



Oh the Prices You'll See: Designing a Fair Exchange System to Mitigate Personalized Pricing

Aditya Karan*
University of Illinois
Urbana-Champaign
Siebel School of Computing and Data
Science
Urbana, Illinois, USA
karan2@illinois.edu

Naina Balepur*
University of Illinois
Urbana-Champaign
Siebel School of Computing and Data
Science
Urbana, Illinois, USA
nainab2@illinois.edu

Hari Sundaram
University of Illinois
Urbana-Champaign
Siebel School of Computing and Data
Science
Urbana, Illinois, USA
hs1@illinois.edu

Abstract

Many online marketplaces personalize prices based on consumer attributes. Since these prices are private, consumers may be unaware that they have spent more on a good than the lowest possible price, and cannot easily take action to pay less. In this paper, we introduce a fairness-centered exchange system that takes advantage of personalized pricing, while still allowing consumers to individually benefit. Our system produces a matching of consumers to promote trading: the lower-paying consumer buys the good for the higher-paying consumer for some fee. We explore various modeling choices and fairness targets to determine which schema will leave consumers best off, while also earning revenue for the system itself. We show that when consumers individually negotiate the transaction price, and our fairness objective is to minimize mean net cost, we are able to achieve the most fair outcomes. Conversely, when transaction prices are centrally set, consumers are often unwilling to transact. When price dispersion (or range) is high, the system can reduce the mean net cost to each individual by 66%, or the mean net cost to a group by 69%. We find that a high dispersion of original prices is necessary for our system to be viable. Higher dispersion can actually lead to decreased net price paid by consumers, and act as a check against extreme personalization, increasing seller accountability. Our results provide theoretical evidence that such a system could improve fairness for consumers while sustaining itself financially.

CCS Concepts

• **Software and its engineering** → *E-commerce infrastructure*; • **Information systems** → *Data exchange; Collaborative and social computing systems and tools*; • **Theory of computation** → *Computational pricing and auctions*; • **Social and professional topics** → *Socio-technical systems*.

Keywords

System Design, Fairness, Personalized Pricing

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License. *FAccT '25, Athens, Greece*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1482-5/25/06
<https://doi.org/10.1145/3715275.3732076>

ACM Reference Format:

Aditya Karan, Naina Balepur, and Hari Sundaram. 2025. Oh the Prices You'll See: Designing a Fair Exchange System to Mitigate Personalized Pricing. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*, June 23–26, 2025, Athens, Greece. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3715275.3732076>

1 Introduction

Suppose you and your friend are interested in buying the same pair of shoes online from the same website. You see the shoes listed for \$60, while your friend sees \$50 due to behavioral profiling [21, 28].

If you knew about this price discrepancy, you could ask your friend to buy the shoes for you, paying them back and throwing in a few extra dollars so you can both buy the shoes at a lower net price.

Many such online marketplaces employ elements of personalized pricing [24]. Differential pricing can vary with protected attributes [49], and fairness to consumers can suffer when personalized pricing is employed [27, 40]. While many consumers are aware of price personalization, it is difficult to take money-saving actions without knowing what other consumers are paying.

Fairer prices could be achieved either by sellers implementing fair pricing algorithms or via consumers' more active involvement in improving pricing outcomes. On the seller's side, prior works examine fairer pricing algorithms [7, 12, 19, 47, 54]. These methods still attempt to build revenue maximizing, sometimes personalized, pricing algorithms while considering some notion of fairness. However, there is often not a way to enforce seller use of fair pricing algorithms. Thus, developing tools that promote consumer agency is key. Using participatory design methods, Richards et al. [43] find that involving consumers in price-setting and negotiation leads to increased perceptions of fairness. Other prior work advises users to change their behaviors, and therefore receive better personalized prices [32]. Another solution in online setting is coupon trading, which has been investigated as one means for consumers to get lower prices [31]. However, none of these methods directly tackle achieve fairer pricing—allowing consumers to directly financially benefit without changing underlying behavior. A fairness-aware trading platform may be able to simultaneously reduce the impact of personalized pricing and improve fairness.

In this work, we design a fairness-centered exchange system that allows consumers to trade to reduce the price they pay for a good online. As part of this system, we design matching and transaction-price setting procedures that aim to improve fairness. Figure 1 depicts how the system, market, and consumers interact.

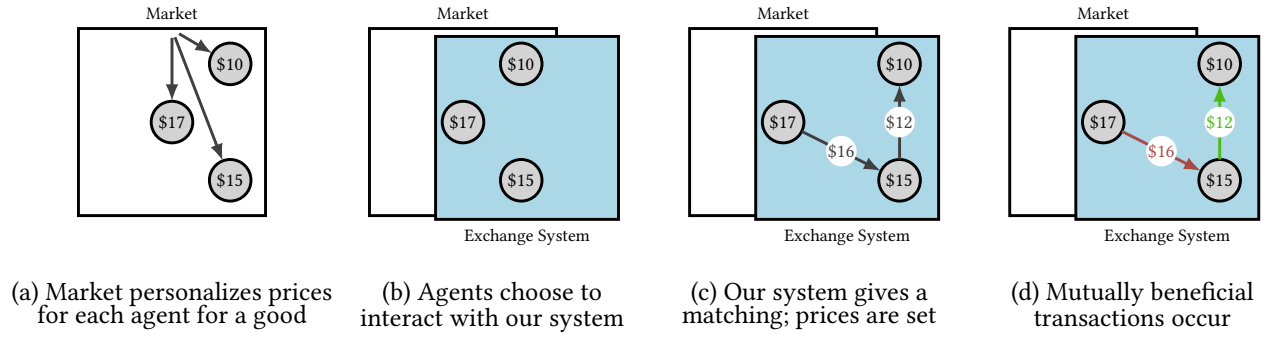


Figure 1: The relationship between the market (white square), the system (blue square), and consumers (gray circles). Consumers enter the market and are offered prices (a). They decide to participate in the exchange system (b). The system assigns a matching, and prices for these transactions are assigned or decided upon (c). The arrow direction represents the transfer of money. Of these matched transactions, only those that are mutually beneficial to both agents occur (d). This is determined by each agent's utility function.

To achieve our goal of designing an effective fairness-improving system, we simulate various design choices. With a given matching, we test two transaction price-setting procedures: centralized, where the system sets price, and decentralized, where consumers negotiate (*RQ 1*). Our system is intended to improve fairness; we consider four fairness targets for each procedure. The four fairness targets cover two objectives (the value measured: mean and variance of cost), and scopes (who it is measured for: individuals and groups). We investigate which fairness target is most feasible, or achievable, in our system (*RQ 2*). To ensure financial viability, we explore how the system's revenue is affected by the number of consumers, the cut the system takes from each transaction, and the dispersion (or range) of prices initially offered by the market (*RQ 3*). To conclude we examine a case study on an empirical pricing distribution modeled in prior work.

We find that decentralized negotiations between consumers and minimizing the average net price are fairness maximizing. We are able to decrease the average price paid by 66% when compared to the market without trades, under a highly dispersed pricing algorithm (producing a wide range of prices). Minimizing average net price paid per group leads to a 69% decrease in the metric, compared to a 66% decrease for average net price paid per individual. Minimizing standard deviation is unsuccessful, leading to increases in the metric. We find that as dispersion (or range) of prices increases, the exchange system acts as a more effective counterweight to unfair pricing. While in the high dispersion case we see a 66% reduction in average price paid, the medium and low dispersion cases are a 15% and 0% reduction respectively. As we increase the cut taken by the system, the system earns more revenue; however, for a sufficiently high percent taken, system revenue falls as consumers find fees too high. Our contributions are as follow:

Exchange system design: We design an exchange system that could feasibly be implemented, financially sustain itself, and does not require the cooperation of the market. Prior work in this area has found personalized pricing in different markets [5, 21, 28], and some has proposed fair pricing algorithms [4, 19, 52–54]. However, none to our knowledge has proposed an exchange system to address unfair personalized pricing. In this work we close that

gap—exploring design choices and fairness targets that will lead to a financially self-sustaining, price-decreasing system.

Opportunities for personalized pricing: We find that even though personalized pricing can cause monetary loss to consumers, when properly taken advantage of, it can provide a money-saving opportunity. While prior work notes personalized pricing can help some consumers by making certain goods more affordable to them [16], we design a system that helps all consumers. We show that, surprisingly, highly dispersed personalized pricing leads to lower costs for consumers *and* higher revenue to the system. This finding emphasizes the strengths of our exchange system design: we do not require sellers to cooperate by fairly pricing goods. In fact, when they don't, it can be beneficial to consumers and creates a measure of accountability to extreme, unfair personalization.

