

Personalized Pricing with Group Fairness Constraint

Xin Chen

Georgia Institute of Technology
Atlanta, GA, USA
xin.chen@isye.gatech.edu

Zishuo Zhao

University of Illinois Urbana-Champaign
Champaign, IL, USA
zishuoz2@illinois.edu

Zexing Xu

University of Illinois Urbana-Champaign
Champaign, IL, USA
zexingx2@illinois.edu

Yuan Zhou

Tsinghua University
Beijing, China
yuan-zhou@tsinghua.edu.cn

ABSTRACT

In the big data era, personalized pricing has become a popular strategy that sets different prices for the same product according to individual customers' features. Despite its popularity among companies, this practice is controversial due to the concerns over fairness that can be potentially caused by price discrimination. In this paper, we consider the problem of single-product personalized pricing for different groups under fairness constraints. Specifically, we define group fairness constraints under different distance metrics in the personalized pricing context. We then establish a stochastic formulation that maximizes the revenue. Under the discrete price setting, we reformulate this problem as a linear program and obtain the optimal pricing policy efficiently. To bridge the gap between the discrete and continuous price setting, theoretically, we prove a general $\mathcal{O}(\frac{1}{l})$ gap between the optimal revenue with continuous and discrete price set of size l . Under some mild conditions, we improve this bound to $\mathcal{O}(\frac{1}{l^2})$. Empirically, we demonstrate the benefits of our approach over several baseline approaches on both synthetic data and real-world data. Our results also provide managerial insights into setting a proper fairness degree as well as an appropriate size of discrete price set.

CCS CONCEPTS

• **Social and professional topics** → **Computing / technology policy**; • **Applied computing** → *Electronic commerce; Operations research.*

KEYWORDS

personalized pricing, group fairness, statistical parity, social welfare

ACM Reference Format:

Xin Chen, Zexing Xu, Zishuo Zhao, and Yuan Zhou. 2023. Personalized Pricing with Group Fairness Constraint. In *2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*, June 12–15, 2023, Chicago, IL, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FAccT '23, June 12–15, 2023, Chicago, IL, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0192-4/23/06...\$15.00

<https://doi.org/10.1145/3593013.3594097>

USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3593013.3594097>

1 INTRODUCTION

The widespread availability of individuals' features and behavioral information has led to the increased adoption of personalized pricing strategies in various industries. Firms can leverage collected data to predict consumers' willingness to pay, ultimately increasing profits by charging each customer exactly their valuation. However, despite its popularity and profitability, concerns about fairness in personalized pricing have arisen due to its potential discriminatory nature. Protected groups, such as Black individuals in race, Asians in national origin, and females in gender identity, may be charged higher prices due to their potentially higher valuations for a product, resulting in disparities against these protected classes [8, 15, 20]. Moreover, this pricing strategy is illegal if it discriminates on the basis of race, religion, nationality, or gender, or if it violates antitrust or price-fixing laws such as the Civil Rights Act of 1964 and the Equal Credit Opportunity Act of 1974. Ensuring fairness is also a concern for the Federal Trade Commission (FTC).

Considering the significance of fairness in the pricing context, we investigate a single-product personalized pricing problem for different groups under group fairness constraints. We propose a constrained revenue maximization framework that balances the seller's profit and the fairness of different groups based on their sensitive attributes.

Recent works have focused on developing fair pricing algorithms to restrict price discrimination [22–24, 27, 41]. A study by [12] considered a simple scenario of a single-product seller facing two consumer groups and proposed pricing approaches with fairness constraints in price, demand, customer surplus, and social welfare. They characterized the impact of imposing different types of fairness. However, their model did not involve features, and they adopted a single pricing strategy where the monopolist offers the product to all customers in one group at the same price. In contrast, we apply a feature-based "contextual" pricing strategy where the monopolist offers a customized price based on the observed feature vectors for each customer. Moreover, their price fairness measure is a looser constraint on price than ours, which only restricts the single prices of the two sensitive groups to be close. Another work [28] studied fairness and ethical issues under personalized pricing settings but only considered parity constraints with respect to the first moment of the price distribution, which is also looser than ours. In addition, these two works only considered limited types

of demand models, while our results can extend to more general demand models with mild assumptions.

In this paper, we generalize the non-contextual pricing framework in [12] to the contextual setting [19] and consider a stronger price fairness constraint than [12, 28]. Our contributions are four-folds:

- We introduce a context-dependent personalized pricing model that considers group fairness by measuring statistical parity, a more suitable metric in the personalized pricing context compared to alternatives like accuracy parity, equalized odds, or predictive value parity [3, 28]. This model captures price sensitivity and customer valuation distributions in a more practical manner for real-world settings. To our knowledge, this is the first work to formulate a personalized pricing problem with fairness constraints as a stochastic programming problem, where the constraint is the pairwise probability distance between the price distributions of different groups.
- To solve the stochastic formulation under discrete price settings, we reformulate it as a maximum flow type of linear program which can be solved efficiently. To bridge the gap between the discrete and continuous price settings, we theoretically prove an $\mathcal{O}(\frac{1}{l})$ gap between the optimal revenue with continuous and discrete price sets of size l under fairness constraints measured by total variation distance (TVD) and earth mover's distance (EMD). Furthermore, under mild assumptions, we improve this bound to $\mathcal{O}(\frac{1}{l^2})$ for the TVD constraint.
- Empirically, we demonstrate the benefits of the fairness constraint under several metrics (i.e., higher customer surplus, higher social welfare, and lower Gini Index) and superior pricing policy compared with several baselines, and provide business insights into the implementation of this algorithm (i.e., setting a proper fairness degree and an appropriate size of discrete price set).

1.1 Related Literature

Our work is mainly related to three streams of literature:

Personalized pricing and price discrimination. With the increasing availability of consumers' data, personalized pricing or price discrimination has become more popular in e-commerce [11, 18], airlines [39], and many other industries [27]. The value of personalized pricing is studied in [4, 19, 34, 38] and there are several methods have been proposed to solve the personalized pricing problem efficiently in the offline setting [7, 9]. There are three types of price discrimination in the classical economic taxonomy [36]. First-degree price discrimination offers individual prices to customers exactly at their willingness to pay, which is assumed to be known. Second-degree price discrimination depends on the quantity purchased but does not differ across consumers, such as bulk discounts. Third-degree price discrimination charges different prices to different groups of consumers. We focus on analyzing contextual prices, which is first-degree price discrimination.

Fairness in personalized pricing. The discriminatory nature of personalized pricing has triggered a heated debate among policy-makers and academics on designing fair policies to restrict price

discrimination [22–24, 27, 41]. People have developed legal constraints on anti-discrimination [2, 29], which aim to protect different subgroups of consumers such as females and blacks. In this paper, similar to [13, 28], we suppose a monopoly with perfect information on consumers' willingness to pay and consider the regulation towards first-degree price discrimination.

