The Case Against Algorithmic Price Discrimination

This paper assesses the risk of algorithmic price discrimination, which occurs when companies use sophisticated machine learning algorithms to offer different prices to different customers based on data such as customers' past purchases, web browsing history, location, or age, race, and gender. I argue that the standard economic account of price discrimination overlooks unexpected behavior by pricing algorithms: for example, pricing algorithms may learn to *nudge* or *inflate* consumer willingness to pay rather than merely *identifying* willingness to pay, contra the assumptions of the standard economic view. This paper highlights the limitations of the standard economic view and argues that we should instead adopt what I call the egalitarian view, which favors equal prices for all consumers.

Introduction

Throughout the 1980s and 1990s, economists developed the *standard economic view*, which generally favored price discrimination. According to the standard economic view, price discrimination increases efficiency because it enables a greater number of market transactions. In addition to the efficiency argument, some economists also argued that price discrimination could improve distributional equity, insofar as companies could charge higher prices to higher-income consumers and lower prices to lower-income consumers. One well-known economics textbook, for example, cited a small-town doctor who charged only what his patients could afford to pay.[1]

Yet the standard economic view remained largely theoretical for decades. Companies lacked the technology needed to calculate personalized prices, and they also lacked detailed personal information about consumers. While companies could engage in broad-grained price discrimination (for example, offering senior discounts or student discounts), most companies could not engage in fine-grained personalization (which economists called "perfect" price discrimination).[2]

Today, machine learning algorithms have transformed price discrimination from an economic theory into a practical business strategy. Companies have adopted sophisticated pricing algorithms that can incorporate information like consumer location, past purchases, browsing history, and other personal information like age, race, or gender. While few companies have disclosed the details of their pricing schemes, algorithmic pricing is now widespread.

In response to this revolution in pricing, many economists and policymakers have adopted the standard economic framework, which favors algorithmic price discrimination. Google's chief economist, for example, described price discrimination as "largely beneficial" because "you charge higher prices to people who can afford to pay higher prices."[3] While some policymakers are more cautious—a White House report noted the "promise and peril" of algorithmic pricing—that report, too, largely adopted the standard economic framework.[4]

---

[1] Hal R. Varian, *Intermediate Microeconomics: A Modern Approach* (1987) at pp 431.

[2] Per the standard economic view, the most efficient form of price discrimination is "perfect" price discrimination, or first-degree price discrimination, when sellers charge each consumer what the consumer is willing to pay. Other forms of price discrimination include second-degree price discrimination, when sellers charge consumers based on the quantity they purchase (i.e. discounts for purchasing in bulk) or third-degree price discrimination (such as student discounts or senior discounts), which Part II discusses in more detail.

[3] Natasha Singer, *The Government's Consumer Data Watchdog*, N.Y. TIMES, May 23, 2015.

[4] Executive Office of the President, *Big Data and Differential Pricing* (2015).

This paper argues that we should jettison the standard economic view and instead adopt the egalitarian view, which favors equal prices for all consumers. The argument proceeds in two parts.

Part I lays out the standard economic view and highlights three of its limitations. First, the standard economic view assumes that algorithms aim to *identify* willingness to pay rather than *nudging* it. As a result, the standard economic view overstates the efficiency gains from price discrimination. Second, the standard economic view smuggles in an implicit assumption about distributional equity—it assumes that algorithms will charge higher prices to higher-income or higher-wealth consumers and lower prices to lower-income or lower-wealth consumers. But in some cases, algorithms have price gouged low-income consumers who have fewer options than high-income consumers. Third, the standard economic view assumes away consumer differences in race, gender, and other protected characteristics. In so doing, it ignores the antidiscrimination concerns that arise when companies charge different prices to different demographic groups.

In Part II, the paper develops and defends the *egalitarian view*, which generally supports equal prices for all consumers. The egalitarian view, however, does allow some forms of price discrimination which are described by two provisos. The first proviso allows broad-grained discounts for groups like seniors or students, while the second proviso allows companies to pass along differential costs to consumers when some consumers are more "costly" than others. While the egalitarian view is not perfect, it faces fewer difficulties than the standard economic view.

I. The Standard Economic View

This section lays out the standard economic view and then describes three of its limitations. According to the standard economic view:

> <u>Standard Economic View:</u> Sellers may charge different prices to different buyers for identical goods and services.

The primary rationale for the standard economic view is *efficiency*. As economists argued in the 1980s and 1990s, price discrimination increases efficiency because it eliminates the monopoly deadweight loss.

Monopoly deadweight loss occurs when goods are priced above the marginal cost of production. In non-monopoly markets, competition drives prices down, so goods are typically priced at the marginal cost (which is most efficient for the market). In monopoly markets, however, the monopolist can charge a price p that is higher than the marginal cost.[5] But when a monopolist offers only a single standardized price *p* for a good *x*, some consumers want to purchase the good *x* (and are willing to pay more than the marginal cost) but cannot pay full price *p*. Transactions that would make everyone better off do not occur—hence, the monopoly deadweight loss. If, however, the company can offer a lower price just to those consumers, then both the consumer and the company benefit: more transactions occur and efficiency increases. By enabling a greater number of mutually beneficial transactions, price discrimination reduces monopoly deadweight loss and increases efficiency.

The standard economic view is thus centered solely on one value—efficiency. But as we will see in the subsequent discussion, algorithmic pricing raises concerns across multiple values, including privacy, autonomy, equity, and antidiscrimination.

Below, the paper discusses three limitations of the standard economic view: first, that it overstates the efficiency gains from price discrimination because it overlooks the existence of nudging algorithms; second, that it mistakenly assumes that price discrimination will benefit low-income consumers; third, that it overlooks antidiscrimination concerns related to consumers' race, gender, and ethnicity. While the standard economic view can attempt to carve out problematic instances of price discrimination, the carve-outs may contradict the profit motive of profit-oriented companies.

1. Overstates the efficiency gains from price discrimination, because it overlooks the existence of algorithms that *nudge* consumer willingness to pay rather than just *identifying* it.