2 Related Work

Fairness and pricing: Various methods and metrics for fair machine learning have been proposed, in particular to deal with challenges of group fairness [1, 2, 9, 50, 51]. In the pricing domain, unfair pricing and ad delivery caused by online behaviors have been found by several studies [5, 8, 21, 28] and can result in structural inequity [60].

This work examines personalized pricing, which stands in contrast to flat or fixed pricing. While personalized pricing *could* lead to consumer benefit [10], it may also cause certain consumers to pay much more than others [15], and lead to a perceived feeling of unfairness [61]. One solution is to redesign pricing algorithms for fairness. Seminal work on fair pricing [23, 44] examines what fair pricing is and how consumers may react to various pricing algorithms by firms. Kallus and Zhou [27] examine how different markets and concerns (e.g., information asymmetry) inform which fairness criteria to consider. Grari et al. [19] study fair pricing under adverse selection, balancing actuarial risk with a demographic parity or equalized odds constraint. Xu et al. [54] explore how imposing restrictions on the degree of personalized pricing (*i.e.*, the price of an object cannot vary more than $x\%$) can be customized to balance the needs of both buyers and sellers. The goal of fair

pricing literature is to create fair algorithms that a seller can use while still achieving high profits.

While this approach *could* achieve lower prices for consumers *if deployed*, this assumes the ability and willingness of a seller to do so. To mitigate this issue, we design a solution to unfair pricing without directly requiring cooperation from the seller. Some work has taken this approach—counterfactual fairness [32] could allow an individual to achieve better pricing without direct access to the underlying pricing system. However, this requires both an accurate causal model and toggles that individuals can act on [29]. On the collective action side, Hardt et al. [22] quantified the effect of coordinated feature changes against an algorithmic system. In our work, collaboration does not aim to change the direct outcome of an algorithmic system, but rather to adjust final outcomes for individuals. Our work does not require a causal model or online behavior change for individuals to benefit. Rather, we focus on consumers who are incentivized by monetary reward to share information. We specifically are interested in cases where the goods are identical and purchased on online marketplaces, where potentially unfair differential pricing can occur without users' knowledge [36, 39].

Federated learning: In federated learning, data is indirectly “shared” via model updates and aggregated (e.g., FedAvg [38]). There is a very large existing body of work investigating fairness and trying to achieve it, on both group [15, 18, 25, 34, 35, 37, 55, 59] and individual [33, 56] levels. In the fair federated learning framework, it is already assumed that individual entities want to participate, as the outcomes are clear. Participation can improve both accuracy in the underlying models as well as some notion of fairness.

In our work, we focus on whether we can incentivize rational agents who act purely selfishly to achieve some notion of fairness. Donahue and Kleinberg [13] extend the federated learning context to consider whether agents should participate in a shared model or rely on only their local information. However, this work differs from ours as our objective is not one of model performance. They further investigate [14] egalitarian and proportional fairness in the context of these model-sharing games. Salehi et al. [46] develop a model-sharing architecture for agents' mental models but do not explicitly consider fairness.

Marketplace mechanisms: Shapley and Shubik [48] and Roth and Sotomayor [45] both examine the assignment game — a two-sided matching with money. In this work, we are interested in situations where everyone has access to the same good at different prices, rather than valuing goods at different prices. Jagadeesan and Teytelboym [26] look at how some markets give rise to universal pricing, while others employ personalized pricing. Babaioff et al. [6] and Branzei et al. [11] design auctions to limit gains from post auction dealings, similar to those we introduce here.

3 Problem Statement

Consider a marketplace \mathcal{M} that offers one type of good g to consumers V , where $N = |V|$. Each consumer $v \in V$ requires exactly one unit of good g ; the supply of g is finite but sufficient to satisfy V . Each consumer $v \in V$ has a vector of attributes r_v , a resource constraint k , and a private utility function f_v . The consumers in V can be partitioned into non-overlapping groups based on r_v .

Pricing algorithm \mathcal{A} takes a vector of properties of the consumer r_v as input, and outputs a personalized price p_v for consumer v . Resulting prices may have varying degrees of personalization—we measure this using the *dispersion* of prices, δ , our measure of pricing spread. In this work, we design an exchange system \mathcal{S} that takes advantage of personalized pricing so consumers profit. System \mathcal{S} pairs up consumers via matching process \mathcal{P} , which outputs a set of pairwise consumer interactions \mathcal{J} . Let $j_{uv} \in \mathcal{J}$ be the (directed) interaction parameterized by consumers u and v , where $p_u > p_v$. In an interaction j_{uv} , consumer u will pay some m dollars to consumer v for good g . The system \mathcal{S} then takes some fraction γ of this m to sustain the trading ecosystem, so the payment v will receive for this transaction is $(1 - \gamma)m$. The interactions in \mathcal{J} are proposed to consumers by the matching process \mathcal{P} , but all interactions need not occur. An interaction is executed if both consumers u and v benefit according to their utility functions f_u and f_v .

We frame the matching process \mathcal{P} as a network problem. Our consumers exist in a directed network $G = (V, E)$ where $(u, v) \in E$ if $p_u > p_v$. Interactions \mathcal{J} are matched directed edges; each consumer v has at most k matched incoming edges. The outgoing edge count is capped at 1, since each consumer can only buy the good g once. This resource constraint k represents how much time a consumer is willing to spend on this system. If an interaction j_{uv} is executed, the edge (u, v) has been transacted on. We call u the buyer and v intermediary; one consumer might be matched as both an intermediary and a buyer on different transactions.

In this paper, we examine which design of exchange system \mathcal{S} will maximally improve our fairness targets \mathcal{F} on market \mathcal{M} , and under what circumstances it is financially feasible to maintain. Specifically we ask:

- RQ 1:** Given a fairness target \mathcal{F} , which of the following methods to set transaction price m will maximize fairness: centralized setting or individual negotiation?
- RQ 2:** In the context of our exchange system \mathcal{S} , are some fairness considerations (i.e., definitions of \mathcal{F}) more feasible to achieve than others? In particular, how does varying \mathcal{F} in scope and objective affect feasibility?
- RQ 3:** How does revenue to the exchange system \mathcal{S} change with the number of consumers N , cut γ taken by the system, and dispersion δ of the pricing algorithm \mathcal{A} ?

4 Fairness

In **RQ 2** we ask whether some fairness criteria \mathcal{F} are more feasible than others to achieve. Inspired by Kallus and Zhou [27], we develop our notions of fairness to be suited to this setting. We started by considering what the “ideal” scenario would be for all consumers. We determined the ideal fair outcome would result in all consumers paying the *same, lowest* price for the good. Hence, we seek to minimize both average price and the standard deviation in prices paid by consumers. However, these ideal outcomes might not be possible in our exchange system. Thus, we demonstrate lower bounds for mean (Theorem 1) and standard deviation (Claim 1) which can be found in Subsection 5.5. We measure the “feasibility” of our fairness metrics by considering how close each mean and standard deviation come to the ideals.

Table 1: Our fairness outcome measures. All measures rely on ω_u , the net cost to consumers u . μ_g represents the average net cost to consumers u in group g . $|\mathcal{G}|$ is the number of groups.

	INDIVIDUAL	GROUP
MEAN	$\mu_I = \frac{\sum_{u \in V} \omega_u}{N}$	$\mu_G = \frac{\sum_{g \in \mathcal{G}} \mu_g}{ \mathcal{G} }$
S.D.	$\sigma_I = \sqrt{\frac{\sum_{u \in V} (\omega_u - \mu_I)^2}{N}}$	$\sigma_G = \sqrt{\frac{\sum_{g \in \mathcal{G}} (\mu_g - \mu_G)^2}{ \mathcal{G} }}$

Thus, we vary objective and scope for our fairness definitions, *i.e.*, *what* we measure, and *who* we measure it for. All objectives take as input the net cost incurred by each consumer from participating in system \mathcal{S} , assuming that each consumer purchases exactly one unit of good g . For example, if consumer u is offered price p_u by the market, but pays m dollars to consumer v to buy good g for price p_v , then consumer u 's net cost is m . Consumer v 's profit from selling the ticket is $(m(1 - \gamma) - p_v)$ and therefore their net cost for their own ticket is $p_v - (m(1 - \gamma) - p_v)$. As a reminder, consumers in our system can engage in up to k transactions as intermediary. We denote ω_u as the net cost to consumer u , where $\omega_u < 0$ implies monetary gain to u . We define two objectives (mean and standard deviation) and two scopes (individual and group-level). We define a group as a collection of individuals who share a demographic attribute (in our case, r_v). We call this set of groups \mathcal{G} . This gives us four different fairness measures; we present them in Table 1.

