying text (discussing *Dinerstein v. Google, LLC*, 73 F.4th 502 (7th Cir. 2023)).

²²⁹ 568 U.S. 398 (2013).

²³⁰ 136 S. Ct. 1540 (2016).

²³¹ See, e.g., COHEN, *supra* note 211, at 150; Daniel J. Solove & Danielle Keats Citron, *Risk and Anxiety: A Theory of Data-Breach Harms*, 96 TEX. L. REV. 737 (2018).

²³² COHEN, *supra* note 212.

law at the moment when those two cases represented the Supreme Court's latest word on the issue, Julie Cohen found that data breach cases increasingly passed muster, with courts "reasoning that 'customers should not have to wait until hackers commit identity theft or credit-card fraud'" and finding that such risks constituted an "objectively reasonable likelihood of future harm."²³³ Interestingly, Cohen saw things moving in a different direction when cases involved the sorts of harms covered in this Article, as courts viewed "risks flowing from increased vulnerability to data-driven predictive profiling [as] too remote and speculative to count as actionable injuries."²³⁴

In the years since these assessments, standing doctrine has grown unequivocally more hostile to risk-based claims. So if the baseline for algorithmic profiling cases was already poor (or at least worse off than the more borderline cases about data breaches), the situation now appears to be dire. In the 2021 case *TransUnion v. Ramirez*,²³⁵ the Court considered whether plaintiffs who had been wrongly labeled as potential terrorists in TransUnion reports (due in part to unlawfully lax data quality controls) could move forward with their Fair Credit Reporting Act claims. Distinguishing between those class members whose incorrect "potential terrorist" labels had, and had not been, already passed on to third parties, the Court held that the latter group's harm was not "concrete" enough to qualify as an injury-in-fact. In so holding, the Court clarified that "the mere risk of future harm, standing alone," could not confer standing in a suit for damages, even if it supported a claim for injunctive relief to combat such a risk.²³⁶

The consequences of this decision are only starting to become clear, but the early signs are worrisome. In a notable recent case, *Dinerstein v. Google LLC et al.*,²³⁷ the Seventh Circuit found that the plaintiff could not sue based on anonymized medical records that were passed on to Google—irrespective of the merits of the "breach of medical

²³³ *Id.*

²³⁴ *Id.*

²³⁵ 594 U.S. 413 (2021).

²³⁶ *Id.* at 433.

²³⁷ 73 F.4th 502 (7th Cir. 2023).

confidentiality” claim — because the risk of future reidentification was not sufficiently concrete under *TransUnion*. (The court also reached the eyebrow-raising conclusion that even actions for breach of contract that would have passed muster at common law might nonetheless fail the Supreme Court’s new concreteness tests if independent harm beyond the breach itself were not alleged.)²³⁸ Combined with the reasoning in *Dukes* on statistical modes of proof, this string of decisions on standing operates as a double whammy for algorithmic profiling plaintiffs who cannot pinpoint precisely how a given data point is being, or might be, used against them (notwithstanding robust data about how it is generally used).

Importantly, the *Dinerstein* court threw out the risk-of-harm theory before anything beyond bare-bones initial discovery had occurred. This aspect of the case points to a broader problem: while *TransUnion* itself dealt with a case that had been tried all the way to a jury verdict, courts applying its language now routinely dismiss cases on Article III grounds much earlier in their life-cycles. As discussed in the following Subpart, plaintiffs may have some opportunities to counter this trend. But as it stands, the Supreme Court has created a situation in which lower courts are dismissing putative class actions arising from potentially risky data flows before any of the information surfacing, collective communication, or norm development can occur.

4. Additional Procedural Challenges

Because aggregation and private enforcement are so central to the democracy-enhancing potential of litigation, I focus above on procedural issues that threaten specifically to erode their proper functioning. Unfortunately for would-be plaintiffs in algorithmic profiling cases, there are other areas of procedure that may pose further challenges to viable suits.

