informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–17
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Price Discrimination with Fairness Constraints

Maxime C. Cohen,[a] Adam N. Elmachtoub,[b,c] Xiao Lei[b,*]

[a] Desautels Faculty of Management, McGill University, Montreal, Quebec H3A 1G5, Canada; [b] Department of Industrial Engineering and Operations Research, Columbia University, New York, New York 10027; [c] Data Science Institute, Columbia University, New York, New York 10027
*Corresponding author

Contact: maxime.cohen@mcgill.ca, https://orcid.org/0000-0002-2474-3875 (MCC); adam@ieor.columbia.edu, https://orcid.org/0000-0003-0729-4999 (ANE); xl2625@columbia.edu, https://orcid.org/0000-0003-1421-749X (XL)

Received: June 27, 2020
Revised: December 13, 2020; August 31, 2021; September 26, 2021
Accepted: October 7, 2021
Published Online in Articles in Advance: March 18, 2022

https://doi.org/10.1287/mnsc.2022.4317

Copyright: © 2022 INFORMS

Abstract. Price discrimination strategies, which offer different prices to customers based on differences in their valuations, have become common practice. Although it allows sellers to increase their profits, it also raises several concerns in terms of fairness (e.g., by charging higher prices (or denying access) to protected minorities in case they have higher (or lower) valuations than the general population). This topic has received extensive attention from media, industry, and regulatory agencies. In this paper, we consider the problem of setting prices for different groups under fairness constraints. We first propose four definitions: fairness in price, demand, consumer surplus, and no-purchase valuation. We prove that satisfying more than one of these fairness constraints is impossible even under simple settings. We then analyze the pricing strategy of a profit-maximizing seller and the impact of imposing fairness on the seller's profit, consumer surplus, and social welfare. Under a linear demand model, we find that imposing a small amount of price fairness increases social welfare, whereas too much price fairness may result in a lower welfare relative to imposing no fairness. On the other hand, imposing fairness in demand or consumer surplus always decreases social welfare. Finally, no-purchase valuation fairness always increases social welfare. We observe similar patterns under several extensions and for other common demand models numerically. Our results and insights provide a first step in understanding the impact of imposing fairness in the context of discriminatory pricing.

History: Accepted by Jayashankar Swaminathan, operations management.
Funding: A. N. Elmachtoub was supported by the Division of Civil, Mechanical and Manufacturing Innovation [Grants 1763000 and 1944428].
Supplemental Material: The data files and online appendix are available at https://doi.org/10.1287/mnsc.2022.4317.

Keywords: fairness • price discrimination • personalization • social welfare

1. Introduction

The increased availability of consumer data in conjunction with the widespread use of e-commerce has led to a proliferation in discriminatory and personalized pricing strategies both in practice (Xue et al. 2015, Ye et al. 2018) and academia (Gallego and Topaloglu 2019, Elmachtoub et al. 2021). Specifically, companies often try to engage in first- or third-degree price discrimination tactics by leveraging the available data on their consumers, such as past purchase behavior, browsing history, and personal attributes, to predict consumer valuations. Although the practice is generally widespread, discriminatory pricing can result in disparate impact against protected groups. Protected groups may have higher (or lower) valuations for a product because of historical disadvantages or unobserved factors, and such differences can result in higher prices (or more limited access) for protected groups. Charles et al. (2008) show that Black individuals receive higher interest rates for auto loans, whereas Alesina et al. (2013) show that women receive higher interest rates for small business loans even after controlling for other consumer features. Fang and Munneke (2020) and Bartlett et al. (2022) show that Black/Latinx and women borrowers pay higher interest rates for mortgage loans while controlling for all possible factors, including risk. In fact, the study in Bartlett et al. (2022) even shows that this discrimination exists for FinTech lenders that make decisions based on AI algorithms. Larson et al. (2015) show that a test preparation provider charged Asian Americans higher prices, even when controlling for income. In all these examples, the number of transactions is at a very large scale. Thus, it suggests that the seller's motivation to carry forward such statistical discrimination is driven by financial gain. In contrast, the U.S. Civil Rights Act of 1964 and the Equal Credit Opportunity Act of 1974 protect against discrimination based on protected attributes, such as race, color, religion, sex, and national origin as well as most recently, sexual orientation and gender identity.[1] In fact, in 2020 the state of New York banned gender-based price discrimination[2] to fight against the increasing trend of large retailers selling the same product at different prices by simply changing the packaging or the product color.[3] Ensuring fairness is a direct concern of the Federal Trade Commission (FTC).

> When we at the FTC evaluate an algorithm or other AI tool for illegal discrimination, we look at the inputs to the model—such as whether the model includes ethnically-based factors, or proxies for such-factors, such as census tract. But, regardless of the inputs, we review the outcomes. For example, does a model, in fact, discriminate on a prohibited basis? Does a facially neutral model have an illegal disparate impact on protected classes? Our economic analysis looks at outcomes, such as the price consumers pay for credit, to determine whether a model appears to have a disparate impact on people in a protected class. —Andrew Smith[4]

Recently, there has been a surge of interest in understanding how to make discriminatory pricing practices that are fair from business (Wallheimer 2018, Weinberger 2019), legal (Gerlick and Liozu 2020, Gillis 2020), and regulatory perspectives (White House 2015, Gee 2018). Moral, legal, and ethical obligations are prompting sellers and regulators to ensure that pricing practices do not unfairly discriminate against protected attributes. Although this general principle is universally accepted, no formal framework prior to this work exists to properly implement or assess the impact of such fairness measures in the context of pricing decisions. In fact, in a recent discussion paper by the United Kingdom's Financial Conduct Authority (Gee 2018), it clearly states the need for conducting research on fairness in pricing.

> [I]t is important that we consider the fairness of pricing in markets we regulate. It is also important to consider the harm that may be caused by particular types of pricing practice … However, fairness issues can often be more complicated and the right course of action for us may be less clear. (Gee 2018, p. 5)

In light of the growing interest on fairness in the context of pricing decisions, we consider the following research questions.

1. How can we model the fairness of decisions made in the context of price discrimination? Is it possible to impose several types of fairness simultaneously?

2. What is the impact of fairness constraints on the seller, customers, and society at large?

In this paper, we propose a formal framework for pricing with fairness, including several definitions of fairness and their potential impact on consumers, sellers, and society. In a first step toward the ambitious agenda of designing pricing strategies that are fair, we consider the simplest scenario of a single-product seller facing consumers who can be partitioned into two groups based on a single binary feature observable to the seller (we then consider the extension with more than two groups in Section 4). For each group, we assume that the seller knows the valuation distribution, which allows us to isolate the effect of fair decision making from the machine learning task of predicting valuations. The seller's goal is to maximize profit by optimally selecting a price for each group, potentially subject to a fairness constraint that may be self-imposed or explicitly enforced by laws and regulations. We highlight that our model assumes that price discrimination occurs because of the difference in valuation distributions of the customer groups and not because of inherent racism or biases of the seller.

In this paper, we propose four definitions of fairness based on several different contexts and motivations. More details and motivation are presented in Section 2.1.

- *Price fairness* enforces that the prices offered to the two groups are nearly equal and is the common focus of the studies referenced.
- *Demand fairness* enforces that the access to the product is as close as possible across groups, meaning that the prices should be set in a way that yields a similar market share for each group. For example, a local college may want to offer tuition loans or scholarships in such a way that each group has an equal probability of enrolling.
- *Surplus fairness* requires that the surplus of the average person in each group is similar. As is standard, surplus is defined as the consumer valuation minus the price paid, and it is zero if no purchase is made.
- *No-purchase valuation fairness* imposes that the average valuation of consumers who do not purchase the product is approximately the same for each group. In other words, the normalized value lost in each group from individuals who could not afford the product should be similar.

With our model and definitions in place, we first show that satisfying all four fairness goals simultaneously is impossible unless the mean valuations are the same for both groups. In fact, even achieving two fairness measures simultaneously cannot be done in simple settings. We then consider the impact of imposing each fairness criterion separately and identify conditions under which the consumer surplus and the social welfare increase or decrease. (Clearly, imposing fairness always results in profit loss for the seller because of the additional constraint.) Note that the impact of price discrimination on social welfare has been studied in economics (Robinson 1969, Varian 1985) but without explicitly considering fairness constraints. For instance, we show that when the demand for each group is linear or exponential in price, a small amount of price or no-purchase valuation fairness will increase social welfare, whereas demand or surplus fairness will decrease social welfare. We also fully characterize all scenarios under linear demand and show, for example, that imposing too much price fairness may lead to a strictly lower social welfare relative to having no fairness constraints. We first focus on the setting with two groups and a single (protected) feature. We then extend our findings to settings with more than two groups and to the case where a second unprotected consumer feature is observed. Finally, we showcase computationally the robustness of our findings for other common demand models, such as exponential, logistic, and log-log.

1.1. Summary and Implications of Our Research

For industry practitioners and policy makers, our paper offers the following takeaways.

a. We show that achieving all four fairness definitions simultaneously is impossible. In fact, even achieving two of these definitions simultaneously is impossible under standard demand models. Thus, one should focus on a single notion of fairness depending on the context.

b. Imposing fairness constraints may not necessarily increase social welfare. The welfare change depends on both the fairness definition and the level of fairness. For price fairness, a little fairness improves social welfare, but too much fairness may lead to a lower welfare relative to imposing no fairness. For demand or surplus fairness, imposing any level of fairness will decrease social welfare. Finally, no-purchase valuation fairness always increases social welfare.

1.2. Related Literature

The concept of fairness has been extensively studied in economics, operations management, and computer science. Broadly speaking, fairness can be modeled either (i) as a utility term that is dependent on a reference point or (ii) as an exogenous constraint that may be imposed by a social planner based on social justice. Our work adopts the second approach, but we still review the literature related to both approaches for completeness.

In the economics literature, fairness is typically modeled as a reference effect, which depends on either a perceived value based on historical information or unequal outcomes across groups of individuals. In such settings, fairness is motivated in light of social comparison. Fairness with respect to perceived value refers to the situation where the price of an item should be close to its "fair" value. More precisely, customers form a reference price (based on historical information), and the demand is affected when the seller sets a price that is far from the reference price. This concept was first proposed by Kahneman et al. (1986), where the authors empirically show that people perceive a price raise as unfair if the surge is driven by shifts in demand. Eyster et al. (2021) then study the pricing problem under this type of fairness. Models based on a reference price were extensively studied in the context of dynamic pricing (Popescu and Wu 2007, Cohen et al. 2020) and for the newsvendor problem (Baron et al. 2015). On the other hand, several papers consider fairness with respect to unequal outcomes across groups of individuals (e.g., race, age, gender). Rabin (1993) and Fehr and Schmidt (1999) are among the first to study game-theoretic models with fairness considerations. Rabin (1993) models fairness as an explicit intention and shows that a fairness equilibrium may be achieved only if the Nash equilibrium also satisfies additional fairness constraints. In Fehr and Schmidt (1999), the need for fairness is modeled as a disutility for any unequal outcome among players. Ho and Su (2009) consider ultimatum games with peer-induced fairness concerns. Using a similar setting, Cui et al. (2007) consider a contract design problem and find that cooperation may be achieved when the manufacturer and retailer are sensitive to unequal outcomes. Li and Jain (2016) study a duopoly market with behavior-based pricing and find that incorporating fairness may increase sellers' profit and decrease consumer surplus.