According to the standard economic view, price discrimination increases efficiency.[6] However, the standard economic view assumed that sellers would aim only to *identify* consumer willingness to pay, and it did not contemplate that sellers might attempt to *nudge* or *inflate* willingness to pay. This difference is significant, because it affects the efficiency analysis: pricing algorithms that nudge willingness to pay may reduce efficiency rather than increasing efficiency. This section explains why algorithmic nudging may decrease efficiency rather than increasing it, argues that algorithms may learn to exploit consumer misperceptions, and finally explains why algorithmic nudging may be difficult to stamp out.

---

[5] Per Mankiw, "price discrimination is not possible when a good is sold in a competitive market… For a firm to price discriminate, it must have some market power." Gregory Mankiw, Principles of Economics (1997) at pp 323.

[6] Economists use different definitions of efficiency, but roughly speaking, an economy is Pareto efficient when "there is no alternative allocation that leaves everyone at least as well off and makes some people strictly better off." Hal R. Varian, *Intermediate Microeconomics: A Modern Approach* (1987) at pp 15.

As we discussed above, price discrimination increases efficiency because it eliminates the monopoly deadweight loss. This happy outcome, however, is less certain when price discrimination aims to *nudge* consumer willingness to pay rather than merely *identify* it. When monopolists nudge willingness to pay, they may induce consumers to enter into transactions that do not benefit them. Oren Bar-Gill cites the example of a gym membership to illustrate how consumers can misperceive the value of their purchases: a consumer might purchase an expensive gym membership thinking she will attend once a week, but in reality she attends only once per month, so she derives less value from the purchase than she had initially perceived. Her transaction hurts her, since she gained less value than she thought she would, but it benefits the gym company (so consumer surplus falls but producer surplus increases). If the consumer's misperception is large enough, however, then her transaction hurts not just her, but also decreases social surplus overall, thus reducing efficiency. Per Bar-Gill, "when the misperception is strong[] … price discrimination definitely decreases efficiency."[7]

The standard economic view did not account for misperceptions and other behavioral biases among consumers. Instead, the standard economic view assumed that consumers entered only into transactions that benefited them, and it assumed that consumers had a fixed, pre-existing willingness to pay. Because of these assumptions, the monopolists in the standard economic view never even tried to *nudge* willingness to pay (indeed, doing so might have been impossible since willingness to pay was assumed to be fixed and pre-existing); instead, monopolists sought only to *identify* the consumer's pre-existing willingness to pay. This assumption was embedded in the standard economic view developed throughout the 1980s and 1990s: one textbook, for example, described a firm that "knows exactly the willingness to pay of each customer,"[8] while another described companies that "assess" a customer's willingness to pay.[9] Nowhere did the textbooks discuss monopolists who attempted to nudge, inflate, or otherwise influence willingness to pay.

Decades later, when researchers began studying algorithmic pricing, they imported the standard view's old assumptions about willingness to pay. For example, a 2021 business journal article described digital tracking that allows companies to "identify[] . . . individual willingness to pay," yet the article entirely overlooked the possibility that companies might inflate willingness to pay rather than

---

[7] Oren Bar-Gill, *Algorithmic Price Discrimination: When Demand is a Function of Both Preferences and (Mis)Perceptions*, 86 UNIV. CHI. L. REV. 217 (2019) at pp 241. Bar-Gill uses the language of "misperceptions," while I use the language of "nudging," but we are largely talking about the same thing.
[8] Gregory Mankiw, Principles of Economics (1997) at pp 327.
[9] Robert Pindyck & Daniel Rubinfeld, Microeconomics (1992) at pp 375.

just identify it.[10] Similarly, a 2020 chapter on algorithmic pricing explained that sellers are "trying to discern consumers' willingness to pay," but it did not contemplate that sellers might inflate willingness to pay not just "discern" it.[11] Policymakers, too, have adopted the erroneous assumptions embedded in the standard view: the White House report noted that sellers will "try to predict how buyers will behave," but the report overlooked how sellers might nudge buyer behavior, not just predict it.[12]

To nudge willingness to pay, algorithms may learn to exploit human behavioral biases. Ryan Calo, for example, wrote that algorithms may "trigger irrationality or vulnerability in consumers" to induce them to buy more products.[13] Bar-Gill, too, found that Facebook and Google may be implementing "cognitive services" to build psychological profiles of their users, and that companies have experimented with "'assessing someone's personality by sifting through their writings.'"[14] Pricing algorithms might learn to vary the content of user's news feeds (as Facebook did during a 2014 experiment about social contagion of emotions), or they might use other methods to inflate willingness to pay.[15] These practices may infringe upon users' privacy and autonomy, values which the standard economic view typically overlooks.

A proponent of the standard economic view might ask—is algorithmic nudging really so bad? After all, it seems pretty similar to advertising, which inflates demand by persuading consumers to buy more goods and services. Like algorithmic nudging, advertising may also decrease economic efficiency,[16] yet advertising is usually considered morally harmless. Most advertisements are perfectly legal, as long as they do not cross the line into outright fraud.

This paper, however, does not need to develop a full account of the relative badness of algorithmic nudging vs. advertising. It is enough to note that, by overlooking the existence of nudging algorithms, the standard economic view has *overstated* the efficiency

---

[10] Peter Seele, Claus Dierksmeier, Reto Hofstetter & Mario D. Schultz, *Mapping the Ethicality of Algorithmic Pricing: A Review of Dynamic and Personalized Pricing*, 170 J. BUS. ETHICS 697, 705 (2021).
[11] Salil K. Mehra, *Algorithmic Competition, Collusion, and Price Discrimination*, CAMBRIDGE HANDBOOK L. ALGORITHMS 199, 207 (2020).
[12] Executive Office of the President, *Big Data and Differential Pricing* (2015) at pp 9 [hereinafter "White House Report"].
[13] Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995, 999 (2014).
[14] Bar-Gill, *supra* note 7, at 231.
[15] Vindu Goel, *Facebook Tinkers With Users' Emotions in News Feed Experiment, Stirring Outcry*, N.Y. Times, June 29, 2014.
[16] Within the economics literature, advertising is sometimes presented as benign or helpful to consumers, and at other times presented as distortive and efficiency-reducing. According to some economists, advertising "distorts the consumer's decisions" and wastes "real economic resources." Kyle Bagwell, 3 HANDBOOK OF INDUSTRIAL ORGANIZATION 1701, 1711.

gains from price discrimination. When algorithms exploit behavioral biases to nudge willingness to pay, price discrimination hurts consumers, and if consumer misperceptions are strong enough, price discrimination may even reduce efficiency. Rather than reliably increasing efficiency, price discrimination may at times reduce efficiency, contra the standard economic view.