We seek to minimize these values. In minimizing individual mean, we want all consumers to minimize the price they're paying (some may even earn money from \mathcal{S}). Minimizing individual standard deviation captures that we want our system's benefit to be distributed equally among consumers; no single consumer should benefit while others do not. In the group fairness cases, we desire similar outcomes, but distill each group by taking the mean over individuals in the group. We recognize that these definitions of fairness are neither standard nor exhaustive. We considered more standard fairness metrics such as demographic parity, predicted parity, and equalized odds [1, 2, 9, 17, 30, 32, 57, 58]. However, we decided that fairness would best be captured in this application by metrics designed from our ideal scenario (where all individuals get the lowest price). We vary our fairness target over these definitions to test **RQ 2**.

5 The Model

5.1 Model overview

We simulate our proposed exchange system via rational agents. We begin with consumers V who desire exactly one unit of good g on market \mathcal{M} . Each consumer $v \in V$ is initialized with attribute vector r_v , utility function f_v , and price p_v for good g according to pricing algorithm \mathcal{A} .

The matching process \mathcal{P} outputs a set of interactions \mathcal{J} between consumers. The transaction price m of each interaction is then set in either a centralized or decentralized fashion (*i.e.*, either set by the system or via consumer negotiations). Given a matched edge (u, v) , buyer u and intermediary v exchange money for good g if u and v both have positive utility for the transaction and are within their resource constraints. Once all consumers make transaction

Algorithm 1: Model overview

Data: V the set of consumers
Result: G the network after trades, including transaction prices

// PRICING ALGORITHM

1 **for** $u \in V$ **do**
2 $p_u \leftarrow \text{getPrice}(u)$; // market assigns personalized price to consumer

// MATCHING PROCESS

3 D an empty map from $V \rightarrow V$;
4 $G \leftarrow (V, E)$ such that directed edge $(u, v) \in E$ exists iff $p_u > p_v$;

5 **for** $u \in V$ **do**
6 $D(u) \leftarrow \text{solveMatching}(u, G, k)$; // produces matching of consumers

// EXCHANGE PROCESS

7 **for** $u \in V$ **do**
8 $v = D(u)$;
9 $m \leftarrow \text{computeM}(j_{uv})$; // gets m in a (de)centralized manner
10 **if** $f_u(j_{uv}) > 0$ and $f_v(j_{uv}) > 0$ **then**
11 consumer u pays m dollars to v to get good g at price p_v ;
12 **else**
13 consumer u pays p_u to the market \mathcal{M} for good g

14 **return** G

decisions, any consumer x who has not yet bought good g will do so at price p_x , the original offered price. Algorithm 1 below references sub-routines *getPrice*, *solveMatching* and *computeM*, which we detail in this section, along with utility function f_u for consumer u .

5.2 Consumer attributes

Consumer properties: We assign each consumer v a vector of consumer properties r_v . We implement r_v as a scalar, but this can be trivially extended as a vector. We consider this consumer property to represent some demographic feature that is used by the pricing algorithm \mathcal{A} to assign price p_v .

Resource constraint: Each consumer $v \in V$ has the same resource constraint k —the number of interactions for which consumer v can serve as intermediary. All consumers serve as buyer for one interaction.

Utility function: Each consumer u has a utility function f_u . The form of f is the same for each consumer. Recall that in a given interaction j_{uv} , consumer u is the buyer while consumer v is the intermediary. The utility function for u can be written as $f_u(j_{uv}) = p_u - m - \epsilon_{u_j}$ and for v as $f_v(j_{uv}) = m(1 - \gamma) - p_v - \epsilon_{v_j}$, where the $(1 - \gamma)$ term accounts for the system receiving γ proportion of the transaction amount m .

Terms ϵ_{u_j} and ϵ_{v_j} are the disutility to u and v from spending time on this interaction. Each consumer u is assigned a truncated Normal distribution \mathcal{E}_u , and ϵ_{u_j} is drawn *i.i.d.* from \mathcal{E}_u .

5.3 Pricing algorithm \mathcal{A} (*getPrice*)

Pricing algorithm \mathcal{A} determines the price for good g offered to each consumer in \mathcal{M} (*getPrice*). To capture a wide range of pricing algorithm behavior, we define a notion of *dispersion* (δ) which represents the spread of prices outputted by algorithm \mathcal{A} when prices are normalized to $[0, 100]$. Mathematically, it is the range within which a large proportion of possible prices fall. We define a pricing algorithm \mathcal{A}_δ which assigns prices with dispersion δ , and has the following properties. 1) The range of feasible prices is in $(0, 100]$. 2) \mathcal{A}_δ is biased based on some immutable attribute r_v of the consumer v . We partition the set of consumers V into non-overlapping groups based on attribute r_v . We call this set of groups \mathcal{G} . 3) Consumer v 's price p_v is drawn from a distribution $\mathcal{D}_{\mathcal{G}_v}$, where \mathcal{G}_v is v 's group. All consumers from the same group have their price drawn from the same distribution. 4) Dispersion δ represents the 2.25σ range of possible prices (i.e., $\max_{g,h \in \mathcal{G}} ((\mu_g - 2.25\sigma_g) - (\mu_h + 2.25\sigma_h))$), where μ_g represents the average net cost to consumers u in group g .

In addition to our simulated prices we also examine an empirical pricing algorithm; we use a pricing model presented in prior work [28], which investigates airline ticket pricing. We assume that prices shared by consumers are collected by an automated system (e.g., browser extension) to ensure that prices available to the algorithm are the same prices available to the consumer.

5.4 Exchange system \mathcal{S} (*solveMatching* and *computeM*)

In **RQ 1** we ask whether the transaction price m should be set centrally or determined by individual negotiations. We experiment with two methodologies for setting the transaction price m :

Centralized: We solve for the optimally fair consumer matching and m -value setting (maximizing fairness criteria \mathcal{F}). The system is not aware of private utility functions of individuals, so recommended matchings at optimal m values may not actually transact. This optimization outputs a set of interactions \mathcal{J} such that each consumer has at most one outgoing interaction (i.e., as buyer), and at most k -many incoming interactions (i.e., as intermediary). The transaction prices m_{uv} for all $j_{uv} \in \mathcal{J}$ are simultaneously set centrally to maximize \mathcal{F} . Each consumer can either accept the transaction price m or refuse.

$$\begin{aligned}
 & \text{minimize} \quad \frac{\sum_{u \in V} \omega_u}{|V|} \\
 & \text{s.t.} \quad p_v / (1 - \gamma) \leq m_{uv} \leq p_u \quad \forall j_{uv} \in \mathcal{J} \\
 & \quad x_{uv} \in \{0, 1\} \quad \forall j_{uv} \in \mathcal{J} \quad // u \text{ buys from } v \\
 & \quad \sum_v x_{uv} \leq 1 \quad \forall u \in V \quad // a \text{ consumer can buy once} \\
 & \quad \sum_u x_{uv} \leq k \quad \forall v \in V \quad // \text{and act as intermediary } k \text{ times} \\
 & \quad \omega_u = \sum_{v \in V} \underbrace{x_{uv} m_{uv}}_{\text{buyer's new price if traded}} + \sum_{v \in V} \underbrace{(1 - x_{uv}) p_u}_{\text{buyer's new price if not traded}} \\
 & \quad - \sum_{v \in V} \underbrace{x_{vu} (m_{uv} (1 - \gamma) - p_v)}_{\text{intermediary's profit if traded}} \quad \forall u \in V
 \end{aligned}$$

Above is the linear program using μ_I , the fairness objective that represents individual mean net cost paid, as our example objective function. The other objectives can be found in Table 1. The goal of this program is to output interactions \mathcal{J} along with transaction prices m_{uv} for all $j_{uv} \in \mathcal{J}$. If no j_{uv} exists in \mathcal{J} for a consumer u , then u will pay the original price it was assigned for good g , which is p_u . Due to the presence of the binary variable x_{uv} , this is a mixed integer program, which is at least NP-hard [3]. Minimizing μ_I and μ_G is a linear objective, while minimizing σ_I or σ_G results in a Mixed Integer Quadratic Program, also NP-hard [42]. In our analysis for the linear objective, we solve to completion. For the quadratic objective, we return the best result given by the solver after a pre-specified amount of time (60 seconds). We use Gurobi [20] as our solver.

Decentralized (individual negotiation): Here, we keep the matching \mathcal{J} given by the centralized process, but allow consumers u and v to negotiate individually for the transaction prices m_{uv} . Rather than simulate bargaining between the consumers, we make an assumption regarding the settled transaction price. In our implementation, we set each m_{uv} to the Nash bargaining solution, which maximizes the product of the net prices paid [41]. We choose this value because it is a likely outcome after individual negotiations. Other values could be used as well, such as the mean of prices.