The second approach in the fairness literature models fairness as exogenous constraints in decision-making or classification problems. The fairness constraints are usually motivated by egalitarianism, where each group of people (or even each individual) should receive the same treatment, or by Rawlsian justice (Rawls 1971), where the social planner aims to make the least advantaged people better off. Decision making under fairness constraints has been seen in the context of stable matching (Sethuraman et al. 2006), transportation systems (Chen and Wang 2018), network design (Rahmattalabi et al. 2019), advertising (Bateni et al. 2016), and dynamic learning (Gupta and Kamble 2019). Levi et al. (2016) investigate conditions under which uniform government subsidies are optimal. Another stream of papers considers the trade-off between fairness and efficiency in resource allocation (Bertsimas et al. 2011, 2012; Hooker and Williams 2012; Donahue and Kleinberg 2020). In our paper, we do not consider resource constraints, and we investigate to what extent our fairness constraints can improve social welfare.

Research on fairness has also been increasing rapidly in the machine learning community, and fairness is also modeled as exogenous constraints. Earlier papers consider classification algorithms under various fairness constraints (Dwork et al. 2012, Hardt et al. 2016, Donini et al. 2018) or the trade-off between different fairness metrics (Chouldechova 2017, Kleinberg et al. 2017). Kallus et al. (2021) provide a framework for assessing fairness without observing the protected attribute. There has also been work on how to design fair policies using causal inference (Nabi et al. 2019, Kasy and Abebe 2020, Viviano and Bradic 2020). In fact, in the context of classification problems, several papers have tried to integrate social welfare into the loss function (Heidari et al. 2018, 2019; Hu and Chen 2018, 2020). Although the fairness definitions in our paper resemble those in the machine learning literature, we consider the problem from a different perspective. The pricing procedure usually includes two steps: valuation prediction and pricing decisions. Machine learning models mainly focus on the first step, and the idea is to distribute the prediction error in a fair manner so that the prediction is unbiased. Our paper assumes that the seller has unbiased and accurate information on customer segmentation and valuations and focuses on how to make fair pricing decisions given such information.

Finally, one can view uniform pricing as a revenue management problem with a (simple) fairness constraint. In this view, our paper contributes to the line of research that compares social welfare under a uniform pricing strategy (i.e., perfect price fairness) versus discriminatory pricing (i.e., no fairness) (see, e.g., Robinson 1969, Schmalensee 1981, Varian 1985). Our paper includes these two extreme cases but also, considers intermediate levels of fairness constraints as well as four different fairness definitions. This literature shows that allowing for price discrimination generally leads to a higher social welfare compared with uniform pricing (ultimately converging to the situation where the seller is able to extract the entire consumer surplus). Surprisingly, our results show that restricting price discrimination by imposing fairness constraints can sometimes increase social welfare. In fact, we identify cases where imposing intermediate levels of fairness results in a social welfare that is higher than both perfect fairness and no fairness scenarios.

2. Framework and Preliminary Results

We consider a single-period setting where a seller offers a single product, with marginal cost $c \geq 0$, to two groups of customers (we consider the extension with more than two groups in Section 4). The seller needs to select a price for each group with the goal of maximizing profit. Specifically, customers are categorized based on an observable binary feature $X \in \{0,1\}$, so that each group $i = 0, 1$ can be offered a different price p_i. In this context, the seller may want to constrain the pricing policy to ensure fairness across the two groups because of a need to either improve customer perception or abide by government regulations. For example, X can correspond to gender, race, operating system, age, or type of device. We let d_i denote the population size of each group i. We assume that customers from group i have valuations for the product denoted by the random variable $V_i \sim F_i(\cdot)$, where $F_i(\cdot)$ is a given cdf. Customers in group i buy the product only if their valuation is at least the offered price p_i. Thus, $\bar{F}_i(p_i) = \mathbb{P}(V_i \geq p_i)$ represents the market share of group i, and $d_i\bar{F}_i(p_i)$ corresponds to the total demand of group i. We assume that the seller has enough supply to fulfill all the demand.

The profit function for group i is then $R_i(p_i) = (p_i - c)d_i\bar{F}_i(p_i)$. The seller's goal is to select p_0 and p_1 to maximize $R_0(p_0) + R_1(p_1)$, potentially subject to some fairness constraints (see more details). We let $p_i^* = \arg\max_p R_i(p)$ denote the optimal price offered by the seller to group i under no fairness constraints: that is, the unconstrained optimal price. We capture consumer welfare by the average consumer surplus given by $S_i(p) = \mathbb{E}[(V_i - p_i)^+]$ (note that we focus on the normalized surplus to account for possible asymmetries in population sizes). We also consider the expected no-purchase valuation, $N_i(p) = \mathbb{E}[V_i \mid V_i < p_i]$, that corresponds to the average valuation of nonbuyers. Finally, the total welfare from group i, $W_i(p_i)$, can be written as the profit plus the consumer surplus: that is, $W_i(p_i) = R_i(p_i) + d_iS_i(p_i)$.

2.1. Fairness Definitions

In the context of pricing, we propose the four following measures of fairness, where smaller quantities imply fairer strategies.

a. Price fairness, which is measured by $|p_0 - p_1|$.

b. Demand fairness, which is measured by $|\bar{F}_0(p_0) - \bar{F}_1(p_1)|$.

c. Surplus fairness, which is measured by $|S_0(p_0) - S_1(p_1)|$.

d. No-purchase valuation fairness, which is measured by $|N_0(p_0) - N_1(p_1)|$.

We also propose a unitless quantity, $\alpha \in [0,1]$, to denote the fairness level. The case of $\alpha = 0$ corresponds to no fairness constraints (i.e., unconstrained discriminatory prices are used), and the case of $\alpha = 1$ corresponds to perfect fairness (i.e., the groups are treated equally with respect to the fairness measure). We emphasize that α is not a decision variable but rather, a parameter that is selected by the seller to meet internal goals or satisfy regulatory requirements. Formally, let $M_i(p_i)$ be the specific fairness measure of interest (price, demand, surplus, or no-purchase valuation) under price p_i, and let $|M_0(p_0^*) - M_1(p_1^*)|$ be the fairness gap under the optimal (unconstrained) pricing strategy. Then, a pricing strategy p_i for $i = 0, 1$ is α fair with respect to $M_i(\cdot)$ if $|M_0(p_0) - M_1(p_1)| \leq (1-\alpha)\,|M_0(p_0^*) - M_1(p_1^*)|$. Imposing a specific amount of fairness for each measure corresponds to selecting a value for α. Specifically, the pricing problem for the seller becomes

$$\begin{aligned}\mathcal{R}(\alpha) := \max_{p_0,p_1\geq 0}\ & R_0(p_0) + R_1(p_1)\\ \text{s.t.}\ & |M_0(p_0) - M_1(p_1)| \leq (1-\alpha)\,|M_0(p_0^*) - M_1(p_1^*)|,\end{aligned} \tag{1}$$

where $\mathcal{R}(\alpha)$ denotes the optimal total profit as a function of the fairness level α. For convenience, we denote

$p_0(\alpha)$ and $p_1(\alpha)$ the optimal prices obtained by solving Problem (1) as a function of the fairness level α. Thus, $\mathcal{R}(\alpha) = R_0(p_0(\alpha)) + R_1(p_1(\alpha))$. We note that $p_i(\alpha)$ may sometimes be less than c in order to meet the fairness constraints. We define $\mathcal{S}(\alpha)$ as the total consumer surplus under the optimal prices with the α-fairness constraint (i.e., $\mathcal{S}(\alpha) = d_0 S_1(p_0(\alpha)) + d_1 S_1(p_1(\alpha))$). Also, we let $\mathcal{W}(\alpha) = \mathcal{R}(\alpha) + \mathcal{S}(\alpha)$ be the social welfare as a function of α.

The fairness definitions are motivated from practical and regulatory considerations. Price fairness is directly motivated by regulations and laws that proscribe price discrimination based on specific attributes, such as the U.S. Civil Rights Act and Equal Credit Opportunity Act. In fact, the U.S. Department of Housing and Urban Development makes it illegal to "impose different terms or conditions on a mortgage loan, such as different interest rates, points, or fees on the basis of race, color, national origin, religion, sex, familial status, or disability."[5] In October 2020, New York state banned gender-based price discrimination after observing that many products and services in brick-and-mortar locations were being sold at different sticker prices for men and women. The idea of imposing a price fairness constraint is mentioned directly by the UK Financial Conduct Authority (Gee 2018) via "relative price caps" that "impose limits on the differences in prices firms can charge to new and longstanding consumer groups" as an option to alleviate unfair pricing in financial services. As we mentioned in Section 1, several studies have found violations of price fairness, even after controlling for all relevant consumer features. In fact, Bartlett et al. (2022) even show that such price discrimination exists when decisions are made by AI algorithms, and we noted that the FTC also explicitly protects against algorithmic bias. All the aforementioned examples are occurring at a fairly large scale by sizeable lenders and retailers, so these practices cannot just be explained by inherent racism or sexism. A fundamental possibility is that the groups of consumers who receive higher prices have a higher valuation on average relative to the other groups. This phenomenon can occur for several reasons. First, it is well recognized that there exist significant gender and racial differences in preferences in terms of products' colors (Madden et al. 2000). Similarly, Byun and Park (2012) show that East Asian Americans are 1.5 times more likely to purchase commercial test preparation services, which indicates that this group tends to have a higher valuation for such services. Second, this phenomenon can occur when a group of people has a higher average search cost or less bargaining intensity (Fang and Munneke 2020), so they are willing to accept higher prices to reduce the searching process (e.g., applying for a new loan). Third, a higher valuation can also occur when a specific group is unable or less likely to know the competitors' prices (Bartlett et al. 2022), which can happen when a group is more likely to be located in a financial desert or is less likely to be eligible for a loan. Fourth, another potential reason for different valuation distributions is the difference in financial literacy across groups (Gillis 2020). We note that many of these factors can potentially be connected to systemic racism and sexism, although this is beyond the scope of our paper.

Demand fairness is well motivated by applications in education and healthcare. For instance, a local college may want to charge tuition in a way such that it ensures a well-represented population of students (i.e., giving an equal opportunity to students coming from all backgrounds and income levels). In the same vein, a healthcare service provider or an insurance company may want to set prices so that every group has an equal chance of affording proper care. It is common for pharmaceutical companies to charge different prices in different countries (depending on the median income). In these types of settings, demand fairness ensures that access to essential products and needs is offered equally among all groups of customers.