In response, a proponent of the standard economic view might propose a modification to pricing algorithms, whereby companies adopted algorithms that exclusively identified willingness to pay rather than nudging it. This modification, however, is difficult for two reasons. First, it runs counter to companies' profit motives. Most companies are oriented toward profit, not toward values like privacy or autonomy. This profit motive means that companies will tend to select algorithms that generate the highest sales, and algorithms that nudge willingness to pay are presumably more profitable than algorithms that merely identify willingness to pay. Companies would have little incentive to carve out instances of algorithmic nudging.[17]

Second, even if companies try to constrain algorithmic nudging, contra their own profit motives, they may find it difficult to stop algorithms from nudging willingness to pay. A company, for example, might try to prevent nudging by reducing the number of "levers" the algorithm can adjust. If an algorithm can adjust only price, and not any other variables such as the user's news feed, then it has fewer ways to nudge consumers. But even if an algorithm has only one lever—price—it can still nudge willingness to pay, because it inevitably has another lever—time. An algorithm would learn to raise prices at times when consumers are most vulnerable, for example, when they are most sleep-deprived or stressed. So even if companies constrain algorithms, giving them only time and price as levers, algorithms can exploit time-based fluctuations in consumers' willingness to pay. While this is slightly conceptually distinct from nudging—the algorithm is not actively moving consumers' willingness to pay, but is instead capitalizing on times when willingness to pay is high—the result is nevertheless a type of algorithmic manipulation that many find troubling.

In sum, the standard economic view overstates the efficiency gains from price discrimination, because it mistakenly assumes that companies aim only to identify willingness to pay rather than nudging it. Algorithmic nudging may be difficult to stamp out, because even if companies tried to constrain their pricing algorithms, algorithms might still find ways to exploit behavioral biases.

---

[17] And this is assuming that it would be possible for companies to identify a neutral "Archimedean point" where willingness to pay stems solely from the consumer and is not influence by nudging.

2. Assumes, implicitly, that pricing algorithms will charge higher prices to people with higher ability to pay, and lower prices to people with lower ability to pay.

In addition to making an *explicit* argument about efficiency, the standard economic view also smuggled in an *implicit* argument about distributional equity. According to some economists, price discrimination allows companies to charge higher prices to higher-income or higher wealth consumers, and lower prices to lower-income or lower-wealth consumers, thus benefiting lower-income consumers. This argument was implicitly wrapped into the economics textbooks of the 1980s and 1990s, and it resurfaced explicitly in the 2010s when economists began to discuss pricing algorithms: according to Google's chief economist, pricing algorithms are "largely beneficial" because "you charge higher prices to people who can afford to pay higher prices."[18] Below, this section unpacks the claim about distributional equity, explains why it sometimes fails to bear out, and discusses how the standard economic view might respond.

In the economics textbooks of the 1980s and 1990s, the standard economic view implicitly assumed that monopolists would charge higher prices to wealthier or higher-income consumers, and lower prices to lower-income or lower wealth consumers. One well-known economics textbook, for example, cited a small-town doctor, who charged only what his patients could afford to pay.[19] Another textbook compared price discrimination to college financial aid programs, writing that financial aid programs were "similar to the behavior of any price-discriminating monopolist."[20]

It is important to note that the above assumption was implicit, not explicit, in the standard economic view. Explicitly, the standard economic view cited only *efficiency* as the justifying rationale for price discrimination. In the efficiency analysis, the benefits to low-income consumers are entirely irrelevant. Implicitly, however, the distributional equity argument was used to improve the public perception of price discrimination. By emphasizing the benefits to low-income consumers, the standard economic view made price discrimination more palatable for policymakers and for the public.[21]

Pricing algorithms, however, do not always follow the beneficent patterns of small-town doctors. Instead, some pricing

[18] Natasha Singer, *The Government's Consumer Data Watchdog*, N.Y. TIMES, May 23, 2015.
[19] Hal R. Varian, *Intermediate Microeconomics: A Modern Approach* (1987) at pp 431.
[20] Gregory Mankiw, *Principles of Economics* (1997) at pp 329.
[21] As I will discuss in more detail in Part II, we have some reason to be skeptical when companies claim that their pricing schemes advance distributional equity. If the goal is to help the worst-off, then wouldn't a tax and transfer scheme achieve that goal more directly? It seems odd to promote distributional equity goals via profit-oriented private companies.

algorithms have charged higher prices to low-income consumers who have fewer options. In 2012, for example, the retailer Staples designed an online pricing algorithm that personalized prices based on the consumer's location. The Staples algorithm tended to charge higher prices in lower-income neighborhoods, and lower prices in higher-income neighborhoods—not because lower-income consumers had more disposable income, but because their neighborhoods had fewer options.[22] According to researchers, "Staples appeared to be calculating prices based on the user's distance from a rival store, but the inadvertent effect was that people in lower-income ZIP codes saw the higher prices."[23] In 2020, researchers published similar findings about rideshare apps like Uber and Lyft, which tended to charge higher prices for rides to and from low-income neighborhoods in Chicago.[24] Researchers thought the disparity was driven by driver behavior and surge pricing, since drivers may "avoid these neighborhoods" due to fear of crime.[25]

How common is it for pricing algorithms to charge higher prices to lower-income customers? We simply do not know, because the data is not publicly available. The rideshare study, for example, was possible only because Chicago changed its regulations to compel companies to publish anonymized data. This data was later used by academic researchers, who otherwise would not have had access to it. Companies typically do not disclose their pricing data, because publishing their data might provide valuable information to their competitors. As a result, we have very few empirics about pricing algorithms.[26]

In response to the above examples, a proponent of the standard economic view might try to defend practices like the Staples algorithm. When a company charges higher prices in locations with few rivals, the higher prices incentivize rivals to enter the market. Over time, new competitors enter underserved locations, driving prices down. While customers might be overcharged in the short term, over the long term, the increased competition generates lower prices for consumers. So short-term higher prices eventually create a longer-term market correction.