We test four fairness definitions, \mathcal{F} , and two price setting processes for m ; in total this gives eight methodologies. For convenience we denote each methodology as a fairness metric, X indexed by a process, Y , i.e., X^Y for $X \in \{\mu_I, \mu_G, \sigma_I, \sigma_G\}$ and $Y \in \{C, D\}$. For example, μ_I^C refers to optimizing μ_I via the centralized methodology. We also notate μ^Y to refer to μ_I^Y, μ_G^Y methods and σ^Y to refer to σ_I^Y, σ_G^Y methods.

5.5 Theoretical claims

THEOREM 5.1. For $\epsilon_u = 0 \quad \forall u \in V$, the mean net cost over all consumers after trading is bounded below by $p_{\min}(\frac{1 - \frac{|B|Y}{1-Y}}{1-Y})$, where B is the set of consumers that will never engage as a buyer in the system, and p_{\min} is the minimum price assigned to any consumer.

PROOF. Without loss of generality, let consumer u buy a ticket from consumer v for a cost of m_{uv} . Let us define gain to a consumer from trade as the cost improvement the consumer experiences as a result of a given transaction. For a fixed γ (the proportion of the trade cost taken by the system), the gain to the intermediary from this interaction is $m_{uv}(1 - \gamma) - p_v$ while the gain to the buyer is $p_u - m_{uv}$. Then, the sum of the gain to consumers u and v from this transaction is $(p_u - m_{uv}) + (m_{uv}(1 - \gamma) - p_v) = p_u - p_v - \gamma m_{uv}$.

In order to minimize mean net cost, we must maximize gains from trade across all consumers for all possible transactions. To maximize gain from a single trade between consumers u and v , m_{uv} must be minimized. The minimum value for m_{uv} is the transaction price $m_{uv} = \frac{p_v}{1-\gamma}$. When $\epsilon_u = 0$, at this price the intermediary v experiences 0 gain; a buyer will accept this price if $p_u \geq \frac{p_v}{1-\gamma}$. Denote A as the set of consumer pairs (u, v) that have traded and B as the set of consumers $u \in V$ that have not acted as a buyer. We note that $|A| + |B| = N$ as a consumer can only be a buyer once. We

can then write the net price paid as shown below, since p_{min} is the lowest price possible.

$$\begin{aligned} \sum_{(u,v) \in A} \frac{p_v}{1-\gamma} + \sum_{u \in B} p_u &\geq |A| \frac{p_{min}}{1-\gamma} + |B| p_{min} \\ &= \frac{(|A| + |B|) p_{min} - |B| \gamma p_{min}}{1-\gamma} = p_{min} \left(\frac{N - |B| \gamma}{1-\gamma} \right) \end{aligned}$$

Dividing by N we are left with the following lower bound:

$$p_{min} \left(\frac{1 - \frac{|B| \gamma}{N}}{1-\gamma} \right)$$

□

CLAIM 1. When $\gamma = 0$ and $k = |V| - 1$, the minimum standard deviation in net cost is 0.

PROOF. When $\gamma = 0$, minimizing standard deviation in net price requires that all consumers trade with the consumer who got the lowest price at $m = p_{min}$. Then, all buyers will pay p_{min} , while the intermediary consumer earns nothing, and still pays p_{min} for their good. The net cost to all consumers is p_{min} and thus the standard deviation is 0. □

CLAIM 2. This system satisfies Individual-Rationality (IR).

PROOF. Buyers transact at price m if $f_u(j_{uv}) = p_u - m - \epsilon_{u_i} \geq 0$. If both consumers have positive utility they will trade and benefit. Otherwise, one or both will reject and both will pay their original price. Hence, they are no worse off from participating. □

6 Results

We implement our model \mathcal{S} with 100 consumers ($N = 100$). We sample consumer disutility from $\mathcal{E}_u \sim \text{truncated } N(\mu_u, 1)$ where $\mu_u \sim U(0, 2)$. We set $|\mathcal{G}| = 5$ where every consumer has an equal probability of being assigned to each group. We choose dispersion values $\delta = \{0.05, 0.25, 0.50, 0.75, 0.95\}$ which maps to five pricing algorithms: $\mathcal{A}_{0.05}, \mathcal{A}_{0.25}, \mathcal{A}_{0.50}, \mathcal{A}_{0.75}, \mathcal{A}_{0.95}$; the exact mapping can be found in Appendix C. In our construction, we change the distribution of prices while ensuring that mean price across pricing algorithms is roughly equal (\$50) which allows us to analyze the impact of dispersion directly. All pricing algorithms are constructed by a set of Normal distributions so prices are within $(0, 100]$. For example, $\mathcal{A}_{0.95}$ involves five Normal distributions with means $\{10, 30, 50, 70, 90\}$ respectively, all with standard deviation $\frac{10}{3}$.

In our analysis for the linear objective, we solve to completion. For the quadratic objective, we return the best result given by the solver after a pre-specified amount of time (60 seconds). We use Gurobi [20] as our solver. In regards to optimization feasibility—we find the duality gaps for μ_I and μ_G are 0. For σ_I , the duality gap is 67%, while for σ_G it is 0.02%. This suggests that μ_I and μ_G are more feasible to optimize, while σ_I is more difficult. We discuss this feasibility further in Section 6.2.

6.1 RQ 1: Centralized versus decentralized m -setting

In **RQ 1** we ask: Given a fairness target \mathcal{F} (Table 1), which of the following methods to set transaction price m will maximize fairness: individual negotiation or central allocation? In Section 5 we described in detail two methods to set transaction price m for an interaction j_{uv} , where u buys from intermediary v . In brief, the centralized setting solves for optimal edges and prices m , from which consumers can choose to accept or reject the given price. The decentralized option uses these same solved edges, but instead allows paired consumers to negotiate for the transaction price m . In our implementation, m values for these interactions are chosen according to the Nash bargaining solution.

To answer **RQ 1**, we simulate transactions on our system \mathcal{S} . We test the centralized and decentralized procedures for setting m , four fairness targets \mathcal{F} (i.e., eight cases), six values for k (the number of interactions for which a consumer can serve as intermediary), with $\gamma = 0.4$ (the cut the system takes from the interaction). Here we present $\mathcal{A}_{0.95}$ as the pricing algorithm (prices are highly dispersed with $\delta = 0.95$). Each simulation is run 100 times; we present the average of these, and bands that show one standard deviation. In Figure 2, we plot all eight cases against four fairness targets.

Notably, the centralized procedures vary minimally and have little impact on fairness targets (Figure 2). Under this procedure, each consumer is required to either accept or reject the price that maximizes fairness globally. As described in Theorem 1, to minimize mean net cost, a consumer acting as intermediary will be asked to engage in the exchange without benefiting monetarily. Because each consumer experiences disutility (ϵ) that is unknown to the system, very few intermediaries engage in exchanges at centrally-set prices. The same is true for other fairness targets.

We contrast this with decentralized procedures where pairings are set, but consumers negotiate the prices. These negotiated prices are more likely to provide utility to both consumers, and hence more trades occur. When optimizing for μ_I (individual mean; shown in Figure 2a) and μ_G (group mean; shown in Figure 2c) this results in the decentralized procedure lowering the net cost to an average of \$17 (Figure 2a)—nearly a 66% reduction in cost, compared to the average assigned price of \$50. However, even when edges are selected to minimize the standard deviation at the individual (σ_I ; Figure 2b) or group (σ_G ; Figure 2d) level, the decentralized procedure can lead to increased price variation as k increases, because consumers selfishly negotiate. As k increases, consumers with lower prices can complete more transactions (benefiting more monetarily), while those with high prices can only benefit by buying once.

6.2 RQ 2: Feasibility of our fairness measures

In **RQ 2** we ask whether some fairness conditions (i.e., definitions of \mathcal{F} , described in Section 4) are more feasible to achieve than others. In particular, how does varying \mathcal{F} in scope (i.e., individual vs. group) and objective (mean vs. variance) affect the ability of the system to achieve that fairness outcome? To answer this question we run simulations as described previously.