Imposing surplus fairness requires the difference in normalized surplus to be small, so that individuals from different groups are similarly satisfied. Consumer surplus is perhaps the most widely used notion in economics and operations management to measure the well-being of customers in the context of retailing (see, e.g., Brynjolfsson et al. 2003, 2019; Cohen et al. 2016). The concept of equal surplus (agents' welfare) is one of the most fundamental principles in economics research (see, e.g., Dworkin 1981, Arneson 1989) and has been extensively studied in resource allocation (Brams and Taylor 1996) and cooperative game theory (Hu et al. 2018). Given the importance of surplus management and the popularity of equal surplus in several economics applications, it is natural to design pricing policies that ensure that the consumer surplus (which can be seen as a proxy for happiness or satisfaction) in each group is relatively similar. Our definition of surplus fairness can be seen in Marcoux (2006), where the author argues that "a unitary price [equal prices] affords unequal degrees of utility enhancement [unequal surplus] to buyers."

Finally, we discuss no-purchase valuation fairness. When defining fairness measures in the context of pricing, it is important to also consider the customers who could not afford the product (because their price exceeds their valuations). Indeed, the nonbuyers are directly affected by discriminatory pricing policies. For example, individuals who need a loan the most may be offered a higher interest rate from banking institutions, which further prevents these individuals from accessing the service. The nonbuyers from one group may feel particularly discriminated against if their willingness-to-pay is higher than the nonbuyers of the other groups, which may lead to potential complaints or lawsuits. The prices offered to each group control the number of nonbuyers as well as how much the average nonbuyer was willing to pay. The report by the Financial Conduct Authority (FCA 2019) mentions that the significance of the harm caused by unfair pricing is not only measured by *how many* people are harmed but also, by *how much* the individual level of harm is: "if a small minority of consumers are affected, but we find that these consumers are a particularly vulnerable group of consumers and the level of individual harm is severe, we would likely be more concerned about the fairness of the pricing practice." Because the utility of nonbuyers in each group is zero, it is thus natural to measure the average level of individual harm among a group by looking at their valuation for the product. Although demand fairness accounts for the fraction of people who cannot afford the product, no-purchase valuation fairness measures the average individual level of harm within nonbuyers. No-purchase valuation fairness aims to ensure that one group of nonbuyers was not more dissatisfied than the other by measuring how much the groups were willing to spend. As we show later, no-purchase valuation fairness tends to provide the largest increase in social welfare (see Section 3.4 for a detailed discussion).

We note that our fairness definitions can be also connected to Rawls' principles of justice (Rawls 1999). In particular, demand fairness can be thought of as a reflection of the equal opportunity principle, whereas price, surplus, and no-purchase valuation fairness can be seen as a reflection of the difference principle (in which any economic inequalities should benefit the least advantaged individuals).

In this paper, we characterize the pricing strategy of a profit-maximizing seller that needs to comply with such fairness constraints. We also discuss the resulting impact on consumer surplus and social welfare. Note that one can come up with alternative fairness definitions beyond the ones we proposed. However, as we show in Section 2.2, our four definitions do not have redundancies in the sense that it is impossible to satisfy all of them perfectly at once. In fact, satisfying any pair of fairness measures perfectly is often not possible.

2.2. Impossibility Results

In an ideal world, regulators would impose perfect fairness (i.e., $\alpha = 1$) along all four definitions, so that customers across both groups will experience the same price, demand, surplus, and no-purchase valuation. The following theorem states that imposing 1-fairness across all four definitions simultaneously requires the necessary (and insufficient) condition that both groups have the same mean valuation. Such an assumption is very restrictive in practice, as different groups often have a different mean valuation. Impossibility results have been shown in the context of fairness for machine learning algorithms (Chouldechova 2017, Kleinberg et al. 2017) but under a setting related to misclassification errors rather than prescriptive pricing.

Theorem 1 (Impossibility of Perfect Fairness). *If* $\mathbb{E}[V_0] \neq \mathbb{E}[V_1]$, *then it is impossible to achieve* 1*-fairness in price, demand, surplus, and no-purchase valuation all simultaneously.*

Proof. Suppose for the sake of contradiction that there exists a pricing strategy that is 1-fair in price, demand, surplus, and no-purchase valuation. 1-fairness in price implies that there exists a price p such that $p_0 = p_1 = p$. 1-fairness in demand implies that $\mathbb{P}(V_0 \geq p) = \mathbb{P}(V_1 \geq p)$. Satisfying 1-fairness in surplus and no-purchase valuation implies that $\mathbb{E}[(V_0 - p)^+] = \mathbb{E}[(V_1 - p)^+]$ and $\mathbb{E}[V_0 \mid V_0 < p] = \mathbb{E}[V_1 \mid V_1 < p]$. By the law of total expectation (combined with adding and subtracting p to one of the conditional expectations), $\mathbb{E}[V_i] = \mathbb{E}[V_i \mid V_i < p]\mathbb{P}(V_i < p) + \mathbb{E}[(V_i - p)^+] + p\mathbb{P}(V_i \geq p)$, and thus, $\mathbb{E}[V_0] = \mathbb{E}[V_1]$, which contradicts our assumption. □

In fact, under common demand models such as linear and exponential, even achieving 1-fairness in two metrics simultaneously is difficult. Specifically, we show that for exponential demand, any pair of 1-fairness constraints cannot coexist unless the price is trivially set to zero. For linear demand, only 1-fairness in price and no-purchase valuation can be achieved simultaneously.

Proposition 1 (Impossibility for Linear and Exponential Demand). *Assume that the demand is as follows.*

a. *Exponential: that is,* $V_i \sim Exp(\lambda_i)$ *with* $\lambda_0 \neq \lambda_1$. *Then, any pair of 1-fairness in price, demand, surplus, and no-purchase valuation cannot coexist under positive prices.*

b. *Linear: that is,* $V_i \sim U(0, b_i)$ *with* $b_0 \neq b_1$. *Then, only 1-fairness in price and no-purchase valuation may coexist, and any other pair of 1-fairness in price, demand, surplus, and no-purchase valuation cannot coexist under positive prices and positive demand.*

Note that even when the mean valuations are equal (i.e., $\mathbb{E}[V_0] = \mathbb{E}[V_1]$), it is also readily possible that satisfying all the 1-fairness constraints simultaneously is impossible unless the prices are trivially set to zero. We illustrate such a case in the following example. Specifically, in Example 1, we provide an example where $\mathbb{E}[V_0] = \mathbb{E}[V_1]$, and only 1-fairness in price and demand can be satisfied simultaneously with positive prices (any other pair of fairness constraints cannot coexist).

Example 1 (Impossibility When Mean Valuations Are Equal). Suppose that $V_0 \sim U(0, 2)$ and $V_1 \sim Exp(1)$.

We find that 1-fairness in price and demand can be simultaneously satisfied when $p = 1.594$. However, for any price $p > 0$, we have $S_0(p) < S_1(p)$ and $N_0(p) > N_1(p)$, so that 1-fairness in price cannot coexist with either 1-fairness in surplus or in no-purchase valuation. Suppose that we have 1-fairness in demand, and let $q \in (0,1)$ be the market share for each group. (Note that $q > 0$ because group 1 follows an exponential demand and that $q < 1$ because $p > 0$.) We then have $p_0 = 2 - 2q$ and $p_1 = -\log q$. Therefore, we obtain $S_0(p_0) = q^2$ and $S_1(p_1) = q$, and thus, $S_0(p_0) < S_1(p_1)$ for any $q \in (0,1)$. Similarly, $N_0(p_0) = 1 - q$ and $N_1(p_1) = 1 + \frac{q\log q}{1-q}$, so that $N_0(p_0) > N_1(p_1)$ for any $q \in (0,1)$. As a result, 1-fairness in demand cannot coexist with either 1-fairness in surplus or in no-purchase valuation. Finally, under 1-fairness in surplus, we have $S_0(p_0) = (2-p_0)^2/4 = e^{-p_1} = S_1(p_1)$, implying that $p_0 = 2 - 2e^{-p_1/2}$. Consequently, $N_0(p_0) = 1 - e^{p_1/2}$ and $N_1(p_1) = 1 - \frac{p_1 e^{-p_1}}{1-e^{-p_1}}$. One can show that $N_0(p_0) < N_1(p_1)$ for any $p_1 > 0$, and thus, 1-fairness in no-purchase valuation is not possible. Hence, only 1-fairness in price and demand can be satisfied simultaneously with positive prices in this example.

In general, the discussion conveys that seeking fairness in multiple dimensions may not be feasible in most cases. Theorem 1 shows that achieving perfect fairness across all four definitions is impossible if the mean valuation of each group is different. Proposition 1 shows that satisfying two fairness definitions simultaneously is not possible even under simple demand models, and Example 1 shows that the same idea can be true (except for one combination) when the mean valuations are the same. These results prompt us to focus on the case where a company or a regulator considers the impact of imposing a single fairness constraint up to a certain level of α, which is easier to achieve. Specifically, we study the impact of fairness on the seller's profit, consumer surplus, and social welfare.

2.3. Imposing a Little Fairness

In this section, we consider imposing a small amount of fairness and examine whether it increases social welfare. Although it is clear that imposing fairness will decrease the seller's profit, we are interested in the impact on social welfare. One may naturally conjecture that one of the motivations behind imposing fairness in pricing is to increase social welfare.

Recall from Problem (1) that $\mathcal{R}(\alpha)$ is the total seller's profit under an α-fairness constraint (where the measure is clear from the context). Recall also that $\mathcal{S}(\alpha)$ is the total consumer surplus under the optimal prices with the α-fairness constraint (i.e., $\mathcal{S}(\alpha) = d_0 S_1(p_0(\alpha)) + d_1 S_1(p_1(\alpha))$), and $\mathcal{W}(\alpha) = \mathcal{R}(\alpha) + \mathcal{S}(\alpha)$ is the social welfare as a function of α.[6] Theorem 2 shows that the impact of imposing a small amount of fairness on social welfare crucially depends on the fairness definition. Mathematically, we are interested in cases where the (right) derivative of the social welfare at $\alpha = 0$, $\mathcal{W}'(0)$, is positive.[7] To gain analytical tractability, we consider two common demand models: linear and exponential.