This defense, however, is less convincing in scenarios where the market does not provide a longer-term correction, for example in the Uber/Lyft rideshare example. In the rideshare example, lower-

---

[22] Jennifer Valentino-DeVries, Jeremy Singer-Vine & Ashkan Soltani, *Websites Vary Prices, Deals Based on Users' Information*, WALL ST. J., Dec. 24, 2012.

[23] Julia Angwin, Surya Mattu & Jeff Larson, *The Tiger Mom Tax: Asians Are Nearly Twice as Likely to Get a Higher Price from Princeton Review*, PROPUBLICA, Sept. 1, 2015.

[24] Akshat Pandey & Aylin Caliskan, *Disparate Impact of Artificial Intelligence Bias in Ridehailing Economy's Price Discrimination Algorithms*, AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AAAI/ACM AIES 2021).

[25] *Id.* at 6.

[26] At the very least, this data should be made available to researchers.

income neighborhoods had higher prices because drivers tended to avoid those areas. Researchers have identified "so-called 'no-go zones,'" in cities like Rio de Janeiro, Johannesburg, and Atlanta, because Uber/Lyft drivers avoid high-crime areas.[27] To expand driver supply in no-go zones, Uber and Lyft use surge pricing that attracts drivers to areas they would otherwise avoid. In the short-term, surge pricing increases driver supply and generates higher prices in no-go zones. In the long-term, however, the market does not correct the error—drivers still avoid no-go zones due to fear of crime. Low-income areas continue to be charged higher prices, and the market does not provide a longer-term correction.

A proponent of the standard economic view might dig in his heels, and insist that the higher Uber/Lyft prices are justified because that is better than no Uber/Lyft at all. In this line of argument, surge pricing in low-income areas, which creates higher prices, is preferable to having no Uber/Lyft service at all in low-income neighborhoods.

That line of argument is certainly open to the standard economic view, and an economist might have no trouble biting that bullet. But I want to point out, at this juncture, how the argument has morphed over time. First, the standard economic view claimed that pricing algorithms were socially beneficial because they helped lower-income consumers. Next, when confronted with evidence that pricing algorithms sometimes overcharge low-income customers in retail deserts, the standard economic view replied that such overcharging was merely a short-term problem that would be corrected by longer-term market forces. Finally, when presented with a problem (crime) that the market cannot correct over the long-term, the standard economic view responded that surge pricing in low-income areas is preferable to no Uber/Lyft service at all. These moves are all open to the standard economic view, but they have pushed it painfully close to a contradiction: we started with a distributional equity claim that said "pricing algorithms are great because they will help the poor" and then arrived at very different claim that "price gouging the poor is okay, because that is better than providing no service at all." The initial claim about distributional equity morphed into a very different claim, one that is far less publicly palatable.

The fundamental worry, here, is that pricing algorithms do not necessarily track a customer's *ability* to pay (signified by her wealth/income) but rather a customer's *willingness* to pay, which may increase when the customer has few options or is vulnerable to exploitation. Some of the most troubling examples of price discrimination occur in the medical field, where willingness to pay rises when patients fall ill. A pricing algorithm might "learn" that some consumers have high willingness to pay for certain medical goods

[27] *Id.*

(because, unbeknownst to the algorithm, those consumers have a serious health condition). The algorithm would then raise prices for consumers with this illness, because it can charge them more than consumers without a medical condition. This kind of medical price gouging has already occurred, albeit non-algorithmically, to HIV patients who needed the drug Daraprim[28] and to people with allergies who needed EpiPens.[29] Algorithms could make the personalization even more fine-grained, to identify people with specific medical conditions. While this is a particularly distasteful case of price discrimination, it seems perfectly plausible, especially if companies are careless with their algorithms.

To avoid these kinds of troubling scenarios, a proponent of the standard economic view might propose a *qualified* standard view that carved out problematic instances of price gouging. In the qualified standard view, algorithmic price personalization would be acceptable except in cases where it charged higher prices to customers with medical conditions, or to low-income customers with fewer options. Companies would have to test their pricing data to ensure that their algorithms avoided price gouging in troubling scenarios.[30]

The difficulty here is the one we have already discussed—the profit motive—coupled with a problem of empirics. To carve out instances of problematic price discrimination, companies would have to work against their own profit motives and invest time and resources into monitoring their algorithms for price gouging. As a matter of empirics, companies might find it difficult to identify when price gouging is occurring, because companies do not typically have information about consumer wealth or income (which bears on ability to pay). And consumers would find it difficult to verify the carve-outs for price gouging, because most companies do not release their pricing data.

In sum, the standard economic view contains an implicit assumption that has proved imprecise. When the view was developed in the 1980s and 1990s, some economists assumed that perfect price discrimination would allow companies to charge higher prices to wealthier or higher-income consumers and lower prices to less wealthy or lower-income consumers. This assumption was not connected to the efficiency analysis, yet it improved the public perception of price discrimination. In practice, however, that assumption did not always

---

[28] *See*, *e.g.*, Heather Long, *Here's what happened to AIDS drug that spiked 5,000%*, CNN BUSINESS, Aug. 25, 2016; Andrew Pollack, *Drug Goes From $13.50 a Tablet to $750, Overnight*, N.Y. TIMES, Sept. 20, 2015.

[29] Toni Clarke, *U.S. lawmakers blast Mylan CEO over 'sickening' EpiPen price hikes*, REUTERS, Sept. 21, 2016.

[30] Price gouging is generally seen as morally problematic, but some ethicists have defended price gouging. Matt Zwolinski, *The Ethics of Price Gouging*, 18 BUSINESS ETHICS QUARTERLY 347 (2008).

bear out. As we saw in the Staples and Uber/Lyft examples, pricing algorithms sometimes charge higher prices to lower-income consumers with fewer options. The standard economic view can attempt to defend these practices, or alternatively to carve out exceptions, but the task is difficult without empirics.