In particular, disputes about “where” the relevant conduct occurs threaten to make it even more difficult for the new harmed groups to make themselves visible or gain redress through litigation. In 2013, scholars were already warning that “critical procedural questions” prompted by the uncertain geographic nexus of cloud computing were being

²³⁸ *Id.*

“overlooked.”²³⁹ Yet more than ten years later, the situation has gotten more unsettled, not less. As a recent article explains, the Supreme Court’s recent jurisprudence on personal jurisdiction has yet to successfully clarify how jurisdictional doctrine should apply to cases involving “online interactions” and “deterritorialized data.”²⁴⁰ This muddle is an opportunity for sophisticated, well-lawyered defendants and a nightmare for plaintiffs. In future research, I intend to show in detail how defendants take advantage of unclear jurisdictional rules in the context of cross-border virtual harms to bog down litigation in months and even years of procedural wrangling.

The list of potential procedural hiccups for would-be algorithmic profiling plaintiffs goes on even further. To name just a few, these include: whether the prevailing interpretations of discovery rules allow plaintiffs’ counsel, especially at the pre-certification stage, to gain meaningful access to the algorithm in question; how the requirement in Fed. R. Civ. P. 11(b) to conduct an “inquiry reasonable under the circumstances” before filing suit applies to theories of harm that depend on proprietary algorithms; and how protective orders make it difficult for plaintiffs to build on each other’s successes in uncovering the new patterns of harm.

C. Green Shoots and Paths Forward

The procedural issues identified in the preceding Subpart may paint a bleak picture. And they do add up to a serious challenge if litigation is to play a meaningful role in helping shed light on the new harmed groups. But there are also procedural arguments and strategies that may help would-be plaintiffs or their advocates fight back in such cases. Below, I aim to develop a preliminary sketch of what this procedural toolkit might look like. Not all of these approaches would be appropriate in every case, nor would they (even in combination) defang the full range of problems identified above. Rather, the aim of this Subpart is twofold: (1) to ensure that those fighting to uncover these patterns of algorithmic

²³⁹ Damon C. Andrews & John M. Newman, *Personal Jurisdiction and Choice of Law in the Cloud*, 73 MD. L. REV. 313 (2013).

²⁴⁰ Paul Schiff Berman, *The Future of Jurisdiction*, 102 WASH. U. L. REV. 1169, 1207–12 (2025).

harm are not leaving useful options on the table, and (2) to lay the groundwork for a research agenda focused on making sure our mechanisms of private enforcement and aggregation can enable meaningful oversight over the emerging technologies with which this Article is concerned.

1. Class Actions

Rule 23 class actions are a particularly fertile ground for procedural moves that plaintiffs may not currently be taking sufficient advantage of. First, rereading *Dukes* in 2024, in the context of the technologies discussed above, suggests that one of its central arguments could be turned on its head. Recall that in *Dukes*, a key problem for the Court was that the decentralized, discretionary, store-by-store decisions about pay and promotions made it impossible to resolve any dispositive issues on a class-wide basis. In some salient senses, the variegated and heterogeneous groups of consumers envisioned in this Article might inspire similar concerns.

But there is a big difference between how automated decision-making usually works and the facts in *Dukes*: the rise of machine learning has often served to remove discretion from consumer-facing employees and instead re-centralize it in the form of a companywide algorithm. The upshot is that even when millions of variables are involved, it may be possible to answer at least *some* dispositive questions about the decision-making calculus used on an entire class, thereby satisfying the commonality requirement of Fed. R. Civ. P. 23(a)(2). As of this writing, I am unaware of a putative class in an algorithmic profiling case that has made this style of argument. I flag it here with the expectation, borne partly of personal litigating experience, that concerns about commonality often stifle such cases before they ever get filed.