Theorem 2 (Impact of Imposing a Little Fairness on Social Welfare). *Assume that the demand is either* (i) *linear (i.e., $V_i \sim U(0,b_i)$ with $b_0 \neq b_1$) or* (ii) *exponential (i.e., $V_i \sim Exp(\lambda_i)$ with $\lambda_0 \neq \lambda_1$). Then, $\mathcal{W}'(0) > 0$ under price or no-purchase valuation fairness, whereas $\mathcal{W}'(0) < 0$ under demand or surplus fairness.*

Theorem 2, proved in Online Appendix A, conveys that for linear or exponential demand, imposing a small amount of fairness in price or no-purchase valuation improves social welfare, whereas imposing a small amount of fairness in demand or surplus decreases social welfare. In fact, one can identify a general necessary condition under which $\mathcal{W}'(0) > 0$ for any demand function that leads to continuous and differentiable $R_i(\cdot)$, $S_i(\cdot)$, and $W_i(\cdot)$ at p_i^* (the exact condition does not provide any further insight and is thus omitted for conciseness; see Lemma A.1 in Online Appendix A for more details). Our result suggests that if a seller is keen on using price discrimination tactics, then it is possible that imposing a small amount of fairness ($\alpha > 0$) can increase social welfare compared with no fairness ($\alpha = 0$). This is a surprising complement to classic economics, which suggests that a seller relaxing the strategy from uniform pricing (one group) to discriminatory pricing (two groups with $\alpha = 0$) will increase social welfare.

To derive additional insights, we next focus on the case of uniform valuations (i.e., linear demand). We then test the robustness of our findings for three alternative demand models in Section 5.

3. Analysis for Linear Demand

In this section, we present a comprehensive analysis for the linear demand model. Specifically, we assume that $\bar{F}_i(p) = \max\{0, 1 - \frac{1}{b_i}p\}$ or equivalently, $V_i \sim U(0,b_i)$. Without loss of generality, we impose $c < b_0 < b_1$. The linear demand model is commonly used in various settings. Not only does linearity make our analysis tractable, it can also be viewed as a near-optimal approximation to more complex demand models (see, e.g., Besbes and Zeevi 2015, Cohen et al. 2021). We consider the case with two groups and study the impact of imposing each type of fairness. We then consider the case with N groups in Section 4.1 and incorporating an unprotected feature in Section 4.2. We consider nonlinear demand in Section 5.

Under linear demand, the market share, profit, and (normalized) consumer surplus for each group $i = 0, 1$ are given by $\bar{F}_i(p_i) = \max\{0, 1 - \frac{1}{b_i}p\}$, $R_i(p_i) = d_i(p_i - c)\bar{F}_i(p_i)$, $S_i(p_i) = \frac{(b_i - p_i)\bar{F}_i(p_i)}{2}$, and $N_i(p_i) = \frac{\min\{b_i, p_i\}}{2}$, respectively. It is well known that the optimal unconstrained price for each group is given by $p_i^* = (b_i + c)/2$. At p_i^*, the demand, consumer surplus, and no-purchase valuation for group i reduce to $\bar{F}_i(p_i^*) = (b_i - c)/2b_i$, $S_i(p_i^*) = (b_i - c)^2/8b_i$, and $N_i(p_i^*) = (b_i + c)/4$, respectively. Because $b_0 < b_1$, all of price, demand, surplus, and no-purchase valuations are lower for group 0 than for group 1. We naturally restrict the prices to be larger than zero, but they may be below the cost c (this captures the situation when it is optimal for the seller to earn a negative profit for one group in order to extract a high positive profit from the other group while enforcing fairness constraints). We next discuss the optimal pricing strategy and the potential impact of imposing each type of fairness constraint for a given α.

The price optimization problem with fairness constraints is not a straightforward extension of the nominal setting (i.e., without fairness constraints). Under linear demand, the profit function for each group is concave for $p \in [0, b_i]$. In our analysis, when imposing fairness constraints, we will show that one of the prices may reach the boundary 0 or b_i, thus potentially making Problem (1) nonconvex. In the left panels of Figure 1, we show an example of the price dynamics, $p_i(\alpha)$, under each of the four fairness constraints. Interestingly, the four fairness constraints lead to totally different pricing strategies. Further, for three of the four constraints (price, surplus, and no-purchase valuation), there are nonlinear price changes. Given that the price strategies vary across different fairness constraints, the impact on profit, consumer surplus, and social welfare is also different (see the right panels of Figure 1). We next provide closed-form expressions for the optimal prices as a function of α under each fairness measure, which allow us to assess the impact on the seller's profit, consumer surplus, and social welfare. All the proofs can be found in Online Appendix B.

3.1. Price Fairness

In this section, we consider imposing price fairness. As α starts to increase, $p_0(\alpha)$ increases, whereas $p_1(\alpha)$ decreases. Consequently, group 0 (group 1) is earning a lower (higher) surplus. Then, when α becomes large enough, it is possible that p_0 is set to be higher than b_0. This implies that it is optimal for the seller to "give up" group 0 (i.e., the demand from group 0 is zero). At this point, it is equivalent to simply set $p_0 = p_1 = p_1^*$. We formally characterize the resulting impact of imposing price fairness in Proposition 2.

Proposition 2. *Let* $\tilde{\alpha}_p = \min\left(\sqrt{\frac{d_0 b_1 + d_1 b_0}{d_1 b_0}\frac{b_0 - c}{b_1 - b_0}}, 1\right)$. *If* $0 \leq \alpha \leq \tilde{\alpha}_p$, *then both the consumer surplus and the social welfare increase with* α. *If* $\tilde{\alpha}_p < \alpha \leq 1$, *then* $p_0(\alpha) = p_1(\alpha) = p_1^*$. *In addition, the profit, consumer surplus, and social welfare are lower relative to the case without price fairness* (*i.e.,* $\alpha = 0$).

The impact of imposing price fairness admits two separate cases.

a. When $\alpha \leq \tilde{\alpha}_p$, the change in consumer surplus is a quadratic function that is increasing with α. The change in welfare is a quadratic function that is concavely increasing for any $\alpha \leq \tilde{\alpha}_p$. Thus, for any $\alpha \leq \tilde{\alpha}_p$ (i.e., before giving up group 0), imposing additional price fairness increases social welfare. Interestingly, both the gain in social welfare and the loss in profit are concave in α. This implies that the marginal effect of imposing additional price fairness on social welfare (profit) decreases (increases) with α. Consequently, imposing a small amount of price fairness yields the highest marginal benefit on social welfare coupled with the lowest profit loss. This insight can help persuade regulators that incorporating a small amount of price fairness is worthwhile.

b. When $\alpha > \tilde{\alpha}_p$, it becomes optimal for the seller to give up group 0 and set both prices at $p_1^* > b_0$. This is assuming that $\tilde{\alpha}_p < 1$ given that if $\tilde{\alpha}_p \geq 1$, the second case does not exist. Consequently, the profit and surplus from group 0 are lost, so that it leads to a lose-lose outcome (i.e., lower seller's profit and lower consumer surplus). In this case, the social welfare drops below $\mathcal{W}(0)$ for any $\alpha > \tilde{\alpha}_p$.

In the top right of Figure 1, we consider a concrete example and show how the profit, consumer surplus, and social welfare vary as a function of α under price fairness. An interesting implication of Proposition 2 is the fact that the social welfare reaches its maximum right before giving up group 0: that is, when $\alpha = \tilde{\alpha}_p$ (in the example of Figure 1, $\tilde{\alpha}_p = 0.22$). The seller's decision to give up group 0 crucially depends on the value of $\sqrt{\frac{d_0 b_1 + d_1 b_0}{d_1 b_0}\frac{b_0 - c}{b_1 - b_0}}$. If it is less than on, then any $\alpha > \tilde{\alpha}_p$ leads to a lose-lose outcome. The square root term, $\sqrt{\frac{d_0 b_1 + d_1 b_0}{d_1 b_0}}$, depends on the relationship between $d_0 b_1$ and $d_1 b_0$ or equivalently, d_1/d_0 and b_1/b_0. For example, when $d_1/d_0 = b_1/b_0$, then $\sqrt{\frac{d_0 b_1 + d_1 b_0}{d_1 b_0}} = \sqrt{2}$. On the other hand, when $d_1/d_0 = 100 b_1/b_0$, then $\sqrt{\frac{d_0 b_1 + d_1 b_0}{d_1 b_0}} = \sqrt{1.01}$. The higher d_1/d_0 is (for a fixed b_0/b_1), the more likely the seller will give up group 0 when imposing fairness (i.e., it will occur for a smaller value of α). This is consistent with the intuition that when the high-valuation group (group 1) dominates the market, the seller is more likely to give up the low-valuation group (group 0). Similarly, the term $(b_0 - c)/(b_1 - b_0)$ conveys that the higher the difference between both groups' valuations is, the more likely the seller is to give up group 0 when imposing fairness, which is also intuitive.

To summarize, imposing price fairness increases social welfare as long as α remains below $\tilde{\alpha}_p$. When the differences in population size and in valuation are significant, setting α to a large value may lead to a lose-lose outcome. Furthermore, the value of α needs to be carefully selected given that the maximum and minimum values of $\mathcal{W}(\alpha)$ are right beside each other.