Algorithmic Fairness. Recently, the volume of literature and public attention on machine learning fairness has been growing significantly [5, 33]. Various fairness notions, including group fairness [25, 30, 42], individual fairness [17, 44], and causality-based fairness notions [10, 31, 32] are proposed to protect different subgroups or individuals. In this work, we adopt statistical parity as our fairness metric.

The rest of this paper is organized as follows. In Section 2, we first provide the preliminaries of the problem setting and define the distance metric and fairness measure. Then, we establish our optimization model with group fairness constraint. In Section 3, we illustrate the linear programming reformulation of the original optimization problem with group fairness constraint. We also provide theoretical guarantees for our algorithm. In Section 4, we demonstrate the benefits of our approach over several baseline approaches on both synthetic data and real-world data. In the main text, we only show the results under the TVD constraint, while all theoretical and numerical results could be extended to the problem under the EMD constraint, which is deferred to the Appendix.

2 MODEL

2.1 Problem Setup

We study a monopolist selling a product with the ability to observe each customer's feature vector, represented as an $(m + q)$ -dimensional random vector $\tilde{\mathbf{X}} := (\mathbf{X}, \mathbf{S}) \in \mathbb{R}^{m+q}$. Here, $\mathbf{S} := (S_1, \dots, S_q) \in \mathbb{R}^q$ denotes the sensitive feature vector (e.g., race, gender, nationality), and $\mathbf{X} := (X_1, \dots, X_m) \in \mathbb{R}^m$ represents the feature vector excluding sensitive attributes.

The problem involves $n = \sum_{u=1}^k n_u$ customers divided into k sensitive groups, with n_u customers in each group u , $\forall u \in [k]$. We assume each customer's feature vector $\tilde{\mathbf{X}}_i := (\mathbf{X}_i, \mathbf{S}_i)$, $\forall i \in [n]$ to be independent and identically distributed. However, each component in $\tilde{\mathbf{X}}_i$ may not be independent, as specific features may be correlated, such as gender affecting height or weight.

To define the fairness constraint, we use a function $g : \mathbb{R}^q \mapsto 1, 2, \dots, k$ to map sensitive features to group labels, where k represents the number of possible groups. In our focus, each sensitive feature takes a finite number of values, and a group is specified by the value of sensitive features. Thus, all customers with the same sensitive features are assigned to the same group. We can also treat each sensitive feature separately or define groups based on combinations of sensitive features.

We assume that the conditional distribution $\mathbf{X}|\mathbf{S}$ is close for each group $u \in [k]$. If not, the difference in price distribution may result from the discrepancy in non-sensitive attributes, deviating from our original purpose of imposing a fairness constraint.

Each customer has a valuation (i.e., willingness-to-pay) for the product denoted by a function $V(\tilde{\mathbf{X}}) \sim F_{\tilde{\mathbf{X}}}(\cdot)$, where $F_{\tilde{\mathbf{X}}}(\cdot)$ is the

feature-dependent cumulative distribution function known by the seller, and $V(\tilde{X})$ is the valuation function of attribute $\tilde{X}$.

We assume that the conditional distribution $X|S$ needs to be close for each group $u \in [k]$. Otherwise, the difference in the price distribution may cause by the discrepancy in their non-sensitive attributes, which deviate from our original purpose of imposing the fairness constraint.

Each customer has a valuation (i.e., willingness-to-pay) for the product denoted by a function $V(\tilde{X}) \sim F_{\tilde{X}}(\cdot)$, where $F_{\tilde{X}}(\cdot)$ is the feature-dependent cumulative distribution function known by the seller and $V(\tilde{X})$ is the valuation function of attribute $\tilde{X}$. We define the pricing policy $\rho : \mathbb{R}^{m+q} \mapsto \mathcal{P}$ mapping the features to the admissible prices, where $\mathcal{P}$ is a bounded price set created by the seller.

Then, the demand function is defined as

$$D(\tilde{X}, \rho(\tilde{X})) = \mathbb{E}_V[a(\tilde{X}) \mathbb{1}[V(\tilde{X}) > \rho(\tilde{X})]], \quad (1)$$

where $\mathbb{1}[\cdot]$ is the 0-1 indicator function and function $a : \mathbb{R}^{m+q} \mapsto \mathbb{R}$ is a feature dependent scalar.

Our demand function differs from the classical valuation-based demand function [13, 28] in that besides the purchasing probability, we also capture the purchasing amount by a multiply of a context-aware function $a(\tilde{X})$ which could better capture the personalized behavior and is appropriate in many business setting. Taking the product candy for example, the intention of buying candies for the elderly and children might be similar but the amount they buy usually differs since children buy more candies than the elderly. In this case, the difference in market size for different customer groups can be captured by our formulation.

Given a pricing policy ρ , the seller's expected revenue is

$$\begin{aligned} R(\rho) &:= \mathbb{E}_{\tilde{X}}[D(\tilde{X}, \rho(\tilde{X}))\rho(\tilde{X})] \\ &= \mathbb{E}_{\tilde{X}, V}[a(\tilde{X}) \mathbb{1}[V(\tilde{X}) > \rho(\tilde{X})]\rho(\tilde{X})], \end{aligned} \quad (2)$$

and the corresponding customer surplus is

$$S(\rho) := \mathbb{E}_{\tilde{X}, V}[a(\tilde{X})(V(\tilde{X}) - \rho(\tilde{X}))^+], \quad (3)$$

where $(\cdot)^+ := \max\{0, \cdot\}$.

For society as a whole, the total welfare is a combination of the revenue obtained by the seller through selling and the surplus $S(\rho)$ gained by the customer via purchase, which is

$$\begin{aligned} W(\rho) &:= R(\rho) + S(\rho) \\ &= \mathbb{E}_{\tilde{X}, V}[a(\tilde{X}) \mathbb{1}[V(\tilde{X}) > \rho(\tilde{X})]V(\tilde{X})]. \end{aligned} \quad (4)$$

We consider general demand models, for instance, linear, exponential, logistic, and log-log demand. Their expressions of the demand D , Revenue R , and consumer surplus S are listed in Table 1. Note that we winsorize the support of the logistic demand and made a slight adjustment to the log-log demand function to ensure that it fits into the random utility framework¹.

¹The common form of log-log demand is $a(\tilde{x}) \left(\frac{c(\tilde{x})}{\rho(\tilde{x})} \right)^{b(\tilde{x})}$, where $b(\tilde{x})$ is the price elasticity and $\rho(\tilde{x})$ is the price given the context $\tilde{x}$. To avoid the demand goes to infinity when the price $\rho(\tilde{x})$ is close to zero, we truncate it by $a(\tilde{x}) \min \left\{ \left(\frac{c(\tilde{x})}{\rho} \right)^{b(\tilde{x})}, 1 \right\}$.