In Figure 2 we show how optimizing for each fairness target \mathcal{F} affects the measure of each outcome. Here we will focus on the decentralized scenario, as centralized price-setting results in

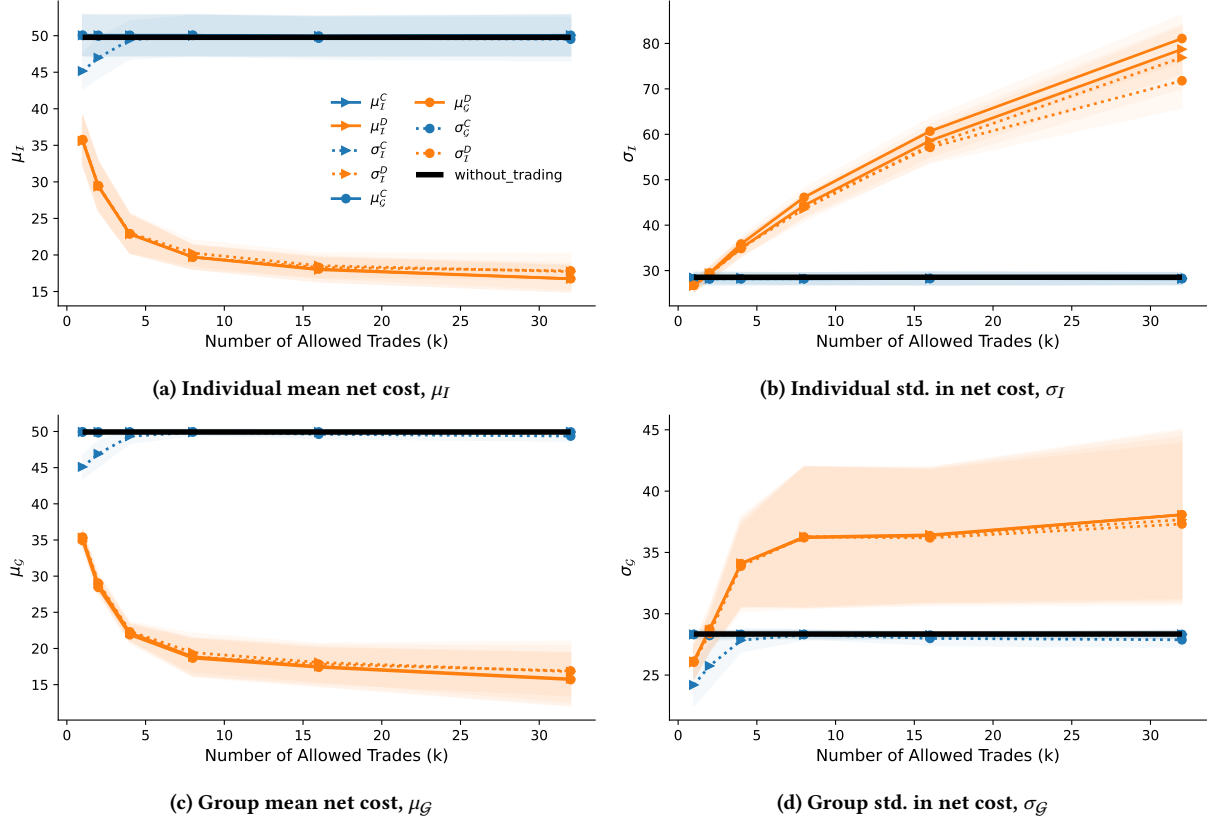


Figure 2: Realization of four fairness definitions (a-d) under eight different optimizations (eight curves representing four fairness targets, two price-setting procedures). We set $N = 100$, $\gamma = 0.4$ and use pricing algorithm $\mathcal{A}_{0.95}$. Centralized methods are blue; decentralized orange. Variance bands show one standard deviation. μ_I and μ_G can be effectively lowered by decentralized methods targeting their respective metrics. σ_I and σ_G are less feasible for all metrics; in particular σ_I for decentralized methods increases with k .

Table 2: Our “feasibility” measure for each fairness metric when $k = 32$, $\gamma = 0.4$. We examine whether optimizing for each metric actually improves said metric. μ^D methodologies are feasible while σ^D methodologies are not.

Metric	Pre-exchange	Post-exchange	% Change
μ_I	49.8	16.7	-66%
σ_I	28.5	76.9	170%
μ_G	49.9	15.7	-69%
σ_G	28.3	37.3	32%

minimal trades. We see that when specifically optimizing for μ_I and μ_G , decentralized price-setting is able to effectively reduce the average price paid. On the other hand, achieving low standard deviation (Figure 2b, 2d) is more difficult. We can investigate how optimizing for each fairness metric compares to the “ideal” scenario. In Table 2 we show numerically that μ -optimizing procedures are more feasible under decentralized settings, while σ -optimizing are less so. During decentralized negotiations, as k (the number of times a consumer can act as intermediary) increases, intermediaries are

able to benefit multiple times while buyers only benefit once. This increases the gap, even when choosing σ -optimizing pairings.

6.3 RQ 3: Revenue to the exchange system \mathcal{S}

In RQ 3 we ask: How does revenue to exchange system \mathcal{S} change due to γ (the cut taken by the system), N (the number of consumers), and δ (the dispersion of pricing algorithm \mathcal{A})? For exchange system \mathcal{S} to be economically viable, 1) the system must be able to recover operation costs and 2) consumers must have an incentive to participate even when the system takes a cut. We compare decentralized methods to determine which earns the system the most revenue. In Figure 3 we show, for different γ values, similar system revenues from each approach. As γ increases, revenue increases until some maximum (approximately $\gamma = 0.8$), where consumers refuse to trade, and revenue decreases. At this point, the system earns $\approx \$2200$ (see the peak in Figure 3). Given that the original seller would have made $\$5000$ when consumers were not trading (with an average good price of $\$50$ for 100 consumers), this represents a substantial cut of revenue being redirected to the exchange. In Figure 3 we show system revenue *and* consumer intermediary profits specifically for μ_I^D . For sufficiently high γ , system revenue





Figure 3: System revenue for four methodologies and intermediaries' profits for μ_T^D . System revenue changes with respect to γ under high dispersion $\delta = 0.95$ and $k = 16$. The system at most earns $\approx \$2200$ in revenue, representing a large cut to the original seller's revenue without trading. For sufficiently high γ , system revenue falls as consumers find fees too high. Intermediary profits decline as a function of γ .

decreases as the system takes more and consumers find it unprofitable to trade. Intermediary consumers consistently lose profits as γ increases.


To study the impact of dispersion δ and number of consumers N on revenue, we again focus on μ_T^D . We examine the importance of dispersion (δ) in Figure 4, where we run our simulation with five different pricing algorithms: $\mathcal{A}_{0.05}, \mathcal{A}_{0.25}, \mathcal{A}_{0.50}, \mathcal{A}_{0.75}, \mathcal{A}_{0.95}$. We construct our dispersion models to hold means constant—this allows for direct comparison of dispersion levels, which can dramatically impact the prices that individuals pay as well as the sustainability of the system. In Figure 4a we show that under system \mathcal{S} , higher dispersion models result in lower average prices for individuals. For high values of k , the highest dispersion model achieves an average net price of \$16 versus \$50 in the lowest dispersion case—a near 66% reduction. This is in contrast to the medium dispersion case ($\delta = 0.5$) with a reduction of $\approx 15\%$ and the low dispersion case ($\delta = 0.05$) with no reduction.

Figure 4b shows that this system earns more under higher dispersion settings; if dispersion is too low, trading with any fees is not viable. We closely examine the high dispersion scenarios: the pricing algorithm with $\delta = 0.75$ narrowly earns more than $\delta = 0.95$ (within variance bounds). When dispersion is sufficiently high, consumers are willing to trade. If dispersion increases (from $\delta = 0.75$ to 0.95), it results in a decrease in the mean Nash bargaining solution for price, so the system earns slightly less. Nevertheless, the revenues are comparable.

We can also use the dispersion δ of the market to determine whether one should invest in this type of system. We see in low dispersion settings ($\delta = 0.05$ or $\delta = 0.25$) that no N value would be able to sustain this system; for $\delta = 0.50$ to $\delta = 0.95$ the revenue scales linearly up to $N = 500$. To justify developing such a system, market prices need to exhibit sufficiently high dispersion—else there is no revenue to be made regardless of N . Whether this system \mathcal{S} is profitable to build depends on the cost structure of the implementation: the fixed cost as well as cost that scales with N .



(a) Effect of k on μ_I under different δ



(b) Effect of N on revenue under different δ

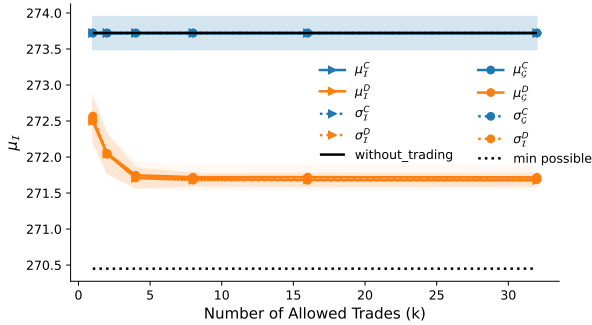
Figure 4: Impact of price dispersion (δ) on outcomes ($\gamma = 0.4$). As k increases, more consumers can access the lower price (a). Note that μ_I without trading is $\approx \$50$. If dispersion is too low for a given γ , trading will decrease (b). Assuming a financially viable γ , system revenue increases with N , and $\delta = 0.75$ maximizes revenue. At $\delta > 0.75$ nearly all consumers have incentive to trade; increasing δ lowers the Nash bargaining solution and thus system revenue. Increasing price dispersion lowers mean cost paid (μ_I) and increases system revenue to a point.