3. Assumes away consumers' differences in race, gender, and ethnicity.

Finally, the standard economic view faces a final, and perhaps insurmountable, difficulty, because it lacks a theory of antidiscrimination that would justify charging different prices to consumers in different race and gender groups. Below, this section explains how the standard economic view assumed away differences in race, gender, and other protected characteristics, it next provides some legal background about antidiscrimination law, and it finally explores how the standard economic view might attempt to supply a theory of antidiscrimination (and why those attempts will likely fail).

The standard economic view assumes away differences in race, gender, and other protected characteristics. In the textbook discussions of price discrimination, consumers differed in terms of their tastes or in terms of their willingness to pay, but consumers were assumed to be homogenous in terms of race, gender, ethnicity, and other protected characteristics.

When policymakers began considering algorithmic pricing in the mid-2010s, they assumed that price discrimination would benefit "historically disadvantaged groups," because members of these groups were likely to be "more price-sensitive than the average consumer.[31] But this did not always hold in practice. Researchers have found that pricing algorithms charged higher prices to Black and Latine mortgage applicants, even when minority mortgage applicants were "risk-equivalent" to non-minority mortgage applicants.[32] Other researchers have found that Uber/Lyft charge higher prices to women than men,[33] and that test prep companies charged higher prices to Asian families, including low-income Asian families.[34]

It is not clear whether race and gender discrimination in consumer pricing is illegal. The Robinson-Patman Act prohibits some forms of price discrimination, but it aims mostly to protect small businesses from large businesses, and it has not been read as protecting

---

[31] White House Report, *supra* note 4, at 17.

[32] Robert Bartlett, Adair Morse, Richard Stanton & Nancy Wallace, *Consumer-Lending Discrimination in the FinTech Era*, 143 JOURNAL OF FINANCIAL ECONOMICS 30 (2022).

[33] Yanbo Ge, Christopher R. Knittel, Don MacKenzie & Stephen Zoepf, *Racial and Gender Discrimination in Transportation Network Companies*, NBER Working paper 22776 (2016).

[34] Julia Angwin, Surya Mattu & Jeff Larson, Julia Angwin, Surya Mattu & Jeff Larson, *The Tiger Mom Tax: Asians Are Nearly Twice as Likely to Get a Higher Price from Princeton Review*, PROPUBLICA, Sept. 1, 2015.

consumers. Some U.S. laws protect consumers against race and gender discrimination in housing, employment, and lending, but these laws cover only specific sectors of the economy and do not apply across all sectors of the economy. And constitutional requirements, such as the Equal Protection Clause and Due Process Clause, typically apply to state actors but not to private actors engaged in free market transactions.

Furthermore, many antidiscrimination laws protect consumers only against disparate treatment, not disparate impact. Disparate treatment, sometimes called intentional discrimination, occurs when businesses explicitly single out a protected class: for example, restaurants with signs that said "No Blacks or Mexicans allowed." Disparate impact, by contrast, occurs when a business does not explicitly single out a race or gender group, but instead adopts a policy that disproportionately impacts a protected class: for example, restaurants that required staff to wear their hair straight.[35]

Algorithmic bias often shows up as disparate impact, rather than disparate treatment, because many machine learning algorithms are trained on biased datasets that reflect historical patterns of discrimination. Algorithms that help companies screen job applicants, for example, were found to be biased against women: in 2014, Amazon adopted a recruiting algorithm that scanned applicants' resumes and flagged top candidates.[36] An internal review discovered that the algorithm downgraded resumés with the word "women," such as "women's chess club captain," and that it downgraded resumes from two all-women's colleges.[37] The algorithm had not been deliberately designed to be gender-biased but instead "learned" to replicate the biases in its dataset. Similar biases have been found in algorithmic job advertising, where Google ads showed high-paying jobs to men much more frequently than to women.[38]

In the face of these difficulties, how can the standard economic view respond?

<u>First</u>, the standard economic view might argue that differential pricing is acceptable as long as algorithms are formally race-blind and gender-blind. In adopting this line of argument, the standard view

[35] RICHARD THOMPSON FORD, DRESS CODES: HOW THE LAWS OF FASHION MADE HISTORY 330–35 (2022).
[36] Jeffrey Dastin, *Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women*, REUTERS, Oct. 9, 2018.
[37] *Id.* ("It penalized resumes that included the word "women's," as in "women's chess club captain. And it downgraded graduates of two all-women's colleges, according to people familiar with the matter. They did not specify the names of the schools.").
[38] Researchers evaluated Google ads with different user settings and found that "setting the gender to female resulted in getting fewer instances of an ad related to high paying jobs than setting it to male." Amit Datta, Michael Carl Tschantz & Anupam Datta, *Automated Experiments on Ad Privacy Settings: A Tale of Opacity, Choice, and Discrimination*, 11 Proceedings on Privacy Enhancing Technologies (2014).

adopts a "color-blind" theory of antidiscrimination. According to this color-blind theory, companies can charge different prices to different customers, as long as pricing algorithms do not explicitly take into account race or gender. If they avoid inputting information about race and gender into their pricing algorithms, then companies can avoid accusations of discrimination.

The problem with this color-blind theory is that, even if algorithms are formally race- and gender-blind, they will likely generate disparate impact via proxy variables. Proxy variables are strongly correlated with race and gender: for example, neighborhood location often correlates strongly with race due to a history of redlining in the United States, and purchase of feminine hygiene products is strongly correlated with gender. Per Solon Barocas, "even in situations where data miners are extremely careful, they can still effect discriminatory results with models that, quite unintentionally, pick out proxy variables for protected classes."[39] Amazon's recruiting algorithm, for example, was formally gender-blind but nevertheless learned to downgrade women's resumes because it was trained on biased data.

Given the failure of its color-blind theory of antidiscrimination, the standard economic view might turn to a second line of argument. <u>Second</u>, the standard economic view might argue that price discrimination is acceptable as long as it charges lower prices to groups historically discriminated against. (This argument, recall, was floated by policymakers in the White House report, when they theorized that price discrimination might benefit "historically disadvantaged groups."[40]).