Also, we require $c(\tilde{x})(b(\tilde{x}) - 1) < \text{Constant} \cdot b(\tilde{x})$ to ensure that the no-purchase probability is positive.

Consider the personalized pricing framework: upon observing $\tilde{X} = \tilde{x}$, the seller offers a price from the price set $\mathcal{P}$ to each individual. Then, the seller obtains each customer's demand in response to the price through the known demand function D defined in Eq. (1).

Finally, the seller's goal is to determine a pricing policy ρ that maximizes the revenue.

However, without any regulation constraint, the monopoly could charge each consumer with his or her willingness to pay exactly (if the monopoly knows them), which is well known as first-degree price discrimination [6]. In this case, consumers get no benefits and the revenue is maximized.

However, in practice, certain price discrimination is prohibited by regulations for the concerns of social well-being. Specifically, in this paper, we propose a personalized pricing framework incorporating group fairness for sensitive consumer groups, which will be defined in Section 2.2.

2.2 Distance Metrics and Fairness Measure

In this section, we introduce the foundation of statistical parity by first defining the distance metric. Similar to [16], we consider two different distance metrics under statistical parity: total variation distance (TVD) and earth mover's distance (EMD), which are commonly used for measuring distribution differences in statistical parity [40].

DEFINITION 1 (TOTAL VARIATION DISTANCE (TVD)). *Given the set M of all probability measures on a countable set Ω . The total variation distance of probability measures $\text{TVD} : M \times M \mapsto [0, 1]$ is defined as*

$$\text{TVD}(Q_1, Q_2) = \frac{1}{2} \sum_{\omega \in \Omega} |Q_1(\omega) - Q_2(\omega)|, \quad (5)$$

where Q_1 and Q_2 denote two probability measures on the finite domain Ω .

The definition of EMD is deferred to Appendix A.1.

In our experiments, we observed that TVD and EMD yield qualitatively similar insights into the fairness of personalized pricing policies. While there are differences in the mathematical properties of these metrics and their respective bounds, the managerial insights derived from both metrics generally align. This similarity in insights might suggest that, regardless of the specific probability metric used, the qualitative findings in terms of fairness in personalized pricing are likely to be consistent.

Group fairness is a requirement that the protected groups should be treated similarly to the advantaged group or the population as a whole. Common notions concerning group fairness are statistical parity [16, 40], disparate impact [26], statistical discrimination [43] and etc. Here, we consider statistical parity measuring the difference in probabilities of an outcome across two groups², which is defined as follows.

DEFINITION 2. Statistical Parity: *Two probability measures Q_1 and Q_2 satisfy statistical parity up to bias $\delta \geq 0$ if*

$$d(Q_1, Q_2) \leq \delta, \quad (6)$$

²This can be generalized to more than two groups. For multiple groups, we compare their probability measures pair-wisely.

Demand model	$D(\tilde{\mathbf{x}}, \rho)$	$R(\tilde{\mathbf{x}}, \rho)$	$S(\tilde{\mathbf{x}}, \rho)$
Linear	$a(\tilde{\mathbf{x}}) \max \left\{ 0, \left(1 - \frac{\rho}{c(\tilde{\mathbf{x}})} \right) \right\}$	$a(\tilde{\mathbf{x}}) \max \left\{ 0, \left(1 - \frac{\rho}{c(\tilde{\mathbf{x}})} \right) \right\} \rho$	$a(\tilde{\mathbf{x}}) \frac{(\max\{0, c(\tilde{\mathbf{x}}) - \rho\})^2}{2c(\tilde{\mathbf{x}})}$
Exponential	$a(\tilde{\mathbf{x}}) e^{-c(\tilde{\mathbf{x}})\rho}$	$a(\tilde{\mathbf{x}}) e^{-c(\tilde{\mathbf{x}})\rho} \rho$	$\frac{a(\tilde{\mathbf{x}})}{c(\tilde{\mathbf{x}})} e^{-c(\tilde{\mathbf{x}})\rho}$
Logistic	$\frac{b(\tilde{\mathbf{x}}) e^{-c(\tilde{\mathbf{x}})\rho}}{1 + b(\tilde{\mathbf{x}}) e^{-c(\tilde{\mathbf{x}})\rho}}$	$\frac{a(\tilde{\mathbf{x}}) b(\tilde{\mathbf{x}}) e^{-c(\tilde{\mathbf{x}})\rho}}{1 + b(\tilde{\mathbf{x}}) e^{-c(\tilde{\mathbf{x}})\rho}} \rho$	$\frac{a(\tilde{\mathbf{x}})}{c(\tilde{\mathbf{x}})} \log(1 + b(\tilde{\mathbf{x}}) e^{-c(\tilde{\mathbf{x}})\rho})$
Log-log	$a(\tilde{\mathbf{x}}) \min \left\{ \left(\frac{c(\tilde{\mathbf{x}})}{\rho} \right)^{b(\tilde{\mathbf{x}})}, 1 \right\}$	$a(\tilde{\mathbf{x}}) \min \left\{ \left(\frac{c(\tilde{\mathbf{x}})}{\rho} \right)^{b(\tilde{\mathbf{x}})}, 1 \right\} \rho$	$\frac{a(\tilde{\mathbf{x}}) b(\tilde{\mathbf{x}})}{b(\tilde{\mathbf{x}}) - 1} c(\tilde{\mathbf{x}}) \cdot \left(\min \left\{ \left(\frac{c(\tilde{\mathbf{x}})}{\rho} \right)^{b(\tilde{\mathbf{x}})}, 1 \right\} \right)^{1 - \frac{1}{b(\tilde{\mathbf{x}})}} - a(\tilde{\mathbf{x}}) \rho \min \left\{ \left(\frac{c(\tilde{\mathbf{x}})}{\rho} \right)^{b(\tilde{\mathbf{x}})}, 1 \right\}$

Table 1: Demand, Revenue, and Surplus for Different Demand Models

where $d(\cdot, \cdot)$ is some distance metric between two distributions.

Statistical parity is a more suitable fairness measure under the personalized pricing setting due to perception and regulations. Firstly, consumers are more concerned with being treated fairly in terms of the prices they are offered, rather than the accuracy or odds of receiving a particular price. Statistical parity focuses on ensuring that different groups receive a similar distribution of prices, thus promoting a perception of fairness among consumers. Secondly, many jurisdictions have regulations prohibiting price discrimination based on certain protected attributes, such as race, gender, or nationality. Statistical parity as a fairness measure aligns with these regulations by ensuring that the distribution of prices offered to different groups is not significantly different, thereby reducing the risk of regulatory violations.