Our results suggest that more dispersed personalized pricing allows the system \mathcal{S} to better help users achieve fairer outcomes while earning good revenue. This provides an opportunity for fair pricing as well as keeping sellers accountable when employing extreme personalization.

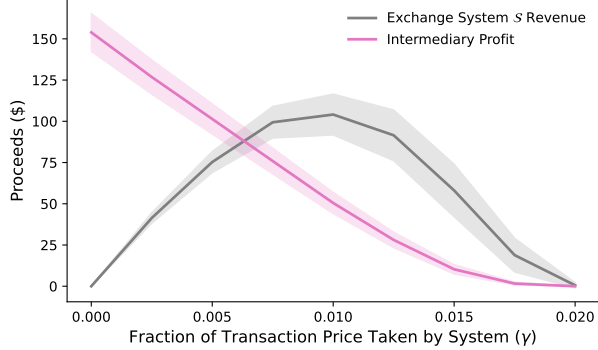
7 Discussion

Analysis of an empirical pricing distribution: Here we show our system \mathcal{S} on a real (rather than simulated) pricing distribution. We use prices as modeled in prior work [28]. Using this pricing distribution (which includes nine non-overlapping groups), we run a simulation with $N = 100$ consumers. Figure 5 shows improvement even when the range in prices is small. Specifically, the range is \$6.15, and the maximum price is \$275.82; this gives a dispersion of ≈ 0.02 after normalizing prices to $(0, 1]$. Compared to the simulated example, the minimum price achievable is not 0 but rather ≈ 270 ;

hence this is the lowest possible expected price. Without trading, the gap between the average price and best price was \$3.23. Trading using the μ_I^D methodology results in \$1.22, a 62% reduction of the gap (Figure 5a) while reducing the sticker price by $\approx 0.75\%$. We set γ values to be much smaller than in the simulated examples due to the higher transaction prices. Even with low γ values, revenues are reasonable, peaking at $\gamma = 0.01$ (Figure 5b) and about \$100 in revenue for $N = 100$. These results indicate the feasibility of implementing our system in a market such as this one.



(a) Welfare improvement for flights, ($\gamma = 0.005$)



(b) γ vs system proceeds

Figure 5: Exchange system on an empirical pricing distribution. In (a) we see again that decentralized methods are able to reduce the gap between the best possible price and the prices paid by 62%. In (b) the cut taken by the system γ must be low; at $\gamma = 0.01$ the system still earns \$100.

What is fair?: We consider various definitions of fairness in this work, outlined in Table 1. Ideally, one would like to maximize all definitions simultaneously for the most “fair” result. In particular, the perfect solution would involve every member of every group simultaneously achieving the lowest price, minimizing the variation and average price paid across individuals and groups. Our results suggest simultaneously achieving low prices and low variation in prices across individuals or groups is difficult. In particular, there is an explicit trade-off between mean and variance. In order to reduce the net price, an intermediary with a low price must be incentivized to engage in trades by earning money. However, at this scale, an intermediary can benefit from multiple trades, while a buyer can

only benefit from a single trade—this can increase variance. We make no claim regarding which is ideal; it is context dependent.

A more informed system: In our simulations, we assume that the system \mathcal{S} has minimal information: it is aware of the prices that the consumers receive, but not each individual consumer’s disutility (ϵ_v) per trade. This limits the ability of the system to effectively match consumers, especially in the centralized case, where the optimal matching requires one of the consumers to transact with minimal personal benefit. Thus, such transactions are not executed. If the exchange system were aware of each consumer’s disutility, it could incorporate that information in determining the optimal prices and result in more successful, fairness-improving transactions. This, however, requires collecting information from consumers that they may not be entirely aware of. Though we model consumer v ’s disutility with a distribution ϵ_v , consumer v “draws” from this distribution for each transaction. This draw represents the current state of consumer v and how long they are willing to spend on this transaction, which is information they may not be able to accurately convey. While having more information could boost the effectiveness of the centralized matching, the details regarding collecting it are better left for future studies.

Closeness of μ_I and μ_G : In our implementation, μ_I and μ_G values track very closely. This is because an average over individuals weights each individual equally, while an average over groups weights individuals in smaller groups more heavily. Given our assumption that group sizes are roughly equal, μ_I and μ_G are algebraically quite close; if group sizes were identical, μ_I and μ_G would be as well. Using different distribution of group sizes can lead to different outcomes. We show an example in Appendix B where in highly skewed cases, a small group getting the best price can receive a far more disproportionate benefit. Future work can look at different and more realistic distributions of pricing buckets.

Interaction between the exchange system and market: In this work, we assume that the exchange system and market are independent; in particular, that the market’s pricing algorithm does not respond to consumer trading. We believe this is a reasonable assumption; if a very small group of individuals is participating in this system, the consumers’ behavior may go unnoticed. We simulate this a single-shot game, with no opportunity for the seller to respond. We recognize that if the behavior is detected, the pricing system could react, possibly raising prices overall and resulting in all users losing access to cheaper pricing [31].

If we allow for a repeated game, there are several ways a marketplace might respond: 1) increasing technical barriers for the consumer that block trades, 2) increasing personalization, 3) decreasing personalization but increasing prices overall. While we cannot predict which of these might occur, the modeling of our system is robust to all of these changes—increased technical barriers can be modeled as additional disutility to each consumer from participating in each transaction; increased personalization would actually increase the fairness of outcomes; decreased personalization, while decreasing fairness, can be easily modeled with a change to pricing algorithm \mathcal{A} .

Another issue that may arise between the system and market is that owners of the exchange system could forgo fairness and accountability in exchange for higher revenue. If the system and market collude, the market could produce prices that, when passed

through the colluding exchange system S , result in worse outcomes for consumers, but higher revenue for the system and market. Considering our finding that higher dispersion can increase revenue and decrease net cost to consumers, existence of such a strategy is possible. Thus, the assumption that the system is independent from the market and has different goals is key.

8 Future Work and Limitations

Scale of transactions: In this work we discuss the sale of one good g . We assume consumers V simultaneously browse the market M for good g , so the system S matches these consumers. This assumption is important so all consumers in V can acquire good g in a timely manner. It is possible, however, that consumers are browsing the same marketplace for multiple goods. In this case, multiple matchings and transactions would occur in parallel, increasing profit but also the cost of the system S to maintain itself. If our system were to be deployed on a real market, the time scale and number of goods would need to be considered in the matching. We leave the multi-good matching problem for future work.

Market structure: Our simulations assume a particular market structure, where a single good is available with enough supply to satisfy all demand. We also assume that each consumer *must* buy the good. However, one rationale for personalized pricing is that some users are less willing to pay for a good than others [16]. Incorporating this requires a more detailed utility function for good purchase, which we defer to future work. We also assume that the resource constraint per user is the same; in practice this may not be true.

System participation: Here we allowed consumers to participate for free but pay a fee per transaction. Another potential structure is to charge a flat fee to participate, but then allow consumers to exchange freely. However, this structure requires convincing users to join the system with an upfront fee. Investigating and comparing these exchange system structures is an avenue for future work.

Human studies: In this work we make simplifying assumptions regarding consumers' actions, e.g., the form of their utility functions, willingness to participate, demographic attributes, and resource constraints. To test these assumptions, future human studies would be ideal. This would allow us to test our design principles on users who may not respond as we originally modeled them. Other design principles could be tested as well (e.g., wording of messaging to consumers) to incentivize trading.

9 Conclusion

In this paper we introduced a fairness-centered exchange system that takes advantage of personalized pricing to improve fairness. We simulated the effect of price dispersion and explored two different transaction price-setting procedures paired with four fairness targets to answer **RQ 1** and **RQ 2**. These modeling choices set the transaction prices from which the system and consumers can profit. To answer **RQ 3**, we showed that our system's revenue is higher when prices are more dispersed, and that consumers are able to achieve lower prices—up to 66% improvement for the mean net cost to individuals and 69% for the mean net cost to a group. We demonstrated that a decentralized negotiation approach is better

able to achieve most notions of fairness compared to a centralized approach. While a designer could choose to focus on either mean or standard deviation as fairness targets, from a financial sustainability perspective, minimizing mean net cost paid by an individual earned more than other targets. Our approach is a consumer-driven solution to personalized pricing that does not rely on fair pricing by the seller, or marketplace regulations. Further directions include considering multiple goods—requiring computing multiple matchings in parallel, and changing the market structure such that all consumers are not required to buy the good. One might also consider varied financial structures to support the exchange system, or conducting human studies to better test the assumptions we made in this work. Our results provide theoretical evidence that such a system could improve fairness for consumers while sustaining itself financially and holding sellers accountable for extreme personalized pricing.