Figure 1. Impact of Fairness Under Linear Demand

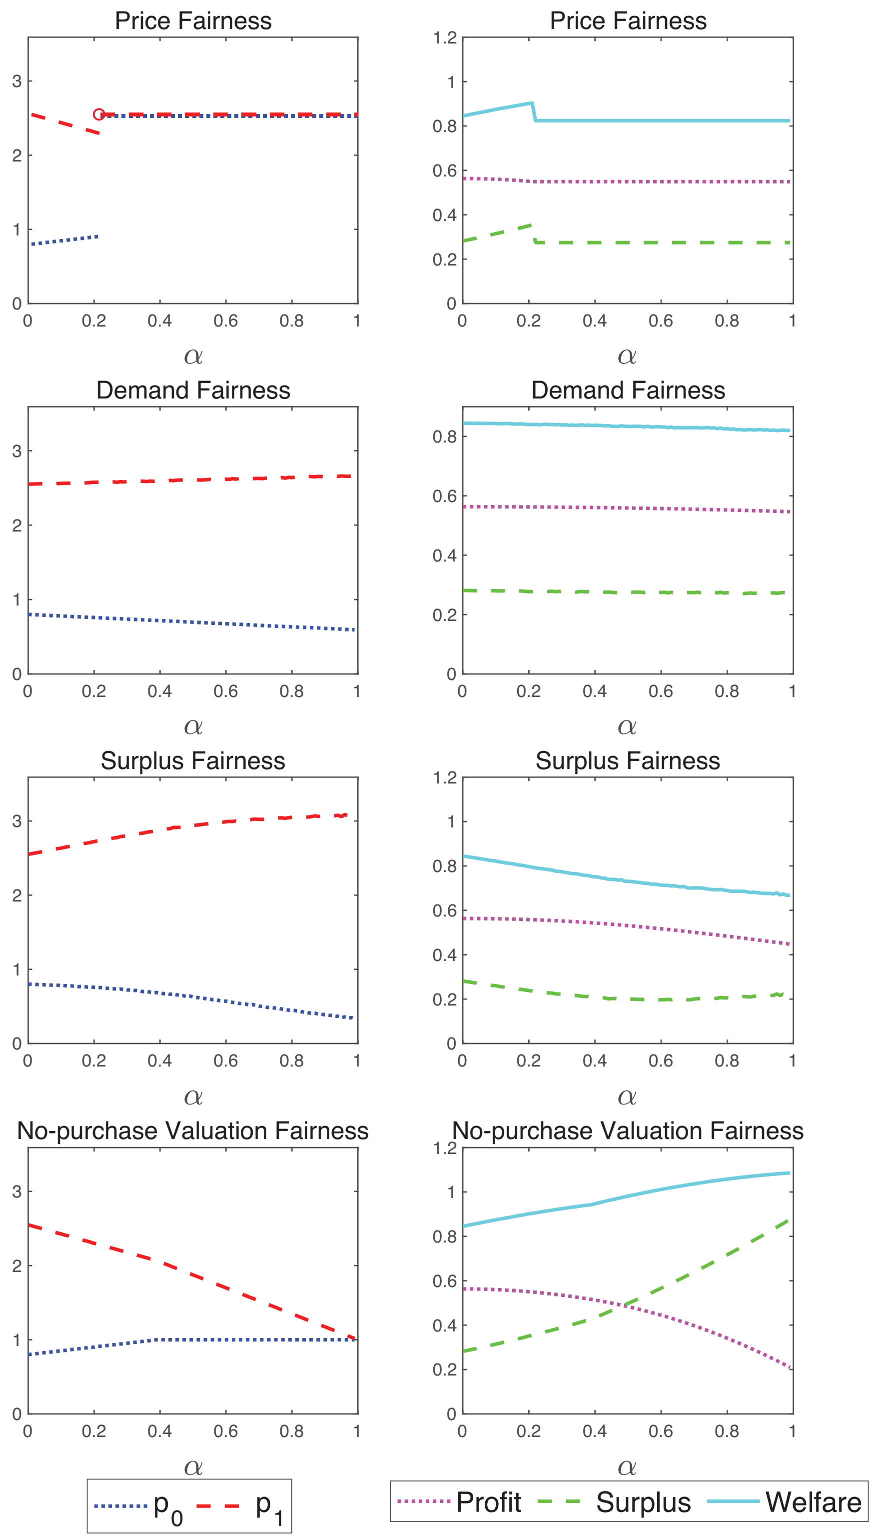

Note. Parameters: $d_0 = 0.35, d_1 = 0.65, b_0 = 1, b_1 = 4.5, c = 0.6$.

3.2. Demand Fairness

We next consider the case of demand fairness. Recall that $\bar{F}_i(p_i^*) = (b_i - c)/2b_i$, so that group 0 has lower demand. Thus, as α increases, $p_0(\alpha)$ decreases to raise demand from group 0, whereas $p_1(\alpha)$ increases to reduce demand from group 1. We characterize the impact of imposing demand fairness in Proposition 3.

Proposition 3. *For demand fairness, the profit, consumer surplus, and social welfare all decrease with* α.

For demand fairness, the change in surplus is always negative and reaches its minimum at $\alpha = 1$. Hence, the change in surplus is monotonically decreasing for $\alpha \in [0,1]$. Consequently, the change in social welfare is also monotonically decreasing, so that any degree of demand fairness reduces social welfare and leads to a lose-lose outcome.

3.3. Surplus Fairness

Recall that the surplus of group i is $S_i(p_i) = \frac{(b_i - p_i)(1 - p_i/b_i)}{2}$ and that $S_0(p_0^*) < S_1(p_1^*)$. Thus, as α starts to increase, $p_0(\alpha)$ decreases to raise the surplus from group 0, and $p_1(\alpha)$ increases to reduce the surplus from group 1. The closed-form expressions for surplus fairness are complicated because of the nonlinearity of the surplus function. However, as we show in Proposition 4, the social welfare is always below $\mathcal{W}(0)$ for any $\alpha > 0$.

Proposition 4. *For surplus fairness,* $\mathcal{W}(\alpha) < \mathcal{W}(0)$ *for any* $\alpha \in (0,1]$.

Hence, regardless of the value of α, imposing surplus fairness always leads to lower social welfare relative to no fairness constraint.

3.4. No-Purchase Valuation Fairness

For no-purchase valuation fairness under linear demand, we have $N_i(p_i) = p_i/2$. Therefore, for small values of α, we obtain the same pattern as for price fairness. However, when α becomes large, the price dynamics under no-purchase valuation fairness follow a different pattern. We formalize the impact of no-purchase valuation fairness in Proposition 5.

Proposition 5. *For no-purchase valuation fairness, both the consumer surplus and the social welfare increase with* α.

For no-purchase valuation fairness, the social welfare always increases with α. When α is small, the dynamics are the same as for price fairness. As in price fairness, both the gain in social welfare and the loss in profit are concave in α, so that imposing a small amount of fairness yields the highest marginal benefit on social welfare coupled with the lowest profit loss. When α is large, however, instead of setting $p_1 = p_0 = p_1^* > b_0$ (so that the demand from group 0 is zero), the seller has to lower p_1 to reduce the gap in the no-purchase valuation between both groups. Indeed, for any $p_0 \geq b_0$, the expected no-purchase valuation of group 0 is equal to $b_0/2$ and cannot be raised by increasing p_0. Thus, the only way to reduce the difference in no-purchase valuations is to decrease p_1, and hence, the social welfare continues to increase (because the social welfare is monotonically decreasing with price). As a result, imposing additional no-purchase valuation fairness always increases social welfare, even though group 0 may be given up (when p_0 is set at b_0). Interestingly, for a large value of α, the only fairness definition that yields a social welfare that is greater than $\mathcal{W}(0)$ is the no-purchase valuation fairness. As a result, no-purchase valuation fairness weakly dominates the other three fairness metrics in terms of social welfare.

3.5. Summary and Discussion

We next summarize the results derived so far.

a. Under price fairness, the social welfare increases with α and reaches its maximum at $\alpha = \tilde{\alpha}_p$. It then drops below $\mathcal{W}(0)$ for any $\alpha > \tilde{\alpha}_p$.

b. Under demand fairness, the social welfare decreases with α—leading to a lose-lose outcome.

c. Under surplus fairness, the social welfare is always below $\mathcal{W}(0)$—leading to a worse outcome relative to imposing no fairness.

d. Under no-purchase valuation fairness, the social welfare always increases with α, but it is possible that the demand of group 0 vanishes.

4. Extensions

In this section, we consider two extensions of our model and show the robustness of our findings from Section 3. In Section 4.1, we consider the case with $n > 2$ groups. In Section 4.2, we study the case where there is an additional feature Y that does not need to be protected.

4.1. Multiple Groups of Customers

We now assume that X is not binary and can take on N values. We index the groups by $0, \ldots, N-1$ and assume that group i has population d_i and parameter b_i. Without loss of generality, we assume that $b_0 \leq \ldots \leq b_{N-1}$. The profit maximization Problem (1), with linear demand, can be generalized as follows:

$$\begin{aligned}
\mathcal{R}(\alpha) := \max_{p_i \geq 0} \ & \sum_{i=0}^{N-1} d_i (p_i - c) \max\left\{0, 1 - \frac{p_i}{b_i}\right\} \\
\text{s.t.} \ & |M_i(p_i) - M_j(p_j)| \leq (1-\alpha) \max_{i,j \in \{0,\cdots,N-1\}} |M_i(p_i^*) - M_j(p_j^*)|, \\
& \forall i,j \in \{0,\cdots,N-1\}, \qquad (2)
\end{aligned}$$

where M_i is the fairness metric under consideration. Although Problem (2) is easy to solve numerically (as we will show in Lemma 2), its closed-form solution and the impact on social welfare are difficult to characterize. In particular, there are potentially many phase changes in the optimal solution as α varies. However, by leveraging the structural properties of the linear demand, we can still derive managerial insights that turn out to be similar to the two-group case studied in Section 3.

We start by investigating the cases of demand, surplus, and no-purchase valuation fairness. Recall that for the setting with two groups, Propositions 3–5 show that imposing fairness is either detrimental (for demand or surplus fairness) or always beneficial (for no-purchase valuation fairness) in terms of social welfare. We next extend these results to the multigroup case, as stated in Proposition 6.

Proposition 6. *Consider any $\alpha \in (0,1]$. Then, the following results hold.*

a. *For demand fairness, $\mathcal{W}(\alpha)$ decreases monotonically with α.*

b. *For surplus fairness, $\mathcal{W}(\alpha) < \mathcal{W}(0)$.*

c. *For no-purchase valuation fairness, $\mathcal{W}(\alpha)$ increases monotonically with α.*

The proof for each part of Proposition 6 relies on different arguments and machinery (see Online Appendix C). For demand and no-purchase valuation fairness, we first show that the prices $p_i(\alpha)$ are monotonic. We then show that the problem can be reduced to an instance with two groups. For surplus fairness, we also find a reduction to an instance of the problem with two groups, but in this case, the social welfare of the two-group instance does not match the social welfare of the N-group problem. Instead, we leverage convexity properties of several relevant functions to arrive at our desired result. Ultimately, Proposition 6 shows that our findings from Section 3 continue to hold for settings with any finite number of customer groups.

We next consider the case of price fairness. Recall that for the setting with two groups, Proposition 2 shows that for small values of α, the social welfare increases with α. However, when α becomes large, it may be optimal for the seller to give up a low-value group. In a setting with more than two groups, the impact of α on prices is more intricate relative to the setting with two groups. In Figure 2, we present an example with three groups. As α increases, the price changes (left panel) admit four linear pieces, and the social welfare function (right panel) includes two drops. Nevertheless, we can still partially characterize the impact on social welfare, as stated in Proposition 7 (the proof can be found in Online Appendix C).

Proposition 7. *For price fairness, we have the following.*

a. *$\mathcal{W}'(0) > 0$: that is, imposing a small amount of price fairness increases social welfare.*

b. *Suppose that all the groups have positive demand for all $\alpha \in [0,1]$ (i.e., $\bar{F}_i(p_i(\alpha)) > 0$); then, $\mathcal{W}(\alpha)$ increases monotonically with α.*

c. *If there exists an α' such that at least one group has zero demand, then $\mathcal{W}(\alpha')$ may be either higher or lower than $\mathcal{W}(0)$.*

When α is small enough, Proposition 7(a) suggests that a little price fairness still increases social welfare for any finite number of groups. As illustrated in Figure 2, when α is relatively large, some groups may be excluded by being offered high prices, thus leading to potentially complicated patterns. Nevertheless, when group exclusion does not happen, Proposition 7(b) conveys that the social welfare increases monotonically with α. On the other hand, when group exclusion does happen, the social welfare could either be higher or be lower than the unconstrained social welfare per Proposition 7(c), which is different from the two-group setting where group exclusion always leads to a lower welfare.