A second strategy that could be a particularly good match for the new harmed groups is the certification of “issue classes” under Fed. R. Civ. P. 23(c)(4), which provides that “[w]hen appropriate, an action may be brought or maintained as a class action with respect to particular issues.” In recent years, scholars have brought attention to this “long-neglected” procedural tool, which offers plaintiffs the opportunity to work together to “decouple[] the inquiry into the defendant’s

conduct from questions regarding the eligibility of individual claimants for relief.”²⁴¹ In so doing, litigants can avoid the certification headaches discussed above, which tend to focus on differences among how individual plaintiffs are situated. When used successfully, this strategy can result in declaratory judgments that individual plaintiffs (or *groups* of plaintiffs) can then use in follow-on suits for damages.²⁴² Although Gilles and Friedman do not specifically mention the particular utility of the 23(c)(4) issue class for broadly dispersed algorithmic harms, the procedural payoffs they identify could hardly be better suited for the groups envisioned in Part II. Once a declaratory judgment establishes that, say, a seller committed an unfair or deceptive practice by increasing prices based on a perceived characteristic or behavioral trait, affected consumers have a tremendous head start for any future actions, while everyone else has the chance to update their empirical knowledge and distributive norms to account for such a finding.

Third, there is what one scholar has referred to as “the overlooked utility of the defendant class action.”²⁴³ Scholars and litigants alike have devoted surprisingly little attention to the fact that Rule 23 allows the aggregation not just of plaintiffs, but also of defendants, providing that “[o]ne or more members of a class may sue or *be sued* as representative parties on behalf of all members.” Fed. R. Civ. P. 23(a) (emphasis added). Once again, the potential utility of this tool for plaintiffs in consumer profiling cases has gone essentially unmentioned. But as discussed in Parts I and II, a key feature of algorithmic decision-making may make this option an excellent fit, namely, the repeat use of similar algorithms or profiling tools by multiple marketplace actors. If the same patterns of harm or unlawful treatment are visited upon a consumer or group of consumers by multiple sellers, the potential value of a combined suit starts to grow significantly, as would the incentive for plaintiffs’ attorneys to file. Moreover, just as joint discovery can help groups of plaintiffs

²⁴¹ Myriam Gilles & Gary B. Friedman, *The Issue Class Revolution*, 101 B.U. L. REV. 133, 136 (2021).

²⁴² *Id.*

²⁴³ Francis X. Shen, *The Overlooked Utility of the Defendant Class Action*, 88 DENV. L. REV. 1 (2010).

discover patterns amongst themselves, so too can discovery in a defendant class action reveal patterns among wrongdoers that might not be apparent if looking at only one defendant in isolation. Although not a consumer profiling case, at least one plaintiff class action involving the unlawful use of an artificial intelligence tool (in this case a facial recognition app) *also* proceeded as a defendant class action, attempting to deal in one fell swoop with all the third-party retailers that purchased and used the app.²⁴⁴

2. Multidistrict Litigation

As with class actions, MDLs offer some useful potential procedural options for would-be consumer profiling plaintiffs seeking to uncover new patterns of shared harm. Importantly, the procedural openness and flexibility identified as sources of mischief in the recent MDL literature²⁴⁵ may also represent untapped opportunities for the groups of plaintiffs with which this Article is concerned. In her earlier work on MDLs, Elizabeth Burch focused on the ways in which groups of plaintiffs could work together in precisely the ways envisioned in Part III.A, *i.e.*, sharing information, developing joint narratives, and making justice claims.²⁴⁶ In particular, Burch advocated for the appointment of a “special officer” whose job it would be to “facilitate deliberation and communication among plaintiffs.”²⁴⁷ As she envisioned it, such special officers would help groups of plaintiffs to “specify, clarify, and prioritize their ends and values;” “determine which ends, claims, and injuries are commensurate,” and “encourage problem-solving and collective decision-making about litigation strategy,” among other difficult tasks of collective communication. Burch also argues persuasively that the same sources of authority for appointing special masters for evidentiary purposes would also allow judges to appoint such

²⁴⁴ *In re Clearview AI, Inc., Consumer Priv. Litig.*, No. 21-CV-0135, 2022 WL 3226777, at *1 (N.D. Ill. Aug. 10, 2022). Disclosure: the author was an attorney for the plaintiffs.