2.3 Stochastic optimization problem with group fairness constraint

In this section, we introduce the stochastic optimization formulation of a personalized pricing model that takes into account group fairness. To simplify the presentation, we first assume that the seller prespecifies a discrete price set of size l , i.e.,

$$\mathcal{P} := \left\{ \{p_j\}_{j=1}^l \mid 0 < \underline{p} \leq p_1 < p_2 < \dots < p_l \leq \bar{p} < \infty \right\}, \quad (7)$$

where $F^u(\tilde{\mathbf{X}})$ denotes the cumulative distribution function of the feature $\tilde{\mathbf{X}}$ in group u . $\forall p_j, j \in [l]$, we define the probability mass of taking price p_j of group u as

$$\begin{aligned} q_u(p_j) &:= \mathbb{E}_{\tilde{\mathbf{X}}} [\mathbb{1}[g(\mathbf{S}) = u] \mathbb{1}[\rho(\tilde{\mathbf{X}}) = p_j]] \\ &= \int_{\tilde{\mathbf{X}}} \mathbb{1}[g(\mathbf{S}) = u] \mathbb{1}[\rho(\tilde{\mathbf{X}}) = p_j] dF^u(\tilde{\mathbf{X}}). \end{aligned} \quad (8)$$

The empirical counterpart of the price distribution in group u for discrete price $p_j \in [\underline{p}, \bar{p}]$, $\forall j \in [l]$ is

$$\hat{q}_u(p_j) := \frac{1}{n_u} \sum_{i=1}^n \mathbb{1}[g(\mathbf{s}_i) = u] \mathbb{1}[\rho(\tilde{\mathbf{x}}_i) = p_j]. \quad (9)$$

Now, we define $Q_u : \mathcal{P} \mapsto \Delta(\mathcal{P})$ as a measurable function, where $\Delta(\mathcal{P})$ is the probability simplex over the set $\mathcal{P}$ represented as $\Delta(\mathcal{P}) := \left\{ \{q_j\}_{j=1}^l \in [0, 1]^l : \sum_{j=1}^l q_j = 1 \right\}$. Then $\forall \mathbf{p} = (p_1, \dots, p_l) \in \mathcal{P}$, we have $Q_u((p_1, \dots, p_l)) := (q_u(p_1), \dots, q_u(p_l))$.

The revenue maximization problem with price parity fairness constraint can be written as

$$\begin{aligned} \max_{\rho} \quad & R(\rho) \\ \text{s.t.} \quad & d(Q_u(\mathbf{p}), Q_v(\mathbf{p})) \leq \delta d(\tilde{Q}_u(\mathbf{p}), \tilde{Q}_v(\mathbf{p})), \quad \forall u, v \in [k]. \end{aligned} \quad (10)$$

In the revenue maximization problem with group fairness constraint (10), the degree of fairness $\delta \geq 0$ determines our tolerance for discrimination. A δ value of 0 enforces strict fairness, while larger δ values allow for looser fairness constraints. $R(\rho)$ is defined in Equation (2) and Q_u is the probability mass given by our pricing policy for group $u \in [k]$ while $\tilde{Q}_u$ is that of the unconstrained problem (i.e. $\max_{\rho} R(\rho)$) for group $u \in [k]$. $d(\cdot, \cdot)$ denotes either of the two different distance metrics: total variation distance and earth mover's distance.

We define the model with mean price constraint as follows:

$$\begin{aligned} \max_{\rho} \quad & R(\rho(\cdot)) \\ \text{s.t.} \quad & |Q_u(\mathbf{p})^T \mathbf{p} - Q_v(\mathbf{p})^T \mathbf{p}| \leq \delta |\tilde{Q}_u(\mathbf{p})^T \mathbf{p} - \tilde{Q}_v(\mathbf{p})^T \mathbf{p}|, \quad \forall u, v \in [k]. \end{aligned} \quad (11)$$

If we restrict the pricing policy to be independent of the feature information, this model reduces to the price parity model in [14] which does not consider feature information for determining the pricing policy.

In the above formulation, we assume that the possible prices are given. It would be interesting to determine the optimal prices as well. Our model formulation under the discrete price setting can be generalized to this setting, which is omitted here. In the next section, we will show that as $l \rightarrow \infty$, the optimal objective under the discrete price setting converges to the continuous price setting.

3 ALGORITHM AND THEORETICAL RESULTS

Under sample average approximation, we reformulate problem (10) for general demand models as a linear program and solve it efficiently.

Here, we only display the LP program and theoretical results for TVD while its EMD counterpart is deferred to Appendix A.

3.1 Linear Program Reformulation under Total Variation Distance

In the linear program, let each customer i belong to a specific group $u \in [k]$, associated with a pair of feature vectors $(\mathbf{x}_i, \mathbf{s}_i)$ - a standard

feature vector $\mathbf{x}_i$ and a sensitive feature vector $\mathbf{s}_i$ corresponding to group u . The seller selects l admissible prices.

To describe the pricing policy by illustrating the connection between customers and prices, we define $f_{ij}^u := \mathbb{P}(\rho(\mathbf{x}_i) = p_j \mid g(\mathbf{s}_i) = u)$ as the probability of customer i in group u being assigned to price p_j , and $\alpha_j^u := \widehat{q}_u(p_j)$ as the proportion of customers in the group u assigned to price p_j by the pricing policy. The variables f and α are subject to the following constraints:

$$\begin{aligned} \sum_{j=1}^l f_{ij}^u &= 1, \quad \forall i \in [n_u], u \in [k], \\ f_{ij}^u &\geq 0, \quad \forall i \in [n_u], j \in [l], u \in [k], \\ \frac{1}{n_u} \sum_{i=1}^{n_u} f_{ij}^u &= \alpha_j^u, \quad \forall j \in [l], u \in [k]. \end{aligned}$$

The group fairness constraint for Total Variation Distance (TVD) is given by:

$$\sum_{j=1}^l |\alpha_j^u - \alpha_j^v| \leq \delta \sum_{j=1}^l |\tilde{\alpha}_j^u - \tilde{\alpha}_j^v|, \quad \forall u, v \in [k], \quad (12)$$

which ensures that the proportion of customers offered price p_j in two groups should not be δ -apart.

The revenue obtained from customers in group u is given by $r^u = \sum_{i=1}^{n_u} \sum_{j=1}^l f_{ij}^u p_j D_{ij}^u$, where D_{ij}^u is the demand for customer i with feature vector $\mathbf{x}_i^u$ in group u under price p_j . The total revenue is then $\sum_{u=1}^k \sum_{i=1}^{n_u} \sum_{j=1}^l f_{ij}^u p_j D_{ij}^u$.