To summarize, for demand, surplus, or no-purchase valuation fairness, we find that our insights from the two-group setting generalize for multiple groups. For price fairness, even though the prices follow a more

Figure 2. Prices (Left Panel) and Profit, Surplus, and Welfare (Right Panel) Under Price Fairness

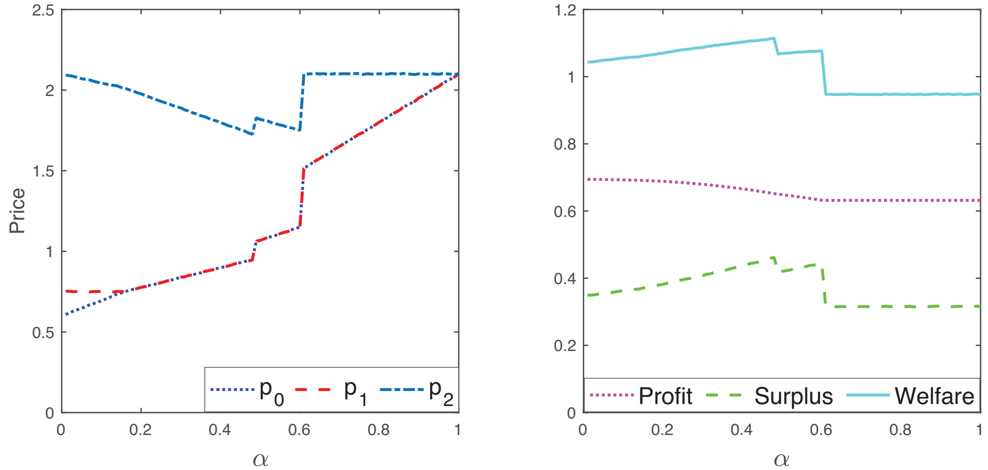

Note. Parameters: $d_1 = 0.1, d_2 = 0.2, d_3 = 0.7, b_0 = 1, b_1 = 1.3, b_2 = 4, c = 0.2$.

complex pattern, we still find that a little price fairness is always beneficial in terms of social welfare.

4.2. Adding an Unprotected Feature

In practice, it is possible that a subset of the features is unprotected, so that the seller is allowed to discriminate freely based on such features (e.g., loyalty status, purchase history, country). For simplicity, we consider the case of two observable features: a binary (protected) feature $X = \{0, 1\}$ on which we would like to impose fairness and a binary unprotected feature $Y = \{0, 1\}$. This gives rise to four groups of customers. We use the subscript xy to denote a specific group; for example, d_{00} is the population of group $(X = 0, Y = 0)$, and p_{10} is the price offered to group $(X = 1, Y = 0)$.

When adding an unprotected feature, our fairness definitions from Section 2 need to be revisited. We propose two refined versions for each fairness definition: conditional fairness and weighted average fairness. For example, consider the price fairness definition. The *conditional α-fairness* is defined such that for any value of Y, the price difference between the group with $X = 0$ and the group with $X = 1$ is small:

$$|p_{0y} - p_{1y}| \le (1 - \alpha)\,|p^*_{0y} - p^*_{1y}|\,,\quad \forall y = 0, 1.$$

The *weighted average α-fairness* is defined such that the weighted average prices (with respect to population sizes) for $X = 0$ and $X = 1$ are close together: that is,

$$|\bar{p}_0 - \bar{p}_1| \le (1 - \alpha)\,|\bar{p}^*_0 - \bar{p}^*_1|\,,$$

where $\bar{p}_i = \frac{d_{i0}p_{i0} + d_{i1}p_{i1}}{d_{i0} + d_{i1}}$ for $i = 0, 1$. The same refinements extend to the other fairness definitions.

For conditional fairness, the problem separates in two parallel subproblems for each value of Y. Each subproblem has two groups, so that the results from Sections 3.1–3.4 naturally apply. On the other hand, the weighted average fairness cannot be solved using the same approach. Specifically, the pricing problem faced by the seller becomes

$$\begin{aligned}\mathcal{R}(\alpha) := \max_{p_{00}, p_{01}, p_{10}, p_{11}} \quad & R_{00}(p_{00}) + R_{01}(p_{01}) + R_{10}(p_{10}) + R_{11}(p_{11}) \\ \text{s.t.} \quad & |\bar{M}_0 - \bar{M}_1| \le (1 - \alpha)\,|\bar{M}^{0*} - \bar{M}^*_1|\,, \end{aligned} \tag{3}$$

where $\mathcal{R}(\alpha)$ denotes the total optimal profit as a function of α, $\bar{M}_i = \frac{d_{i0}M_{i0}(p_{i0}) + d_{i1}M_{i1}(p_{i1})}{d_{i0} + d_{i1}}$ is the weighted average measure of group i with respect to Y, and $\bar{M}^*_i$ is the weighted average measure of group i under the optimal prices when the problem is unconstrained (i.e., no fairness constraints). For convenience, we denote $p_{xy}(\alpha)$ the optimal prices to Problem (3) as a function of α. Note that $\mathcal{R}(\alpha) = R_{00}(p_{00}(\alpha)) + R_{01}(p_{01}(\alpha)) + R_{10}(p_{10}(\alpha)) + R_{11}(p_{11}(\alpha))$. For simplicity of exposition, we focus on the situation where all the groups have positive prices and demand (i.e., $p_{xy}(\alpha) \in (0, b_{xy})$). Proposition 8 shows that our insights from Sections 3.1–3.4 still hold under weighted average fairness (the proof is in Online Appendix D).

Proposition 8. *Assume that the demand is linear so that the valuations for a group xy are uniform between zero and* b_{xy}*, where* $x, y \in \{0, 1\}$*. For all* α *such that* $p_{xy}(\alpha) \in (0, b_{xy})$ *and for any* $x, y \in \{0, 1\}$*, the following holds for weighted average α-fairness.*

a. *For price fairness,* $\mathcal{W}(\alpha)$ *increases with* α.
b. *For demand fairness,* $\mathcal{W}(\alpha)$ *decreases with* α.
c. *For surplus fairness,* $\mathcal{W}'(0) < 0$.
d. *For no-purchase valuation fairness,* $\mathcal{W}(\alpha)$ *increases with* α.

Proposition 8 shows that all the qualitative results from Sections 3.1–3.4 still hold for weighted average fairness (with the exception of surplus fairness for which we now have a slightly weaker claim). Specifically, for small values of α such that $p_{xy}(\alpha) \in (0, b_{xy})$, imposing additional price or no-purchase valuation fairness increases social welfare. On the other hand, imposing demand or surplus fairness has a negative impact on social welfare. Interestingly, conditional fairness and weighted average fairness may lead to different directions of price changes (even under the same fairness metric), but the impact on social welfare is similar. For example, if $d_{00} = 0.9, d_{01} = 0.1, d_{10} = 0.1, d_{11} = 0.9$ and $b_{00} = 1, b_{01} = 2, b_{10} = 4, b_{11} = 3$, then under conditional price fairness, p_{01} decreases and p_{11} increases as $b_{01} > b_{11}$. However, for weighted average price fairness, we have $\bar{p}_0 = 1.3$ and $\bar{p}_1 = 2.9$, so that $\bar{p}_0 < \bar{p}_1$. In this case, p_{01} increases and p_{11} decreases with α. Although the direction of price changes is different, surprisingly both types of price fairness increase social welfare.

Finally, we consider a scenario where the seller may not be able to (or may not want to) price discriminate based on the protected feature, implying that $p_{0y} = p_{1y}$. Meanwhile, the fairness is still measured based on the difference between $\bar{M}_0$ and $\bar{M}_1$, which is in general nonzero. In other words, the seller optimizes prices based only on the unprotected feature Y, whereas the fairness is imposed with respect to the protected feature X. This corresponds to solving Problem (3) with an added price constraint. In this case, one can easily verify that our result in Proposition 8(a) still holds for price fairness. However, for the other fairness definitions, the social welfare can increase or decrease. In fact, Example 2 describes an instance where it is not possible to improve demand fairness at all in this setting.

Example 2. Let $d_{00} = 0.5, d_{01} = 0.5, d_{10} = 0.6, d_{11} = 0.4$ and $b_{00} = 1, b_{01} = 3, b_{10} = 1.2, b_{11} = 2.4$. Suppose that the seller maximizes profit subject to the constraint $p_{0y} = p_{1y}$. Then, the optimal prices are $p_{00} = p_{10} = 0.65$ and $p_{01} = p_{11} = 1.45$. The resulting weighted average demand for groups $X = 0$ and $X = 1$ are 0.417 and 0.567, respectively. We next seek to impose demand fairness across both groups. When we vary $p_{00} = p_{10}$ by any Δp_{00}, the weighted average demand for both groups will change by $-0.5\Delta p_{00}$. Similarly, when we vary $p_{01} = p_{11}$ by any Δp_{01}, the weighted average demand for both groups will change by $-0.167\Delta p_{01}$. Namely, in both cases, no matter how we vary the prices, the difference in weighted average demand will remain the same. Thus, it is impossible to reduce the demand difference across groups (unless the demands of groups are set to zero).

5. Computations for Nonlinear Demand Functions

In this section, we investigate computationally the impact of fairness for alternate demand functions. Specifically, we consider the following three demand models: exponential, logistic, and log-log. We report the expressions of the demand $\bar{F}_i(p)$, consumer surplus $S_i(p)$, and mean valuation $\mathbb{E}[V_i]$ in Table 1. Note that we made a slight adjustment to the log-log demand function to ensure that it fits into the random utility framework.[8]

5.1. Setting with Two Groups

We first consider the problem with two groups of customers. We find the optimal pricing strategy by searching for the optimal $\bar{F}_i(p)$ between zero and one using 10^{-4} increments. For the log-log demand, it is possible that the market share reaches one, which corresponds to any price between zero and a_i. In this case, we also search for the optimal p_i between zero and a_i. We report the results for one representative instance of the logistic demand in Figure 3 (see Figures A.4 and A.5 in Online Appendix F for the exponential and log-log demand functions). Specifically, we show how the prices evolve as a function of α (left panels) as well as the profit, consumer surplus, and social welfare (right panels). We also list all the tested instances in Online Appendix F.1.

By conducting extensive computational tests, we find that imposing each of the four types of fairness constraints yields similar insights as in the case of linear demand. For exponential and logistic demand, we observe that under either price or no-purchase valuation fairness, the social welfare first increases as a function of α, whereas for either demand or surplus fairness, the social welfare decreases monotonically with α.

Under price fairness, it is still possible that p_1 changes nonmonotonically with α (see the top left panel of Figure 3) and that group 0 is (approximately) excluded by setting both prices close to p_1^* (with nearly zero demand from group 0). As we have shown in Section 3, such cases occur when the population of the high-valuation group is large. As a result, under price fairness, we retrieve the result that the social welfare first increases with α and then decreases. A major difference between the linear demand and the nonlinear models considered in this section emerges from the no-purchase valuation fairness. More specifically, the social welfare is not always increasing as it was the case for linear demand. Thus, under price or no-purchase valuation fairness, even though a small value of α increases social welfare, the specific value of α needs to be carefully selected.