This second argument provides a different theory of antidiscrimination, one which might be termed a "color-conscious" theory that attempts to remedy historical disadvantage. But this color-conscious theory also immediately runs into difficulties. For example, an algorithm might be accused of "reverse discrimination" if it charged higher prices to white men compared to other groups. And even when customers benefit from lower prices, they might feel singled out along race or gender lines: customers might feel grateful that the algorithm helped them (by charging them a lower price), yet simultaneously feel uncomfortably singled out on the basis of their group membership. A "color-conscious" pricing algorithm might help some customers financially but also harm their dignity, simultaneously.

Moreover, the standard economic view has not justified *why* its color-conscious theory is the right theory of antidiscrimination for pricing. If we are trying to remedy historical injustice, wouldn't a government program, or a tax scheme, be a better way to go? It seems

---

[39] Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CAL. L. REV. 671, 675 (2016).
[40] *See* supra note 31.

odd to shoehorn antidiscrimination goals into a pricing scheme. Markets usually treat consumers impersonally, as Max Weber and many others have observed,[41] so the color-conscious theory seems an odd fit for market norms.

The point is, any company that uses algorithmic price discrimination will have to justify why it is charging different prices to consumers in different race and gender groups. And the standard economic view has not provided an adequate theory of antidiscrimination. This difficulty may be insurmountable for the standard economic view.

* * *

Is algorithmic price discrimination the kind of problem that requires legal intervention? I am inclined to say yes—as we saw in the prior discussion, algorithmic price discrimination raises concerns about algorithmic nudging, privacy, antidiscrimination, and the monopoly power of large tech companies. These concerns would be difficult for any given consumer to address on her own, via a boycott or via her individual buying choices, so a collective response from the government seems warranted.

It is not the task of this paper, however, to propose a specific legal intervention. Instead, this paper develops the underlying framework that would justify and guide any such legal intervention. Below, the paper proposes and defends the *egalitarian view* on price discrimination.

II. The Egalitarian View

This section discusses the egalitarian view, which includes two provisos that at times allow sellers to charge differential prices. After laying the egalitarian view, this section discusses various objections and proposes a corollary to the egalitarian view.

According to the egalitarian view:

> Egalitarian View: Sellers should not charge different prices to different buyers, for identical goods and services, *absent a compelling justifying reason*.

Compelling justifying reasons fall into two broad categories:

1) Concern for a group that is especially needy or deserving, or that is perceived as such;
2) Differential costs to the seller, which the seller passes along to buyers.

Note that the first proviso focuses on the buyer (the buyer is especially needy or deserving), while the second proviso centers on the

---

[41] Max Weber, *Economy and Society* (1922).

seller (the seller incurs differential costs). Below, we consider each of the provisos in turn.

1. Concern for a group that is especially needy or deserving

To understand the first proviso—concern for a group that is especially needy, or that is perceived as such—we can look to discount programs such as student discounts, senior discounts, or veterans' discounts.

Many sellers routinely offer discounts to seniors, students, and veterans. These discounts are a form of price personalization, because buyers can access a different price based on their personal characteristics—their age, student status, or veterans' status. Other sellers offer discounts to people who clip coupons, or to buyers who are willing to wait until items go on sale. These discount programs tend to increase efficiency because they allow markets to clear: for example, movie theater seats that might otherwise go unfilled can be sold to a senior or student at a discount.

An egalitarian might justify discount programs on prioritarian grounds, because discounts benefit groups that are on average needier or more vulnerable than the general population. Buyers who clip coupons, or who wait until items go on sale, are often needier than other buyers. Seniors and students often have small or fixed incomes, and veterans may also be perceived as more needy or deserving of government aid if they have just returned from war. (The GI Bill, for example, offered college loans and housing loans to veterans returning from World War II.) While the neediness of the group may not hold for every individual member—for example, some seniors are wealthy and some veterans have already successfully reintegrated into high-paying jobs—an egalitarian would argue that the groups on average merit special concern.

Yet discount programs are not uncontroversial. Some authors have criticized coupons for creating a tax on the time of the poor. Discounts may also push prices higher for the rest of the public, because in lowering costs for the discounted group, the seller may be forced to raise costs for everyone else. Finally, some discount programs may generate race or gender bias: more women than men are seniors, for example, while the majority of veterans are men, so discount programs may generate disparate impact along race or gender lines.

An egalitarian could acknowledge the drawbacks of discount programs, while maintaining that discounts are nevertheless justified because society as a whole is well-placed to shoulder the burdens. Yes, discount programs might generate slightly higher costs for other customers—but the increased costs are relatively small, plus the costs are diffused across society as a whole. And while discounts may indeed create disparate impact, the disparate impact seems less pernicious than the algorithmic disparate impact discussed earlier in this paper. Discounts use only one variable (age, student status, veterans' status) and follow fixed rules, unlike pricing algorithms which use many variables and actively "learn" from their data. As a result, pricing algorithms can generate disparate impact in surprising and hard-to-predict ways—a worry which is less salient for discounts.

Objections:

A critic might argue that the first proviso arbitrarily singles out groups rather than individuals. According to the first proviso, a "compelling justifying reason" for differential pricing is "concern for a group that is especially needy."[42] But if egalitarians are concerned about needy *groups*, then why wouldn't they also be concerned about needy *individuals*? Per this critique, broad-grained discount programs should not be restricted to groups, but should instead become more and more fine-grained to offer discounts to specific needy individuals. Eventually, group-based discounts collapse into individualized, need-based discounts, and we end up at need-based price personalization for individual buyers.

The groups vs. individuals critique illuminates a tension within the egalitarian view. On the one hand, egalitarians want equal prices for everyone. On the other hand, egalitarians may want to express concern for specific needy groups, so they offer group-based discounts. These two principles are in tension, because the first principle requires the same price for everyone (including members of needy groups) while the second principle requires discounts for needy groups. As the critic points out, the egalitarian's concern for needy *groups* also seems to apply just as strongly to needy *individuals*, creating a slippery slope between group-based discounts and individualized discounts.