Combining all the above, our problem is formulated as,

$$\begin{aligned} \max_{f, \alpha} \quad & \sum_{u=1}^k \sum_{i=1}^{n_u} \sum_{j=1}^l f_{ij}^u p_j D_{ij}^u \\ \text{s.t.} \quad & \sum_{j=1}^l f_{ij}^u = 1, \quad \forall i \in [n_u], u \in [k], \\ & \frac{1}{n_u} \sum_{i=1}^{n_u} f_{ij}^u = \alpha_j^u, \quad \forall j \in [l], u \in [k], \\ & \sum_{j=1}^l |\alpha_j^u - \alpha_j^v| \leq \delta \sum_{j=1}^l |\tilde{\alpha}_j^u - \tilde{\alpha}_j^v|, \quad \forall u, v \in [k], \\ & f_{ij}^u \geq 0, \quad \forall i \in [n_u], j \in [l], u \in [k]. \end{aligned} \quad (13)$$

Existing linear programming algorithms, such as the Simplex Algorithm, Interior Point Algorithms, and the Ellipsoid Method [1] can be applied to solve it efficiently.

However, the limited discrete price options may not exhaust the optimal prices. To address the issue, we consider the LP (13)³ with

continuous price set defined as

$$\begin{aligned} \max_{f, \alpha} \quad & \sum_{u=1}^k \sum_{i=1}^{n_u} \int_{\underline{p}}^{\bar{p}} f_i^u(p) p D_i^u(p) dp \\ \text{subject to:} \quad & \int_{\underline{p}}^{\bar{p}} f_i^u(p) dp = 1, \quad \forall i \in [n_u], u \in [k], \\ & \frac{1}{n_u} \sum_{i=1}^{n_u} f_i^u(p) = \alpha^u(p), \quad \forall p \in [\underline{p}, \bar{p}], u \in [k], \\ & \int_{\underline{p}}^{\bar{p}} |\alpha^u(p) - \alpha^v(p)| dp \leq \delta \int_{\underline{p}}^{\bar{p}} |\tilde{\alpha}^u(p) - \tilde{\alpha}^v(p)|, \quad \forall u, v \in [k], \\ & f_i^u(p) \geq 0, \quad \forall i \in [n_u], u \in [k], \end{aligned} \quad (14)$$

where the corresponding decision variables are $f_i^u(p)$ and $\alpha^u(p)$. $D_i^u(p)$ denotes the demand for customer i in group u w.r.t. the continuous price p .

Problem (14) is an infinite dimension and is not solvable in practice. Instead, we use the finite dimension LP (13) to approximately solve it. To show the efficacy of the proposed method, we provide a theoretical guarantee to show the convergence rate of the gap between the optimal revenue under a discrete price set and its continuous counterpart (14) in Section 3.2.

3.2 Theoretical Results

In this section, we bridge the gap between the discrete and continuous price setting by providing a general $\mathcal{O}(\frac{1}{l})$ bound on the difference between optimal revenue r^* under the price set $[\underline{p}, \bar{p}]$ and the optimal revenue r^l when the price set contains l equally-spaced prices. Furthermore, under the total variation distance constraints and under mild assumptions, this revenue gap can be improved to $\mathcal{O}(\frac{1}{l^2})$.

For all demand functions $D_i^u(p)$ satisfying Lipschitz continuity on $[\underline{p}, \bar{p}]$ for $\forall u, i$, we prove in Theorem 1 that the optimal revenue for the problem with a discrete price set converges to the optimal revenue for the problem with a continuous price set at a rate of $\mathcal{O}(\frac{1}{l})$.

THEOREM 1. Suppose r^* is the optimal revenue for continuous price set problem (14) and r^l is the optimal revenue for its discrete price setting counterpart (13). Then

$$r^* - r^l \leq \mathcal{O}(\frac{1}{l}).$$

Furthermore, as long as the demand function $D_i^u(p)$ is twice continuously differentiable on $[\underline{p}, \bar{p}]$ for $\forall u, i$ with bounded first and second derivatives (which is satisfied in most demand models), we can improve the revenue gap to $\mathcal{O}(\frac{1}{l^2})$, as we show in Theorem 2.

THEOREM 2. Suppose r^* is the optimal revenue for the problem (14) with continuous price set and r^l is the optimal revenue for its discrete price setting counterpart (13). If $\forall u, i$, $D_i^u(p)$ is twice continuously differentiable on $[\underline{p}, \bar{p}]$, and there exist constants $C_{cont1}, C_{cont2} > 0$ s.t. $\forall p \in [\underline{p}, \bar{p}]$,

$$\left| \frac{dD_i^u(p)}{dp} \right| \leq C_{cont1}, \quad \left| \frac{d^2D_i^u(p)}{dp^2} \right| \leq C_{cont2},$$

³Note that problem (13) and (14) are for the empirical model while not the one with distribution $F(\mathbf{X})$.

then we have

$$r^* - r^l \leq \mathcal{O}\left(\frac{1}{l^2}\right).$$

All the proofs are deferred to Appendix B.

4 NUMERICAL EXPERIMENT

We demonstrate the benefits of imposing fairness constraints on both synthetic data and real-world data. We evaluate the performance of our model by customer surplus, social welfare, and the Gini index.

Here, we only illustrate the results under linear demand with TVD constraint and for a group size of 2 while the results under other demand models in Table 1, with the EMD constraint and for multiple groups are similar in general, and are deferred to Appendix C.

4.1 Model Evaluation

We measure the performance of our model based on the trade-off between the tightness of the fairness constraint and several evaluation metrics.

We illustrate the impact of our proposed model by customer surplus and social welfare defined in (3) and (4). The higher the customer surplus and social welfare, the better the societal well-being.

In addition, we assess our algorithm's fairness using the Gini index, a widely employed metric in economics [21]. The Gini index serves as an indicator of inequality in wealth distribution, measuring the deviation of a country's wealth or income distribution from perfect equality. In our context, α_j^u denotes the probability density for each admissible price p_j in the group u under our pricing policy. We define the Gini coefficients for each admissible price p_j as follows:

$$G_j := \frac{\sum_{u=1}^k \sum_{v=1}^k |\alpha_j^u - \alpha_j^v|}{\sum_{u=1}^k \alpha_j^u + \sum_{v=1}^k \alpha_j^v}, \quad \forall j \in [l]. \quad (15)$$

The Gini index is calculated as the sum of all Gini coefficients, i.e., $G := \sum_{j=1}^l G_j$. A lower Gini index suggests a fairer pricing policy.