²⁴⁵ Burch & Gluck, *supra* note 186, at 228.

²⁴⁶ *Litigating Groups*, *supra* note 204; *Litigating Together*, *supra* note 204.

²⁴⁷ *Litigating Together*, *supra* note 204, at 113.

special officers under current law.²⁴⁸

While Burch's proposal does not seem to have taken off, the time is ripe to revive it. A special officer would be especially helpful in situations where the incentives of attorneys appointed by the court to other MDL leadership roles are not oriented toward uncovering invisible patterns of harm, *e.g.*, because they are interested in the fastest settlements, most straightforward legal claims, or fact patterns that are already well understood.²⁴⁹ (On this logic, a special officer might have utility in *most* cases of diffuse, poorly understood algorithmic profiling.) Unlike other members of the appointed MDL leadership, the special officer would be charged explicitly with building channels for intra-plaintiff communication and with enabling collective decision-making.²⁵⁰ To the extent individual plaintiffs or groups of plaintiffs come to court with litigation goals that include the sorts of democratic deliberation functions contemplated in this Article, it may therefore be worthwhile to petition for an appointment of a Burch-style special officer.

3. Article III Standing

Finally, notwithstanding the developments in Article III standing doctrine that make private enforcement in general (and consumer profiling suits in particular) especially difficult, standing doctrine also suggests a path forward for the new harmed groups. The doctrine of associational standing allows advocacy groups to bring suit on behalf of their members so long as at least one member has suffered a cognizable injury, the suit would help vindicate an interest germane to the organization's interests, and the participation of the members themselves in the adjudication would not be necessary.²⁵¹ As critics of this doctrine have pointed out, the "at least one

²⁴⁸ *Id.* (citing Fed. R. Civ. P. 53, Fed. R. Evid. 706, and "judges' inherent equitable authority").

²⁴⁹ *See supra* Part III.B.

²⁵⁰ David L. Noll, *What Do MDL Leaders Do?: Evidence from Leadership Appointment Orders*, 24 LEWIS & CLARK L. REV. 2, 29–30 (2020) (describing a much less robust set of defined duties vis-a-vis plaintiffs as the norm in appointment orders).

²⁵¹ Michael T. Morley & F. Andrew Hessick, *Against Associational Standing*, 91 U. CHI. L. REV. 1539, 1560–62 (2024).

member” prong is “not onerous,” as a “group is not required to show that a significant portion of its membership has been harmed; injury to a single member suffices.”²⁵² If this sounds like an end run around both the requirements of Rule 23 (such as commonality and predominance) and Article III itself (by empowering an entity other than the harmed party to litigate the case), critics certainly think it is.²⁵³ I share their analysis, but not their normative prescription. Given the serious risks that procedural rules will leave diffuse, atomized plaintiffs to fend for themselves — letting algorithmic profiling tools engage in unchecked social sorting — associational standing may prove a powerful means to level the playing field.

Conclusion

At this stage, what we don’t know about the new patterns of algorithmic harm dwarfs what we do know. But as this Article argues, enough information has emerged to cause genuine concern. More than the dollars and cents of any particular transaction, the deeper problem is that relative advantages and disadvantages are being visited upon groups of citizens without any opportunity for us to evaluate whether we approve of such social sorting. The question of which distributive logics or outcomes are compatible with our distributive norms — and in which contexts — is a fundamentally political one that must be subject to the democratic process. This can only happen if our system of aggregate litigation is able to bring the invisible to light and thereby engender meaningful opportunities for deliberation and collective action. Hopefully, by laying out a procedural roadmap plaintiffs can use while also highlighting areas where reform efforts may be necessary, this Article can help spur a much-needed effort to align our procedural law with these new technological and social realities.

²⁵² *Id.* at 1561.

²⁵³ *Id.* at 1563 (highlighting tensions with Article III purposes); *id.* at 1572–82 (arguing that the doctrine is inconsistent with Rule 23).