Acknowledgments

We would like to thank Andy Lee and Tanvi Bajpai for providing feedback on initial drafts, especially on theoretical contributions. We would also like to thank the various anonymous reviewers for their suggestions.

References

- [1] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. 2018. A reductions approach to fair classification. In *International conference on machine learning*. PMLR, 60–69.
- [2] Alekh Agarwal, Miroslav Dudík, and Zhiwei Steven Wu. 2019. Fair regression: Quantitative definitions and reduction-based algorithms. In *International Conference on Machine Learning*. PMLR, 120–129.
- [3] Sanjeev Arora and Boaz Barak. 2009. *Computational complexity: a modern approach*. Cambridge University Press.
- [4] Khursheed Aurangzeb, Sheraz Aslam, Syed Muhammad Mohsin, and Musaed Alhussein. 2021. A fair pricing mechanism in smart grids for low energy consumption users. *IEEE Access* 9 (2021), 22035–22044.
- [5] Juan Pedro Aznar, Josep Maria Sayeras, Guillem Segarra, and Jorge Claveria. 2018. Airbnb landlords and price strategy: Have they learnt price discrimination from the hotel industry? Evidence from Barcelona. *International Journal of Tourism Sciences* 18, 1 (2018), 16–28.
- [6] Moshe Babaioff, Nicole Immorlica, Yingkai Li, and Brendan Lucier. 2021. Making auctions robust to aftermarkets. *arXiv preprint arXiv:2107.05853* (2021).
- [7] Siddhartha Banerjee, Kamesh Munagala, Yiheng Shen, and Kangning Wang. 2024. Fair price discrimination. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2679–2703.
- [8] Joachim Baumann, Piotr Sapiezynski, Christoph Heitz, and Anikó Hannák. 2024. Fairness in Online Ad Delivery. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1418–1432.
- [9] Suman K. Bera, Deeparnab Chakrabarty, Nicolas J. Flores, and Maryam Negahbani. 2019. Fair Algorithms for Clustering. (2019), 1–21. arXiv:1901.02393 <http://arxiv.org/abs/1901.02393>
- [10] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. 2015. The limits of price discrimination. *American Economic Review* 105, 3 (2015), 921–957.
- [11] Simina Branzei, Mahsa Derakhshan, Negin Golrezaei, and Yanjun Han. 2024. Learning and Collusion in Multi-unit Auctions. *Advances in Neural Information Processing Systems* 36 (2024).
- [12] Maxime C Cohen, Adam N Elmachtoub, and Xiao Lei. 2022. Price discrimination with fairness constraints. *Management Science* 68, 12 (2022), 8536–8552.
- [13] Kate Donahue and Jon Kleinberg. 2021. Model-sharing games: Analyzing federated learning under voluntary participation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 5303–5311.
- [14] Kate Donahue and Jon Kleinberg. 2023. Fairness in Model-Sharing Games. In *Proceedings of the ACM Web Conference 2023 (Austin, TX, USA) (WWW '23)*. Association for Computing Machinery, New York, NY, USA, 3775–3783. doi:10.1145/3543507.3583483
- [15] Wei Du, Depeng Xu, Xintao Wu, and Hanghang Tong. 2021. Fairness-aware agnostic federated learning. In *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*. SIAM, 181–189.
- [16] Jean-Pierre Dubé and Sanjog Misra. 2017. *Personalized Pricing and Consumer Welfare*. Working Paper 23775. National Bureau of Economic Research. doi:10.3386/w23775

- [17] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. *ITCS 2012 - Innovations in Theoretical Computer Science Conference* (2012), 214–226. doi:10.1145/2090236.2090255 arXiv:1104.3913
- [18] Yahya H Ezzeldin, Shen Yan, Chaoyang He, Emilio Ferrara, and A Salman Avestimehr. 2023. Fairfed: Enabling group fairness in federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 7494–7502.
- [19] Vincent Grari, Arthur Charpentier, and Marcin Detyniecki. 2022. A fair pricing model via adversarial learning. *arXiv preprint arXiv:2202.12008* (2022).
- [20] Gurobi Optimization, LLC. 2023. Gurobi Optimizer Reference Manual. <https://www.gurobi.com>
- [21] Aniko Hannak, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson. 2014. Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 conference on internet measurement conference*. 305–318.
- [22] Moritz Hardt, Eric Mazumdar, Celestine Mendler-Dünnér, and Tijana Zrnic. 2023. Algorithmic Collective Action in Machine Learning. arXiv:2302.04262 [cs.LG]
- [23] James E Heyman and Barbara A Mellers. 2008. Perceptions of fair pricing. *Handbook of consumer psychology* (2008).
- [24] Kevin Hogan. 2018. *Consumer Experience in the Retail Renaissance*. Technical Report. Deloitte Digital, SalseForce. <https://www.deloittdigital.com/us/en/blog-list/2018/consumer-experience-in-the-retail-renaissance--how-leading-brand.html>
- [25] Wei Huang, Tianrui Li, Dexian Wang, Shengdong Du, and Junbo Zhang. 2020. Fairness and accuracy in federated learning. *arXiv preprint arXiv:2012.10069* (2020).
- [26] Ravi Jagadeesan and Alexander Teytelboym. 2021. *Matching and prices*. Technical Report. Working paper, Stanford Univ.
- [27] 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.
- [28] Aditya Karan, Naina Balepur, and Hari Sundaram. 2023. Your Browsing History May Cost You: A Framework for Discovering Differential Pricing in Non-Transparent Markets. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 717–735. doi:10.1145/3593013.3594038
- [29] Amir-Hossein Karimi, Bernhard Schölkopf, and Isabel Valera. 2021. Algorithmic Recourse: From Counterfactual Explanations to Interventions. 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, 353–362. doi:10.1145/3442188.3445899
- [30] Jon Kleinberg, Sendhil Mullainathan, Manish Raghavan, and Alexandra Chouldechova. 2017. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data* 6, 2 (2017), 1–23. doi:10.4230/LIPIcs.ITCS.2017.43 arXiv:1703.00056
- [31] Georgia Kosmopoulou, Qihong Liu, and Jie Shuai. 2016. Customer poaching and coupon trading. *Journal of Economics* 118 (2016), 219–238.
- [32] Matt Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual fairness. *Advances in Neural Information Processing Systems* 2017-Decem, Nips (2017), 4067–4077. arXiv:1703.06856
- [33] Jie Li, Tianqing Zhu, Wei Ren, and Kim-Kwang Raymond. 2023. Improve individual fairness in federated learning via adversarial training. *Computers & Security* 132 (2023), 103336. doi:10.1016/j.cose.2023.103336
- [34] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. 2021. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning*. PMLR, 6357–6368.
- [35] Tian Li, Maziar Sanjabi, Ahmad Beirami, and Virginia Smith. 2019. Fair resource allocation in federated learning. *arXiv preprint arXiv:1905.10497* (2019).
- [36] Kasper Lippert-Rasmussen and Lauritz Aastrup Munch. 2023. 467Price Discrimination in the Digital Age. In *Oxford Handbook of Digital Ethics*. Oxford University Press. doi:10.1093/oxfordhb/9780198857815.013.24 arXiv:https://academic.oup.com/book/0/chapter/323167072/chapter-ag-pdf/56886722/book_37078_section_323167072.ag.pdf
- [37] Lingjuan Lyu, Xinyi Xu, Qian Wang, and Han Yu. 2020. Collaborative fairness in federated learning. *Federated Learning: Privacy and Incentive* (2020), 189–204.
- [38] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 54)*, Aarti Singh and Jerry Zhu (Eds.). PMLR, 1273–1282. <https://proceedings.mlr.press/v54/mcmahan17a.html>
- [39] Akiva A Miller. 2014. What do we worry about when we worry about price discrimination-the law and ethics of using personal information for pricing. *J. Tech. L. & Pol'y* 19 (2014), 41.
- [40] Jeffrey Moriarty. 2021. Why online personalized pricing is unfair. *Ethics and Information Technology* 23, 3 (2021), 495–503.
- [41] Martin J Osborne. 1990. Bargaining and markets. (1990).
- [42] Alberto Del Pia, Santanu S Dey, and Marco Molinaro. 2017. Mixed-integer quadratic programming is in NP. *Mathematical Programming* 162 (2017), 225–240.
- [43] Timothy J Richards, Jura Liaukonyte, and Nadia A Streletskaia. 2016. Personalized pricing and price fairness. *International Journal of Industrial Organization* 44 (2016), 138–153.
- [44] Julio J Rotemberg. 2011. Fair pricing. *Journal of the European Economic Association* 9, 5 (2011), 952–981.
- [45] Alvin E Roth and Marilda Sotomayor. 1992. Two-sided matching. *Handbook of game theory with economic applications* 1 (1992), 485–541.
- [46] Sajjad Salehi, Fattaneh Taghiyareh, Mohammad Saffar, and Kambiz Badie. 2012. A context-aware architecture for mental model sharing in intelligent agents. In *2012 IEEE 10th International Symposium on Applied Machine Intelligence and Informatics (SAMI)*. 313–318. doi:10.1109/SAMI.2012.6208979
- [47] Nripsuta Ani Saxena, Wenbin Zhang, and Cyrus Shahabi. 2024. Unveiling and mitigating bias in ride-hailing pricing for equitable policy making. *AI and Ethics* (2024), 1–12.
- [48] Lloyd S Shapley and Martin Shubik. 1971. The assignment game I: The core. *International Journal of game theory* 1, 1 (1971), 111–130.
- [49] Benjamin Shiller. 2013. First Degree Price Discrimination Using Big Data. (2013).
- [50] Edward Small, Kacper Sokol, Daniel Manning, Flora D Salim, and Jeffrey Chan. 2024. Equalised odds is not equal individual odds: Post-processing for group and individual fairness. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1559–1578.
- [51] Seamus Somerstep, Ya'acov Ritov, and Yuekai Sun. 2024. Algorithmic Fairness in Performative Policy Learning: Escaping the Impossibility of Group Fairness. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 616–630.
- [52] Qian Wang, Qihang Sun, Kui Ren, and Xiaohua Jia. 2016. THEMIS: Collusion-resistant and fair pricing spectrum auction under dynamic supply. *IEEE Transactions on Mobile Computing* 16, 7 (2016), 2051–2064.
- [53] Jianyu Xu, Dan Qiao, and Yu-Xiang Wang. 2023. Doubly Fair Dynamic Pricing. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 9941–9975.
- [54] Renzhe Xu, Xingxuan Zhang, Peng Cui, Bo Li, Zheyang Shen, and Jiazheng Xu. 2022. Regulatory instruments for fair personalized pricing. In *Proceedings of the ACM Web Conference 2022*. 4–15.
- [55] Han Yu, Zelei Liu, Yang Liu, Tianjian Chen, Mingshu Cong, Xi Weng, Dusit Niyato, and Qiang Yang. 2020. A fairness-aware incentive scheme for federated learning. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 393–399.
- [56] Xubo Yue, Maher Nouihered, and Raed Al Kontar. 2023. GIFAIR-FL: A Framework for Group and Individual Fairness in Federated Learning. *INFORMS Journal on Data Science* 2, 1 (April 2023), 10–23. doi:10.1287/ijds.2022.0022
- [57] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. 2017. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. *26th International World Wide Web Conference, WWW 2017* (2017), 1171–1180. doi:10.1145/3038912.3052660 arXiv:1610.08452
- [58] Richard Zemel, Yu Wu, Kevin Swersky, Toniann Pitassi, and Cynthia Dwork. 2013. Learning fair representations. *30th International Conference on Machine Learning, ICML 2013* 28, PART 2 (2013), 1362–1370.
- [59] Yuchen Zeng, Hongxu Chen, and Kangwook Lee. 2021. Improving fairness via federated learning. *arXiv preprint arXiv:2110.15545* (2021).
- [60] Aurora Zhang and Anette Hosoi. 2024. Structural Interventions and the Dynamics of Inequality. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1014–1030.
- [61] Frederik Zuiderveen Borgesius and Joost Poort. 2017. Online price discrimination and EU data privacy law. *Journal of consumer policy* 40 (2017), 347–366.