4.2 Simulation

The simulation experiments are designed mainly to test the performance of our revenue maximization problem with group fairness constraint under different fairness degrees δ and the impact of price discretization.

The simulation study is under the assumption of the linear demand function, that is, the customer valuation is simulated from a uniform distribution.

4.2.1 Data. We simulate 2000 customers' samples, each with a two-dimensional feature vector $\tilde{x} = [x, s]$, with one normal feature $x \sim \text{Gaussian}(0, 1)$, and the sensitive attribute $s \in \{0, 1\}$. In accordance with the sensitive feature, we split all customers into two groups with 1000 customers in each group. We assume price $p \in [1, 3]$ with 10 discrete realizations, i.e., $\{1.0, 1.2, 1.4, 1.6, 1.8, 2.0, 2.2, 2.4, 2.6, 2.8, 3.0\}$.

We assume the valuation function follows uniform distribution $V(\tilde{x}) \sim U(0, c(\tilde{x}))$, where $c(\tilde{x}) = (c_1, c_2, c_3)(1, x, s)^T$ with

$(c_1, c_2, c_3) \in [1, 2] \times [1, 2] \times \{0, 1\}$. Then, we result in the linear demand model defined in Table 1, where $a(\tilde{x}) = (a_1, a_2, a_3)(1, x, s)^T$ with $(a_1, a_2, a_3) \in [1, 2] \times [1, 2] \times \{0, 1\}$.

4.2.2 The impact of fairness degree. Figure 1 illustrates how the optimal objective from the constrained linear programming problem (13) under different values of fairness degree δ . X-axis represents fairness degree δ and the y-axis represents various objectives and evaluation metrics. Under TVD constraint, Figure 1 shows that when we decrease δ , the revenue decreases which is due to the stricter constraint on price fairness.

From the results under different parameter settings, we remark that when the sensitive features are more significant than the normal features, we can obtain higher customer surplus, social welfare, and lower Gini index as we impose stricter fairness regulation. This result demonstrates one of the benefits of imposing group fairness constraints that fairness can improve social welfare.

To see how the objective and metrics change within each group, we look at Figure 2a. This figure demonstrates how the optimal objective from the constrained linear programming problem (13) under different values of fairness degree δ for each group respectively. X-axis represents fairness degree δ and the y-axis represents various objectives and evaluation metrics. We can interpret from the Figure 2a that when we decrease the value of δ , group 2 gets higher welfare and customer surplus while group 1 behaves the opposite. From this interesting observation, we can tell that group 2 benefits from imposing fairness constraints.

Also, this analysis provides insight for decision-makers to select the proper fairness degree δ to balance between the decrease in revenue and the increase in social welfare.

4.2.3 Effect of price discretization. To see whether offering more price options is beneficial and meaningful, we evenly discretize prices from 1 to 3, with discrete price sets ranging from 2 to 20.

Figure 3 shows how the revenue, customer surplus, and welfare change respectively as we increase the size of the price set. From the results shown in Figure 3, we can see that the optimal revenue converges faster than $\mathcal{O}(\frac{1}{l})$ rate when we increase the size of the discrete price set, which is verified in Theorem 2. Moreover, we can tell that under fairness constraints with TVD, the further we discretize the prices, the higher the revenue, consumer surplus, and social welfare. However, when the size of the discrete price set reaches 8, the increase in the revenue, consumer surplus, and social welfare become less significant and will cause more cost for creating price labels. Thus, we conclude that it is appealing to set the size of the price set moderately while aggressive discretization may not be that rewarding.

Figure 3c shows how the inverse of the gap between the objective of the discrete price set and the continuous price set changes with respect to the number of prices. We use price set size $l = 30$ to approximate the continuous price set since the change of objective for $l > 8$ is not significant. We could tell that the gap converges at a rate faster than linear, which is proved in Theorem 2.

4.3 Real Data

We use a real-world data set from the e-commerce company JD.com [37] to demonstrate the benefits of our model.




Figure 1: Impact of fairness under TVD constraint with linear demand using synthetic data with model parameters $[a_0, a_1, a_2] = [1, 1.4, 3]$, $[c_0, c_1, c_2] = [1, 1, 1]$.




Figure 2: Comparison of the impact of fairness degree with linear demand between two groups




Figure 3: Impact of price discretization on various metrics under TVD constraint using synthetic data with model parameters $[a_0, a_1, a_2] = [1, 2, 2]$, $[c_0, c_1, c_2] = [1, 2, 2]$

4.3.1 Dataset Introduction. The datasets provided by JD.com grant a comprehensive view of the activities related to all SKUs within an anonymous consumable category during March 2018. This category could be beauty care (e.g., face moisturizers) or men’s grooming (e.g., electric shavers). Due to confidentiality, the specific category

remains undisclosed. We employ the transaction-level data labeled as i) skus, ii) users, and iii) orders. The data description and basic statistics can be found in Appendix C.3.

In this experiment, we use one top-selling product, including 2854 customer purchasing records. We select six normal features

(i.e., user level, age, marital status, education, city level, purchase power) and one sensitive feature (i.e., gender). We assign group membership based on the sensitive attribute (i.e., gender), with 1931 in group 1 (i.e., male) and 923 in group 2 (i.e., female). The price set is scaled to be $\{1.0, 1.2, 1.4, 1.6, 1.8, 2.0, 2.2, 2.4, 2.6, 2.8, 3.0\}$.

4.3.2 Preliminary Analysis. We have checked that the distribution of the normal features, conditioned on the sensitive attributes, is similar. To verify this, we perform a Kolmogorov-Smirnov (KS) test on the distribution of each normal feature, such as age, conditioned on the sensitive attribute (gender). Table 2 presents the p-values obtained from the KS test for each normal feature.

Normal Feature	p-value from KS test
User Level	0.85
Age	0.87
Marital Status	0.92
Education	0.88
City Level	0.91
Purchase Power	0.82

Table 2: p-values from the Kolmogorov-Smirnov test for normal features conditioned on the sensitive attribute (gender)

In Table 2, the p-values from the KS test are all larger than a typical significance level (e.g., 0.05). This indicates that we cannot reject the null hypothesis that the distributions of the normal features are the same between the two groups, conditioned on the sensitive attribute (gender). As a result, we can conclude that the distributions of the normal features, conditioned on the sensitive attributes, are similar. In Figure 5a, we can tell that the original pricing distributions between two groups are biased w.r.t. sensitive attribute (i.e., gender). Group 2 (i.e., female) has been charged much higher prices compared with group 1 (i.e., male). Given that the distribution of normal features is similar between the two groups, it is necessary to implement our fairness algorithm to correct this bias and improve social well-being.