5.2. Setting with Multiple Groups

We next test the performance of the nonlinear demand models when there are $n > 2$ groups of customers. Because Problem (2) is nonconvex and there are N decision variables, using a search heuristic can be burdensome. Interestingly, by exploiting the structure of the problem, we find that the optimal solution can be found efficiently by reducing the N-group pricing Problem (2) to a one-dimensional optimization problem. We leave the detailed discussion in Online Appendix E.

We next discuss the results for 20 randomly generated instances with $n = 5$ groups. Representative figures and the details of the instances are reported in Online Appendix F.2. Although the way the prices vary with α is more intricate than before, most of the analytical results we derived for the case of linear demand still hold (computationally) for the nonlinear demand models we considered. For all demand models, imposing price fairness is always beneficial at first. However, increasing the level of price fairness too much may prompt the seller to exclude low-value groups via a price surge and thus, can lead to a lose-lose outcome. For demand fairness, we observe that it always reduces social welfare under exponential and logistic demands, but for log-log demand, it can go

Table 1. Demand, Surplus, and Mean Valuation for the Different Demand Models

Demand model/metric	$\bar{F}_i(p)$	$S_i(p)$	$\mathbb{E}[V_i]$
Exponential	$e^{-\lambda_i p}$	$\frac{1}{\lambda_i}e^{-\lambda_i p}$	$\frac{1}{\lambda_i}$
Logistic	$\frac{k_i e^{-\lambda_i p}}{1+k_i e^{-\lambda_i p}}$	$\frac{1}{\lambda_i}\log(1+k_i e^{-\lambda_i p})$	$\frac{1}{\lambda_i}\log(1+k_i)$
Log-log	$\min\left\{\left(\frac{a_i}{p}\right)^{\beta_i}, 1\right\}$	$\frac{\beta_i}{\beta_i-1}a_i(\bar{F}_i(p))^{1-1/\beta_i} - p\bar{F}_i(p)$	$\frac{\beta_i}{\beta_i-1}a_i$

Figure 3. Impact of Fairness Under Logistic Demand (Two Groups)

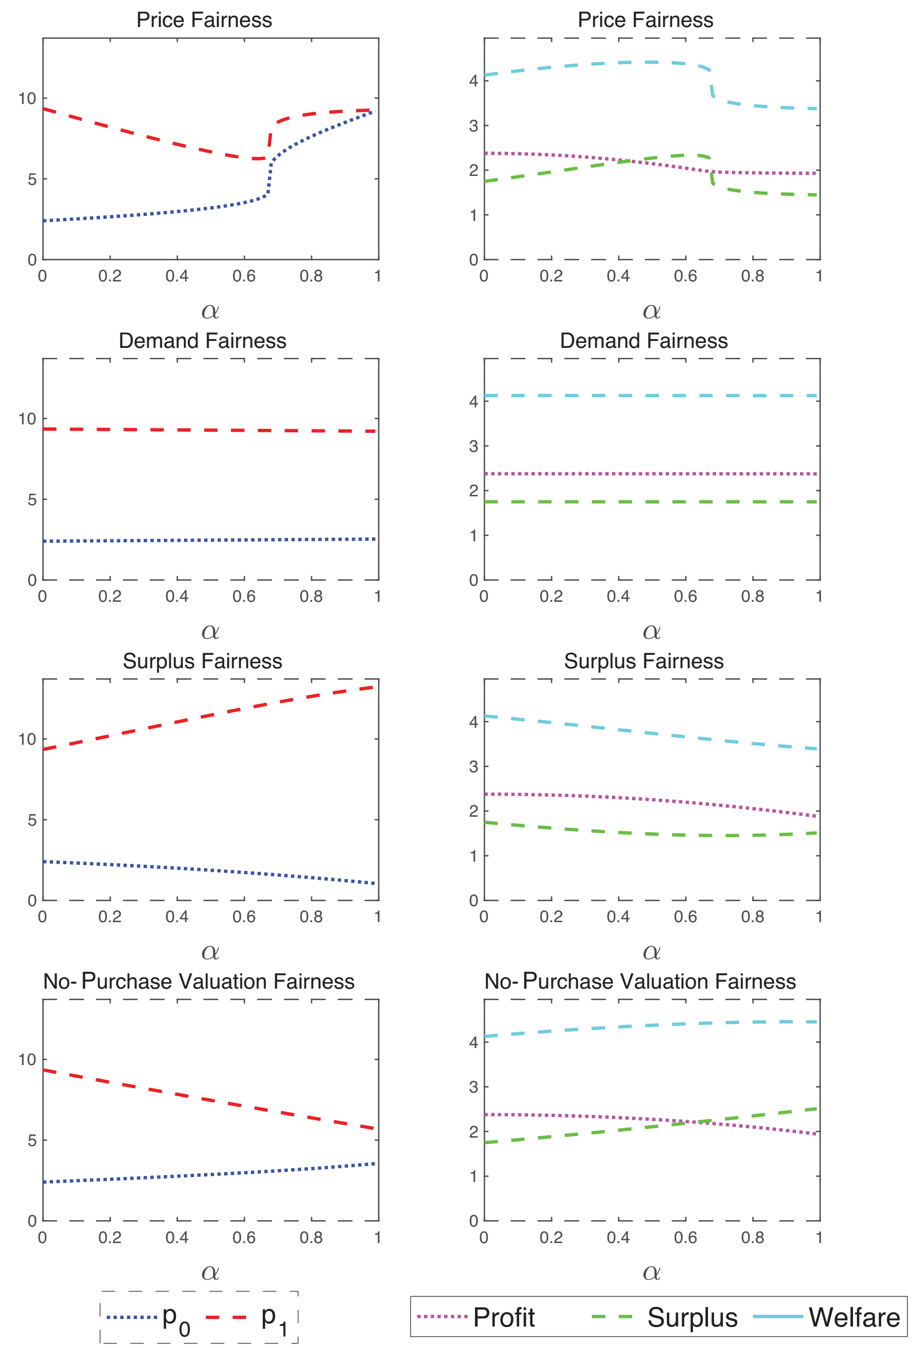

Note. Parameters: $d_0 = 0.5, d_1 = 0.5, \lambda_0 = 1, \lambda_1 = 0.2, k_0 = 10, k_1 = 5, c = 0.5$.

either way as in the two-group case. Imposing a small amount of surplus fairness decreases social welfare for all demand models. Finally, imposing no-purchase valuation fairness increases social welfare when α initially increases from zero under exponential and logistic demand. Under log-log demand, the social welfare may go either direction just as in the two-group case. Such findings increase our confidence that our managerial insights are robust and continue to hold for nonlinear demand models.

6. Conclusion

As discussed in Lobel (2021), although price discrimination has become common practice, it raises important questions in terms of fairness, which have been mostly unexplored. This paper offers a first step in understanding fairness in the context of pricing. We propose four possible fairness definitions—fairness in price, demand, consumer surplus, and no-purchase valuation—and investigate the impact of imposing fairness constraints on social welfare. We first show that imposing simultaneously several fairness metrics is generally impossible, hence reflecting the complexity of achieving perfect fairness in reality. We then focus on each fairness metric separately and characterize the optimal solutions in closed form under a linear demand model. We show that imposing a small amount of price fairness increases social welfare but that imposing too much price fairness may lead to a lose-lose outcome (i.e., both the seller and the consumers are worse off). Imposing either demand or surplus fairness always reduces the social welfare. Finally, imposing no-purchase valuation fairness increases the social welfare monotonically with the fairness level. Our findings also persist for a general setting with more than two customer groups, and most of our results hold computationally for three nonlinear demand models. Our insights have the potential to inform regulatory entities that are concerned with imposing fairness constraints on pricing.

Admittedly, much more research needs to be done on this topic. First, incorporating these fairness definitions into algorithms is an interesting avenue for future research and can potentially have great practical impact. Second, one can consider the role of inventory or capacity constraints in this setting, which may potentially evolve over time. Given the extensive research on fairness related to resource allocation, it would be interesting to develop a combined framework for fairness in both inventory allocation and pricing, which might require a different notion of surplus (Cohen et al. 2017). Another interesting extension of our model would be to consider pricing decisions that can only be made with partial information, such as the mean and variance of customer valuations (Chen et al. 2022). We also recognize that there might be competition between multiple sellers, another dimension that is unexplored in this paper. Finally, running behavioral surveys to learn how the different fairness definitions are perceived by consumers will help better understand how to properly define a fair pricing policy.

Acknowledgments

The authors thank the editorial team for their thoughtful feedback, which has helped improve the clarity and contribution of the paper. The authors also thank Vivek Farias for a valuable discussion of the paper at the 2021 INFORMS Revenue Management and Pricing Conference. The authors are listed alphabetically. The third author is a student author.

Endnotes

[1] Bostock v. Clayton County was decided on June 15, 2020.

[2] See https://www.natlawreview.com/article/new-york-implements-pink-tax-ban.

[3] See https://www.huffpost.com/entry/pink-tax-examples_l_5d24da77e4b0583e482850f0.

[4] See https://www.ftc.gov/news-events/blogs/business-blog/2020/04/using-artificial-intelligence-algorithms.

[5] See https://www.hud.gov/sites/documents/FAIR_LENDING_GUIDE.PDF.

[6] We highlight that the calligraphic quantities $\mathcal{R}(\cdot)$, $\mathcal{S}(\cdot)$, and $\mathcal{W}(\cdot)$ denote functions of α, whereas $R(\cdot)$, $S(\cdot)$, and $W(\cdot)$ denote functions of p.

[7] The social welfare (right) derivative at $\alpha = 0$ is defined as $\mathcal{W}'(0) = \lim_{\alpha \to 0^+} \frac{\mathcal{W}(\alpha) - \mathcal{W}(0)}{\alpha}$.

[8] The common form of the log-log demand is $p = a_i q^{-1/\beta_i}$, where q is the demand and β_i is the price elasticity. Because the demand goes to infinity when the price approaches zero, we truncate the demand at one: that is, we impose $\bar{F}_i(p) = \min\left\{\left(\frac{a_i}{p}\right)^{\beta_i}, 1\right\}$. We also require that $a_i(\beta_i - 1) < c\beta_i$, so that $\bar{F}_i(p_i^*) < 1$ (otherwise, all the customers are buying, hence leading to an unrealistic situation).

References

Alesina AF, Lotti F, Mistrulli PE (2013) Do women pay more for credit? evidence from italy. *J. Eur. Econom. Assoc.* 11(suppl 1):45–66.

Arneson RJ (1989) Equality and equal opportunity for welfare. *Philos. Stud.* 56(1):77–93.

Baron O, Hu M, Najafi-Asadolahi S, Qian Q (2015) Newsvendor selling to loss-averse consumers with stochastic reference points. *Manufacturing Service Oper. Management* 17(4):456–469.

Bartlett R, Morse A, Stanton R, Wallace N (2022) Consumer-lending discrimination in the FinTech era. *J. Financial Econom.* 143(1): 30–56.