An egalitarian might respond to this critique in two ways. First, the egalitarian can draw a line between groups and individuals by invoking privacy. Second, the egalitarian can bite the bullet, and capitulate in theory but not in practice. Below, I discuss each response.

First, drawing a line between individuals and groups. While this line might be messy and not entirely satisfying, the egalitarian can defend group discounts, while rejecting individualized discounts, by appealing to privacy.

[42] Note that the "concern" in the first proviso need not refer to actual concern felt by the seller toward the buyer. A seller can be a red-blooded capitalist who is motivated only by profit, not by concern, yet the seller may nevertheless offer senior or student discounts because discounts expand his market and boost his profits. This is still an "acceptable justifying reason" per the first proviso, because the "concern" in the first proviso refers to an *egalitarian's* concern for more vulnerable groups, not to the *seller's* concern.

Group discounts tend to have a minimal impact on privacy. Most group discounts use only one variable, such as age, student status, or veterans' status, so they collect a relatively small amount of personal information. By contrast, pricing algorithms may collect hundreds or thousands of variables, creating a much greater impact on privacy.

More significantly, group discounts allow buyers to *affirmatively* disclose their personal information. Unlike pricing algorithms, which gather information while users passively browse the web, group discounts require an active disclosure, because the buyer must provide his driver's license, student identification card, or other document that proves his membership in the group. Because the disclosure is active and affirmative, group discounts allow buyers to preserve their privacy: the senior who wishes to avoid senior discounts, for whatever reason, can keep his age to himself.

An egalitarian can acknowledge that broad-grained group discounts are less precise than fine-grained individualized discounts (recall that some seniors and students are wealthy, and not all veterans are needy), while still defending a line between needy individuals and needy groups. Admittedly, the egalitarian's line-drawing is messy and not wholly satisfying. How big or small can the needy group be? Do customers have to be able to self-select into the group? What exactly is the connection between egalitarianism and privacy?[43] These questions may complicate the egalitarian's first proviso, but that might not stop the egalitarian from drawing a line between group-based discounts and individual discounts.

Alternatively, the egalitarian might be dissatisfied with the messy line-drawing required, and instead decide to bite the bullet. Below, we discuss the second option.

<u>Second,</u> biting the bullet. The egalitarian might concede that yes, individualized, need-based pricing would indeed satisfy the egalitarian view, at least in theory. But in biting the bullet in theory, the egalitarian need not concede much in practice, especially with respect to algorithmic price personalization by for-profit companies. As we discussed earlier, algorithmic price personalization is not need-based, and algorithms sometimes exploit lower-income consumers, consumers with fewer options, or consumers who are particularly naive or susceptible to algorithmic nudging. Companies have little incentive to stop this exploitation, because companies will tend to

[43] This last question—about the connection between egalitarianism and privacy—is particularly intriguing and may be a subject for future research. While egalitarians have not traditionally been much concerned with privacy, one side-effect of egalitarianism (at least in the market) is that it tends to protect consumer privacy: because egalitarians treat all customers the same, there is no need for customers to provide personal information that would qualify them for differential treatment.

choose algorithms that generate the highest profit or revenue, not algorithms that prioritize the needy.

So in biting the bullet, the egalitarian in fact concedes very little. In a fantasyland of perfectly beneficent tech companies, with need-based algorithms that prioritized low-income or low-wealth customers, then price discrimination might indeed satisfy the egalitarian view. (Although even need-based pricing algorithms would still raise privacy concerns, given the large amount of information they gather). But this fantasyland does not exist, and pricing algorithms in practice do not satisfy the egalitarian view.

2. Differential costs to the seller

Next, we discuss the second proviso. According to the second proviso, a "compelling justifying reason" for differential pricing is "differential costs to the seller, which the seller passes along to buyers." Below, we explore the second proviso and some objections.

The second proviso describes scenarios where the seller incurs higher costs for some buyers compared to other buyers. The archetypical example is an expensive location vs. a cheaper location: if a seller has a physical storefront in an expensive city and a cheaper city, she typically charges higher prices in the more expensive location where she pays higher rent. This kind of differential pricing is permitted under the second proviso, because the seller incurs differential costs across different buyers, and then passes along those costs to buyers.

The buyer's location, however, is a fairly broad-grained variable. Could a seller incur differential costs in a more fine-grained manner, across specific buyers? And would fine-grained differential costs put pressure on the egalitarian view?

In fact, fine-grained differential costs are common in *risk-based pricing*, which typically occurs in insurance markets. In risk-based pricing, sellers charge higher prices to borrowers deemed more "risky." Health insurance plans, for example, charge higher prices to older people than to younger people, and mortgage loans charge higher rates to applicants with bad credit scores. Risk-based pricing involves fine-grained differential costs across buyers, unlike location-based pricing where the differential costs of serving different customers are typically broad-grained across buyers.

Risk-based pricing presents an interesting test case for the egalitarian view, because it does indeed put pressure on the second proviso. Under the terms of the second proviso, sellers are allowed to charge different prices to different buyers when some buyers "cost" more than others. But in fine-grained scenarios, the second proviso opens the door to morally troubling pricing practices. Health insurance companies could, for example, charge higher prices to sick applicants who are more costly, or refuse to insure patients with expensive pre-existing conditions. (These practices were common in the United States until they were banned by the Affordable Care Act).

In addition to enabling problematic health insurance practices, the second proviso might also open the door to racial discrimination. In U.S. housing markets, for example, minority buyers are often "riskier" than non-minority buyers, because redlining and historical patterns of discrimination create higher risks of default in many minority neighborhoods. Given the racial disparities in the underlying risk distribution, risk-based pricing may hurt minority applicants who are more "costly" to lenders.

In the face of these challenges, the egalitarian can propose a modification to the second proviso:

Corollary to the Second Proviso: Sellers who incur differential costs for different buyers may pass along those costs to buyers, except when it would be unfair to do so.