4.3.3 Analysis on the impact of fairness degree. The result is consistent with our simulation study on the impact of fairness. Figure 4 shows how the revenue, customer surplus, and social welfare change respectively as we increase the size of the price set under TVD fairness measures, respectively. From Figure 4, we can tell that as we impose less fairness degree (i.e., larger δ), we get lower customer surplus and social welfare as well as a larger Gini index.

From the preliminary analysis in Figure 5a, we find that the female group has a higher valuation on the product and gets charged a higher price compared with the male group. Thus, we treat the protected group as the group having a higher valuation for a specific product, which is female in this case. We can tell from the figure 2b that when we decrease the value of δ , group 2 (i.e., female) gets higher welfare and customer surplus while that group 1 (i.e., male) behaves the opposite. From this interesting observation, we can tell that the protected group (i.e., female) benefits from imposing a fairness constraint, indicating that the protected group gets better off when we impose a price parity constraint.

In Figure 6, we aim to demonstrate the influence of fairness on social welfare and the Gini index under various pricing constraints.

The x-axis represents the fairness degree δ , while the y-axis indicates social welfare (left plot) and the Gini index (right plot). As illustrated in the left plot, our proposed model consistently achieves higher social welfare compared to other baseline models, namely those with mean price constraints and single price constraints. In terms of the Gini index, as shown in the right plot, our model generally exhibits lower values than the mean price constraint model, indicating a more equitable distribution of prices. However, the single price constraint model exhibits a distinct behavior, with its Gini index values not consistently lower than those of our model. Overall, our model outperforms the baseline models in terms of social welfare while maintaining a relatively equitable price distribution, as evidenced by the Gini index.

4.3.4 Analysis on Fair Pricing Policy. To demonstrate the advantages of our proposed fairness constraint, we conduct a comparative analysis of pricing policies derived from our fair algorithm and various baseline models.

The pricing policies for two groups under TVD, single-price constraint, mean-price constraint, and the unconstrained model are presented in Figure 5. The x-axis represents a range of 10 prices from 1 to 3, while the y-axis depicts the associated price distributions, which represent the probability of customers being assigned to a specific price. The light blue bars correspond to the pricing policy for Group 1, and the navy blue bars display the policy for Group 2.

As illustrated in Figure 5d, the optimal policy derived from the mean-price baseline shows significant variation in price distributions for both groups, while their average prices remain relatively close. In contrast, the optimal policy derived from the single-price baseline, as shown in Figure 5d, concentrates on a single price, leading to a less diverse policy that may reduce revenue without reaping the benefits of personalization.

Our results indicate that the pricing distribution achieved through our model is more equitable and fairer than the policies obtained through the mean-price constraint baseline and the unconstrained baseline. Furthermore, our approach maintains policy diversity while minimizing revenue loss.

From these results, we can infer that the optimal pricing distribution between the two groups obtained by our model is more closely aligned (i.e., fairer) than the mean-price constraint baseline and the unconstrained baseline without the loss of diversity.

In summary, our numerical studies suggest that for society as a whole, implementing fairness regulations using statistical parity is more advantageous than not doing so.

5 CONCLUSION

In this paper, we study the personalized pricing problem for different groups under fairness constraints measured by total variation distance and earth mover distance. We propose a stochastic programming formulation that maximizes the revenue with statistical parity constraint. Under the discrete price setting, we reformulate this problem as a linear program that can be solved efficiently. To bridge the gap between the discrete and continuous price settings, theoretically, we prove an $\mathcal{O}(\frac{1}{l})$ gap between the optimal revenues with continuous and discrete price sets of size l under fairness constraints measured by total variation distance (TVD) and earth




Figure 4: Impact of fairness under TVD constraint with linear demand using JD.com data




Figure 5: Price distribution from the JD.com data; Comparison of Pricing policy under TVD, single-price constraint, mean-price constraint, and unconstrained model using JD.com data

mover's distance (EMD). Furthermore, under some mild assumptions, we improve this bound to $\mathcal{O}(\frac{1}{\sqrt{p}})$ for the TVD constraint. We implement our model on both synthetic data and real-world data. We demonstrate the benefits of imposing fairness constraints, including higher customer surplus, higher social welfare, lower Gini index, and superior pricing policy compared with several baselines. Our results also provide managerial insights on setting a proper fairness degree as well as an appropriate size of discrete price set.

There are various potential directions for future research. Firstly, alternative distance metrics, such as $\mathcal{H}$ -divergence or KL-divergence [35], could be explored. Secondly, while our current theoretical results provide an upper bound on the gap between optimal revenue with continuous and discrete price sets, the numerical results in Figure 3c suggest that the convergence rate of the gap may be faster than $\mathcal{O}(\frac{1}{\sqrt{p}})$. Investigating the lower bound of this gap would be intriguing. Furthermore, future work could examine other fairness




Figure 6: Comparison of the Impact of Fairness on Linear Demand with TVD Constraint using JD.com Data

notions, extend the framework to accommodate multiple products, and consider the interplay between fairness and other societal objectives like environmental sustainability or customer satisfaction. This could ultimately lead to the creation of more comprehensive and ethically responsible pricing algorithms.