Bateni MH, Chen Y, Ciocan DF, Mirrokni V (2016) Fair resource allocation in a volatile marketplace. *Proc. 2016 ACM Conf. Econom. Comput. (EC '16)* (Association for Computing Machinery, New York), 819.

Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Oper. Res.* 59(1):17–31.

Bertsimas D, Farias VF, Trichakis N (2012) On the efficiency-fairness trade-off. *Management Sci.* 58(12):2234–2250.

Besbes O, Zeevi A (2015) On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Management Sci.* 61(4):723–739.

Brams SJ, Taylor AD (1996) *Fair Division: From Cake-Cutting to Dispute Resolution* (Cambridge University Press, Cambridge, UK).

Brynjolfsson E, Collis A, Eggers F (2019) Using massive online choice experiments to measure changes in well-being. *Proc. Natl. Acad. Sci. USA* 116(15):7250–7255.

Brynjolfsson E, Hu Y(J), Smith MD (2003) Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Sci.* 49(11): 1580–1596.

Byun S-y, Park H (2012) The academic success of East Asian American youth: The role of shadow education. *Sociol. Ed.* 85(1):40–60.

Charles KK, Hurst E, Stephens M (2008) Rates for vehicle loans: Race and loan source. *Amer. Econom. Rev.* 98(2):315–320.

Chen Y, Wang H (2018) Why are fairness concerns so important? Lessons from pricing a shared last-mile transportation system. Working paper, Temple University, Philadelphia.

Chen H, Hu M, Perakis G (2022) Distribution-free pricing. *Manufacturing Service Oper. Management.* Forthcoming.

Chouldechova A (2017) Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data* 5(2):153–163.

Cohen MC, Perakis G, Pindyck RS (2021) A simple rule for pricing with limited knowledge of demand. *Management Sci.* 67(3): 1608–1621.

Cohen MC, Perakis G, Thraves C (2022) Consumer surplus under demand uncertainty. *Production Oper. Management* 31(2):478–494.

Cohen MC, Gupta S, Kalas JJ, Perakis G (2020) An efficient algorithm for dynamic pricing using a graphical representation. *Production Oper. Management* 29(10):2326–2349.

Cohen P, Hahn R, Hall J, Levitt S, Metcalfe R (2016) Using big data to estimate consumer surplus: The case of Uber. NBER Working Paper No. 22627, National Bureau of Economic Research, Cambridge, MA.

Cui HT, Raju JS, Zhang ZJ (2007) Fairness and channel coordination. *Management Sci.* 53(8):1303–1314.

Donahue K, Kleinberg J (2020) Fairness and utilization in allocating resources with uncertain demand. *Proc. 2020 Conf. Fairness Accountability Transparency (FAT '20)* (Association for Computing Machinery, New York), 658–668.

Donini M, Oneto L, Ben-David S, Shawe-Taylor JS, Pontil M (2018) Empirical risk minimization under fairness constraints. *Proc. 32nd Internat. Conf. on Neural Inform. Processing Systems (NIPS'18)* (Curran Associates Inc., Red Hook, NY), 2796–2806.

Dwork C, Hardt M, Pitassi T, Reingold O, Zemel R (2012) Fairness through awareness. *Proc. 3rd Innovations Theoret. Comput. Sci. Confer. (ITCS '12)* (Association for Computing Machinery, New York), 214–226.

Dworkin R (1981) What is equality? Part 1: Equality of welfare. *Philos. Public Affairs* 10(3):185–246.

Elmachtoub AN, Gupta V, Hamilton ML (2021) The value of personalized pricing. *Management Sci.* 67(10):6055–6070.

Eyster E, Madarasz K, Michaillat P (2021) Pricing under fairness concerns. *J. Eur. Econom. Assoc.* 19(3):1853–1898.

Fang L, Munneke HJ (2020) Gender equality in mortgage lending. *Real Estate Econom.* 48(4):957–1003.

FCA (2019) Fair pricing in financial services: Summary of responses and next steps. Financial Conduct Authority, United Kingdom. Accessed on March 13, 2022, https://www.fca.org.uk/publication/feedback/fs19-04.pdf.

Fehr E, Schmidt KM (1999) A theory of fairness, competition, and cooperation. *Quart. J. Econom.* 114(3):817–868.

Gallego G, Topaloglu H (2019) Online learning. *Revenue Management and Pricing Analytics, International Series in Operations Research & Management Science* (Springer, New York), 275–289.

Gee C (2018) Fair pricing in financial services. Financial Conduct Authority, United Kingdom. Accessed on March 13, 2022, https://www.fca.org.uk/publication/discussion/dp18-09.pdf.

Gerlick JA, Liozu SM (2020) Ethical and legal considerations of artificial intelligence and algorithmic decision-making in personalized pricing. *J. Revenue Pricing Management* 19(2020):85–98.

Gillis TB (2020) False dreams of algorithmic fairness: The case of credit pricing. Working paper, Columbia University, New York.

Gupta S, Kamble V (2019) Individual fairness in hindsight. *Proc. 2019 ACM Conf. Econom. Comput. (EC '19)* (Association for Computing Machinery, New York), 805–806.

Hardt M, Price E, Srebro N (2016) Equality of opportunity in supervised learning. *Proc. 30th Internat. Conf. Neural Inform. Processing Systems (NIPS'16)* (Curran Associates Inc., Red Hook, NY), 3315–3323.

Heidari H, Ferrari C, Gummadi K, Krause A (2018) Fairness behind a veil of ignorance: A welfare analysis for automated decision making. *Proc. 32nd Internat. Conf. Neural Inform. Processing Systems (NIPS'18)* (Curran Associates Inc., Red Hook, NY), 1273–1283.

Heidari H, Loi M, Gummadi KP, Krause A (2019) A moral framework for understanding fair ml through economic models of equality of opportunity. *Proc. Conf. Fairness Accountability Transparency (FAT '19)* (Association for Computing Machinery, New York), 181–190.

Ho T-H, Su X (2009) Peer-induced fairness in games. *Amer. Econom. Rev.* 99(5):2022–2049.

Hooker JN, Williams HP (2012) Combining equity and utilitarianism in a mathematical programming model. *Management Sci.* 58(9):1682–1693.

Hu L, Chen Y (2018) Welfare and distributional impacts of fair classification. Working paper, Harvard University, Cambridge, MA.

Hu L, Chen Y (2020) Fair classification and social welfare. *Proc. 2020 Conference on Fairness, Accountability, and Transparency (FAT '20)* (Association for Computing Machinery, New York), 535–545.

Hu X-F, Li D-F, Xu G-J (2018) Fair distribution of surplus and efficient extensions of the Myerson value. *Econom. Lett.* 165:1–5.

Kahneman D, Knetsch JL, Thaler R (1986) Fairness as a constraint on profit seeking: Entitlements in the market. *Amer. Econom. Rev.* 76(4):728–741.

Kallus N, Mao X, Zhou A (2021) Assessing algorithmic fairness with unobserved protected class using data combination. *Management Sci.*, ePub ahead of print April 2, https://doi.org/10.1287/mnsc.2020.3850.

Kasy M, Abebe R (2020) Fairness, equality, and power in algorithmic decision making. *Proc. 2021 ACM Conf. Fairness Accountability Transparency (FAT '21)* (Association for Computing Machinery, New York), 576–586.

Kleinberg J, Mullainathan S, Raghavan M (2017) Inherent trade-offs in the fair determination of risk scores. *Proc. 8th Innovations in Theoretical Computer Science Conference (ITCS 2017)* (Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik), 43:1–23.

Larson J, Mattu S, Angwin J (2015) Unintended consequences of geographic targeting. *Tech. Sci.*, https://techscience.org/a/2015090103/.

Levi R, Perakis G, Romero G (2016) On the effectiveness of uniform subsidies in increasing market consumption. *Management Sci.* 63(1):40–57.

Li KJ, Jain S (2016) Behavior-based pricing: An analysis of the impact of peer-induced fairness. *Management Sci.* 62(9):2705–2721.

Lobel I (2021) Revenue management and the rise of the algorithmic economy. *Management Sci.* 67(9):5389–5398.

Madden TJ, Hewett K, Roth MS (2000) Managing images in different cultures: A cross-national study of color meanings and preferences. *J. Internat. Marketing* 8(4):90–107.

Marcoux AM (2006) Much ado about price discrimination. *J. Markets Morality* 9(1):57–69.

Nabi R, Malinsky D, Shpitser I (2019) *Learning optimal fair policies*. *Proc. 36th Internat. Conf. Machine Learning, Long Beach, California, PMLR* 97:4674–4682.

Popescu I, Wu Y (2007) Dynamic pricing strategies with reference effects. *Oper. Res.* 55(3):413–429.

Rabin M (1993) Incorporating fairness into game theory and economics. *Amer. Econom. Rev.* 83(5):1281–1302.

Rahmattalabi A, Vayanos P, Fulginiti A, Rice E, Wilder B, Yadav A, Tambe M (2019) Exploring algorithmic fairness in robust graph covering problems. *Proc. 33rd Internat. Conf. Neural Inform. Processing Systems* (Curran Associates Inc., Red Hook, NY), Article 1413, 15776–15787.

Rawls J (1971) *A Theory of Justice* (Harvard University Press, Cambridge, MA).

Rawls J (1999) *A Theory of Justice*, revised ed. (Harvard University Press, Cambridge, MA).

Robinson J (1969) *The Economics of Imperfect Competition* (Springer, New York).

Schmalensee R (1981) Output and welfare implications of monopolistic third-degree price discrimination. *Amer. Econom. Rev.* 71(1):242–247.

Sethuraman J, Teo C-P, Qian L (2006) Many-to-one stable matching: Geometry and fairness. *Math. Oper. Res.* 31(3):581–596.

Varian HR (1985) Price discrimination and social welfare. *Amer. Econom. Rev.* 75(4):870–875.

Viviano D, Bradic J (2020) Fair policy targeting. Working paper, University of California at San Diego, La Jolla, CA.

Wallheimer B (2018) Are you ready for personalized pricing? *Chicago Booth Rev.* (February 26), https://review.chicagobooth.edu/marketing/2018/article/are-you-ready-personalized-pricing.

Weinberger D (2019) How machine learning pushes us to define fairness. *Harvard Bus. Rev.* (November 6), https://hbr.org/2019/11/how-machine-learning-pushes-us-to-define-fairness.

White House (2015) Big data and differential pricing. White House Report. Accessed on March 13, 2022, https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/docs/Big_Data_Report_Nonembargo_v2.pdf/.

Xue Z, Wang Z, Ettl M (2015) Pricing personalized bundles: A new approach and an empirical study. *Manufacturing Service Oper. Management* 18(1):51–68.

Ye P, Qian J, Chen J, Chen-hung W, Zhou Y, De Mars S, Yang F, Zhang L (2018) Customized regression model for airbnb dynamic pricing. *Proc. 24th ACM SIGKDD Internat. Conf. Knowledge Discovery & Data Mining (KDD '18)* (Association for Computing Machinery, New York), 932–940.