To determine when it would be "unfair" to pass along costs to riskier buyers, the egalitarian might consider the buyer's vulnerability, the buyer's ability to bear the higher cost relative to other actors, and whether the buyer is responsible for creating her own risk. Health insurance companies, for example, often charge higher premiums to smokers than to nonsmokers, in part because smokers created a higher risk for themselves. On the other hand, health insurance companies are prohibited from raising rates when patients fall ill, because that is the time when patients become most vulnerable. Rather than concentrating the higher cost upon vulnerable, sick patients, the health insurance system spreads the cost among other actors who are better positioned to bear it: healthy patients subsidize the cost of health insurance for sicker patients, and health insurance companies can negotiate for government subsidies for insuring the most expensive patients. In this way, the health insurance system shifts costs away from vulnerable, sick patients and towards other actors who are better placed to bear the differential costs.

This framework for "unfair" may not satisfy every egalitarian. Some egalitarians may be more concerned about moral hazard, and they may fear that riskier individuals will continue to engage in high-risk behavior, knowing that insurance will pay the cost. Other egalitarians may be less concerned about moral hazard, and instead fear the costs of failing to extend universal coverage to everyone. The "fairness" evaluation may also depend upon the kind of behavior at issue: for example, smoking creates higher risk (and we generally allow health insurance companies to charge higher premiums to smokers), but pregnancy also creates higher risk (and we generally do not allow companies to raise premiums when someone becomes pregnant). All in all, the fairness assessment may be highly context-sensitive, and the egalitarian will have to draw some messy lines.[44]

In sum, the egalitarian view can add a corollary to the second proviso, to prevent sellers from passing along "unfair" differential costs to buyers. Costs might be unfair if they are based on risks the buyer did not create herself, if other actors are better placed to shoulder the extra costs, or if the buyer is uniquely vulnerable at the time she becomes "riskier." By adding the corollary, the egalitarian can block some of the more problematic instances of fine-grained differential pricing in risk-based markets.

*　　*　　*

At first glance, it may appear as though the standard economic view and the egalitarian view represent two competing values: efficiency (in the standard economic view) and equity (in the egalitarian view).

But as this discussion illustrates, algorithmic price discrimination is more complex than just efficiency vs. equity. The standard economic view may prize efficiency, but it also includes some implicit claims about distributional equity: when some economists claimed that price discrimination would help the poor, they were making a claim about distributional equity. Similarly, the egalitarian view centers equity but may also incorporate market-oriented values like efficiency: student discounts and senior discounts help needy groups while also promoting efficiency. And moving beyond just equity and efficiency, we find that price discrimination involves additional values like privacy, autonomy, and antidiscrimination.

If we compare the two qualified views, side by side, we see:

Qualified Standard Economic View: Sellers may charge different prices to different buyers, for identical goods and services;

[efficiency]

1) Except when nudging algorithms target consumer misperceptions;

[efficiency, autonomy, privacy]

2) Except when pricing algorithms create price gouging of the poor, sick, or otherwise vulnerable;

[equity]

[44] In the housing context, for example, the Fair Housing Act prohibits lenders from considering a borrower's race or residence in a minority neighborhood, even though lenders might be able to make better risk calculations if they could consider race-related variables. For a longer discussion, *see* Robert Bartlett, Adair Morse, Richard Stanton & Nancy Wallace, *Consumer-Lending Discrimination in the FinTech Era*, 143 JOURNAL OF FINANCIAL ECONOMICS 30, (2022).

3) Except when pricing algorithms discriminate along the lines of race, gender, ethnicity, or other protected characteristics.

[antidiscrimination]

Similarly, the qualified egalitarian view, too, contains a wide array of values:

> *Qualified* Egalitarian View: Sellers should not charge different prices to different buyers, for identical goods and services, absent a compelling justifying reason;
>
> **[equity]**
>
> "Compelling justifying reasons" are:
>
> 1) Concern for a group that is especially needy or deserving, or that is perceived as such; or
>
> **[efficiency, privacy, antidiscrimination]**
>
> 2) Differential costs to the seller, which the seller passes along to buyers, *except* when it would be unfair for sellers to pass along differential costs to buyers.
>
> **[efficiency, equity, antidiscrimination]**

Each qualified view contains a mix of values, among them efficiency, equity, privacy, autonomy, and antidiscrimination.

In balancing these values, the egalitarian view faces fewer challenges than the standard economic view. Even when qualified, the standard view has not really solved its antidiscrimination problem. When algorithms are formally race-blind, companies may be accused of generating disparate impact (so the color-blind theory likely fails). When algorithms are color-conscious, companies may be accused of reverse discrimination (so the color-conscious theory does not get them very far either). The egalitarian view, by contrast, largely sidesteps the antidiscrimination worry. In the egalitarian view, the default option is to charge the same price to everyone. While the egalitarian view allows some exceptions (such as discounts and risk-based pricing), these exceptions are cabined and tend to raise fewer antidiscrimination concerns.

Most significantly, the carve-outs in the qualified standard view run against companies' profit motives. Companies have little incentive to carve out algorithmic price gouging, algorithmic nudging, or other practices which might be quite profitable, because companies are typically oriented toward profit, not toward concerns about autonomy or privacy. Unless consumers actively object, or unless the government intervenes, companies may not want to invest time and resources to monitor their algorithms. And consumers often lack the empirics that would allow them to detect problematic instances of price gouging or algorithmic nudging: in the rideshare study, for example, the city of Chicago forced ride-share companies to disclose their data, which then allowed researchers to identify price gouging in Uber/Lyft rides. Yet empirics are far more difficult to obtain in other situations. Most companies do not disclose their data, so we have no way to know when price gouging occurs, or when to carve out exceptions.

Finally, the standard economic view has a privacy problem, insofar as it is problematic for pricing algorithms to use large amounts of personal data gathered from consumers.

All in all, the standard economic view faces some serious difficulties. While perfect price discrimination might have seemed like a great idea in theory, it is much less attractive in practice.

Conclusion

This paper assessed the standard economic view, which was developed in the 1980s and 1990s, and which has been used to support modern-day algorithmic price discrimination. After critiquing the standard economic view, this paper proposed and defended the egalitarian view, which opposes algorithmic price discrimination and favors equal prices for all.

While the egalitarian view is not perfect, it provides a workable theory that balances values like efficiency, equity, privacy, autonomy, and antidiscrimination. On the whole, we would do better to adopt the egalitarian view and to limit the use of price personalization algorithms.