REFERENCES

- [1] 2008. *Linear Programming Algorithms*. Springer Berlin Heidelberg, Berlin, Heidelberg, 71–97. https://doi.org/10.1007/978-3-540-71844-4_4
- [2] Massachusetts Post Audit. 1997. Shear discrimination: Bureau finds wide price bias against women at Massachusetts hair salons despite anti-discrimination laws.
- [3] Gah-Yi Ban and N Bora Keskin. 2017. Personalized dynamic pricing with machine learning (may 23, 2017).
- [4] Richard E Barlow, Albert W Marshall, and Frank Proschan. 1963. Properties of probability distributions with monotone hazard rate. *The Annals of Mathematical Statistics* (1963), 375–389.
- [5] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning*. fairmlbook.org. <http://www.fairmlbook.org>.
- [6] William E. Becker. 1999. Carl Shapiro and Hal R. Varian. Information Rules: A Strategic Guide to the Network Economy. Boston, Massachusetts: Harvard Business School Press, 1999. x + 352 pp. *The Journal of Economic Education* 30, 2 (1999), 189–190. <https://doi.org/10.1080/00220489909595956> arXiv:<https://doi.org/10.1080/00220489909595956>
- [7] Max Biggs, Wei Sun, and Markus Ettl. 2021. Model Distillation for Revenue Optimization: Interpretable Personalized Pricing. In *International Conference on Machine Learning*. PMLR, 946–956.
- [8] Alexander W Butler, Erik J Mayer, and James P Weston. 2022. Racial Disparities in the Auto Loan Market. *The Review of Financial Studies* (05 2022). <https://doi.org/10.1093/rfs/hhac029> arXiv:<https://academic.oup.com/rfs/advance-article-pdf/doi/10.1093/rfs/hhac029/44179813/hhac029.pdf> hhac029.
- [9] Xi Chen, Zachary Owen, Clark Pixon, and David Simchi-Levi. 2021. A statistical learning approach to personalization in revenue management. *Management Science* (2021).
- [10] Silvia Chiappa. 2019. Path-specific counterfactual fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 7801–7808.
- [11] Stephanie Clifford. 2012. Shopper alert: Price may drop for you alone. *New York Times* 9 (2012).
- [12] Maxime Cohen, Adam Elmachtoub, and Xiao Lei. 2019. Pricing with Fairness. (Sept. 2019). <https://doi.org/10.2139/ssrn.345928>
- [13] Maxime C. Cohen, Adam N. Elmachtoub, and Xiao Lei. 2022. Price Discrimination with Fairness Constraints. *Management Science* 0, 0 (March 2022), null. <https://doi.org/10.1287/mnsc.2022.4317> arXiv:<https://doi.org/10.1287/mnsc.2022.4317>
- [14] Maxime C Cohen, Ilan Lobel, and Renato Paes Leme. 2020. Feature-based dynamic pricing. *Management Science* 66, 11 (2020), 4921–4943.
- [15] Pablo de Andrés, Ricardo Gimeno, and Ruth Mateos de Cabo. 2021. The gender gap in bank credit access. *Journal of Corporate Finance* 71 (2021), 101782. <https://doi.org/10.1016/j.jcorpfin.2020.101782>
- [16] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Rich Zemel. 2011. Fairness Through Awareness. (Nov. 2011).
- [17] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*. 214–226.
- [18] Anne D’Innocenzio. 2017. Neiman marcus focuses on exclusives, personalized offers; ends merger talks. <https://www.usatoday.com/story/money/business/2017/06/13/neiman-marcus-focuses-exclusives-personalized-offers/102814962>.
- [19] Adam N Elmachtoub, Vishal Gupta, and Michael L Hamilton. 2021. The value of personalized pricing. *Management Science* (2021).
- [20] Lu Fang and Henry J. Munneke. 2020. Gender Equality in Mortgage Lending. *Real Estate Economics* 48, 4 (2020), 957–1003. <https://doi.org/10.1111/1540-6229.12198> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/1540-6229.12198>
- [21] Frank A. Farris. 2010. The Gini Index and Measures of Inequality. *The American Mathematical Monthly* 117, 10 (2010), 851–864. <http://www.jstor.org/stable/10.4169/000298910x523344>
- [22] Chris Gee. 2018. Fair pricing in financial services. <https://www.fca.org.uk/publication/discussion/dp18-09.pdf>.
- [23] Joshua A Gerlick and Stephan M Liozu. 2020. Ethical and legal considerations of artificial intelligence and algorithmic decision-making in personalized pricing. *Journal of Revenue and Pricing Management* (2020), 1–14.
- [24] Talia B Gillis. 2020. False dreams of algorithmic fairness: The case of credit pricing. Available at SSRN 3571266 (2020).
- [25] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems* 29 (2016), 3315–3323.
- [26] Hoda Heidari and Andreas Krause. 2018. Preventing Disparate Treatment in Sequential Decision Making. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 2248–2254. <https://doi.org/10.24963/ijcai.2018/311>
- [27] White House. 2015. Big data and differential pricing. https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/docs/Big_Data_Report_Nonembargo_v2.pdf.
- [28] Nathan Kallus and Angela Zhou. 2021. Fairness, Welfare, and Equity in Personalized Pricing. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT ’21). Association for Computing Machinery, New York, NY, USA, 296–314. <https://doi.org/10.1145/3442188.3445895>
- [29] Nathan Kallus and Angela Zhou. 2021. Fairness, Welfare, and Equity in Personalized Pricing. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 296–314.
- [30] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International Conference on Machine Learning*. PMLR, 2564–2572.
- [31] Niki Kilbertus, Mateo Rojas-Carulla, Giambattista Parascandolo, Moritz Hardt, Dominik Janzing, and Bernhard Schölkopf. 2017. Avoiding discrimination through causal reasoning. *Advances in Neural Information Processing Systems* 30 (2017).
- [32] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual Fairness. *Advances in Neural Information Processing Systems* 30 (2017).
- [33] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- [34] Andres Munoz and Sergei Vassilvitskii. 2017. Revenue Optimization with Approximate Bid Predictions. *Advances in Neural Information Processing Systems* 30 (2017), 1858–1866.
- [35] Yves Rychener, Bahar Taskesen, and Daniel Kuhn. 2023. Metrizing Fairness. arXiv:2205.15049 [cs.LG]
- [36] Carl Shapiro, Shapiro Carl, Hal R Varian, et al. 1998. *Information rules: a strategic guide to the network economy*. Harvard Business Press.

- [37] Max Shen, Christopher S. Tang, Di Wu, Rong Yuan, and Wei Zhou. 2020. JD.com: Transaction-Level Data for the 2020 MSOM Data Driven Research Challenge. *Manufacturing & Service Operations Management* 0, 0 (2020), null. <https://doi.org/10.1287/msom.2020.0900> arXiv:<https://doi.org/10.1287/msom.2020.0900>
- [38] Omer Tamuz. 2013. A lower bound on seller revenue in single buyer monopoly auctions. *Operations Research Letters* 41, 5 (2013), 474–476.
- [39] Brad Tuttle. 2013. Flight prices to get personal? airfares could vary depending on who is traveling. <http://business.time.com/2013/03/05/flight-prices-to-get-personal-airfares-could-vary-depending-on-who-is-traveling>.
- [40] Sahil Verma and Julia Rubin. 2018. Fairness Definitions Explained. In *Proceedings of the International Workshop on Software Fairness* (Gothenburg, Sweden) (FairWare '18). Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3194770.3194776>
- [41] Brian Wallheimer. 2018. Are you ready for personalized pricing? <https://review.chicagobooth.edu/marketing/2018/article/are-you-ready-personalized-pricing>.
- [42] Renzhe Xu, Peng Cui, Kun Kuang, Bo Li, Linjun Zhou, Zheyang Shen, and Wei Cui. 2020. Algorithmic Decision Making with Conditional Fairness. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2125–2135.
- [43] Ruofei Zhao Ya'acov Ritov, Yuekai Sun. 2017. On conditional parity as a notion of non-discrimination in machine learning. (June 2017).
- [44] Mikhail Yurochkin and Yuekai Sun. 2020. Sensei: Sensitive set invariance for enforcing individual fairness. *arXiv preprint arXiv:2006.14168* (2020).

Received 6 February 2023; revised 6 April 2023; accepted 10 May 2023