A Notation

The following table serves as a reference for notation used throughout the paper.

Table 3: Summary of Notation

NOTATION	DESCRIPTION
$G = (V, E)$	the set of agents V connected by edges E
\mathcal{M}	the marketplace
g	the one type of good being offered in \mathcal{M}
\mathcal{S}	the exchange system
\mathcal{A}	the pricing algorithm
\mathcal{P}	the agent-matching process
\mathcal{J}	the set of agent interactions defined by the matching process
\mathcal{F}	the fairness metric that our process aims to optimize
r_v	the vector of consumer v 's properties
δ	the dispersion of pricing algorithm \mathcal{A}
p_v	the price offered to agent v by algorithm \mathcal{A}
N	the number of agents V
$f_v(j_{uv})$	utility function f for agent v that takes an interaction and price as input
m	the transaction price of an interaction
k	resource constraint for all agents
γ	the proportion of transaction price m that the system \mathcal{S} takes

B Impact of different group size distributions

As mentioned in the discussion, in the main body of the paper we study the case when group sizes are roughly equal. As an additional measure, we briefly examined a case where the group sizes are unequal ($N = 100$, $\delta = 0.95$, $\gamma = 0.4$). Specially, we consider groups sizes sampled from a power law distribution with $\beta = 2$. In one instance, the group size is assigned directly proportional to the price (*i.e.*, the largest group has the highest price) while the other instance is reversed (*i.e.*, the largest group has the lowest price).

In Figure 6a, we see that in μ_I , while the percentage reduction varies slightly with a lower number of trades k , as k increases, μ_I is reduced by a similar amount for the skewed and equally sized groups.

In contrast, in Figure 6b we see that μ_G , the mean net cost paid by groups, is reduced much further in the scenario where the small group is offered the lowest price. As k gets larger, we see that the percentage reduction is above 100%, because this small group earns enough outsized benefit to drive the μ_G negative. While all groups can benefit in this scenario, the small, privilege group has the opportunity to earn much more and drive down the average group price more substantially.

C Pricing algorithm details

As described in Algorithm 1, we construct a family of pricing distributions \mathcal{A}_δ parameterized by δ , where δ represents the 2.25σ range of possible prices. We consider $\delta = \{0.05, 0.25, 0.5, 0.75, 0.95\}$ with $|\mathcal{G}| = 5$. For a given pricing algorithm, we have $\mu_1, \mu_2, \dots, \mu_5$ for

each group and a fixed σ for all groups. For each group member in group $g \in \mathcal{G}$, we sample a price from $N(\mu_g, \sigma)$. Here we detail for all five pricing algorithms \mathcal{A}_δ , the respective group's mean as well as the σ used. By construction, the price distributions produced by these algorithms have the same mean (\$50), which allows us to specifically focus on understanding the role of dispersion within our proposed system.

Table 4: Parameters that go into \mathcal{A}_δ for the respective pricing algorithm

δ	μ_1	μ_2	μ_3	μ_4	μ_5	σ
0.95	10	30	50	70	90	$\frac{30}{9}$
0.75	20	35	50	65	80	$\frac{30}{9}$
0.5	30	40	50	60	70	$\frac{20}{9}$
0.25	40	45	50	55	60	$\frac{10}{9}$
0.05	50	50	50	50	50	$\frac{10}{9}$

D Flight Pricing Simulation

D.1 Flight Price Model

We use a pricing model from prior work to produce our empirical pricing distribution [28]. In this market, the average price was \$270.45. The work presents pricing models from various sellers in the market, we use the model presented for "Third Party 1". In this work, the authors present the offsets from the base price of the flight ticket : [\$4.55, \$1.46, \$5.29, \$3.55, \$6.15, \$2.91, \$2.36, \$5.05].

From these mean values we created nine non-overlapping groups of consumers, where the price for each group is the offset from the base price. We directly assign the price from one of the nine mean values listed purely based on the group.

D.2 Consumer Utilities

The utility structure of the consumers is the same as described in the main body of the paper. For the individual disutility, each consumer u is assigned a truncated Normal distribution \mathcal{E}_u with varying means (drawn from $U[0, 1]$) but the same standard deviation (0.5) where the numbers are chosen with respect to the prices from the flight price model.

Table 5: Prices for $\mathcal{A}_{\text{flight}}$. No variance is used for this algorithm.

Group	Price
μ_1	\$270.45
μ_2	\$271.91
μ_3	\$272.46
μ_4	\$273.01
μ_5	\$274.21
μ_6	\$275.42
μ_7	\$275.82
μ_8	\$276.20
μ_9	\$276.60




Figure 6: Impact of group size distribution on the reduction in k (number of allowed trades) increases. We optimize for μ_I and μ_G in the decentralized case, and set $\gamma = 0.4$, $\delta = 0.95$. Three group size distributions are shown: one where all groups are equally sized, one where the largest group is offered the lowest price, and where the largest group is offered the highest price. Here the percentage reduction in μ_I , μ_G is shown compared to a no-trading situation.