Designing Fair Systems for Consumers to Exploit Personalized Pricing

Aditya Karan[*], **Naina Balepur**[*], **Hari Sundaram**

University of Illinois Urbana-Champaign
Urbana, IL 61820 USA
karan2@illinois.edu, nainab2@illinois.edu, hs1@illinois.edu

Abstract

Many online marketplaces personalize prices based on consumer attributes. Since these prices are private, consumers will not realize if they spend more on a good than the lowest possible price, and cannot easily take action to get better prices. In this paper we introduce a system that takes advantage of personalized pricing so consumers can profit while improving fairness. Our system matches consumers for 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, they are able to achieve the most fair outcomes. Conversely, when transaction prices are centrally set, consumers are often unwilling to transact. Minimizing the average price paid by an individual or group is most profitable for the system, while achieving a 67% reduction in prices. We see that a high dispersion (or range) of original prices is necessary for our system to be viable. Higher dispersion can actually lead to increased consumer welfare, and act as a check against extreme personalization. Our results provide theoretical evidence that such a system could improve fairness for consumers while sustaining itself financially.

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 (Karan, Balepur, and Sundaram 2023; Hannak et al. 2014). If you knew about this price discrepancy, you could ask your friend to buy the shoes for you, paying them back, maybe throwing in a few extra dollars so you both benefit. Many online marketplaces employ elements of personalized pricing — surveys indicate that many companies intend to use AI or automated systems to personalize prices (Hogan 2018). While many consumers understand that they may see personalized prices, it is difficult to take money-saving actions without knowing what other consumers are paying. As differential pricing can vary with protected attributes, facilitating this type of exchange can help reduce disparate pricing outcomes across protected classes. This exchange, however, must be designed to simultaneously improve fairness and be financially viable — this can greatly depend on price dispersion in the market. In this work we address challenges regarding design of this platform.

Prior work on decreasing disparity across groups suggests fairer pricing algorithms (Xu et al. 2022), or advises users to change their behavior to get better pricing (Kusner et al. 2017). Other work has looked at the benefit of trading data, primarily for the purpose of improving downstream models (Fernandez, Subramaniam, and Franklin 2020). Business literature (Kosmopoulou, Liu, and Shuai 2016; Gans and King 2007) has investigated the effect of coupon trading on a firm's ability to personalize. However, using trading as a way to reduce the impact of personalized pricing and improve fairness has not explicitly been studied.

In this work, we develop a system that allows consumers to reduce the price they pay for a good online by trading. Figure 1 outlines how our system, the market, and consumers interact. First, agents are assigned personalized prices by the market (a). Some agents agree to participate in this trading system, which allows the system to see only these prices (b). The system assigns a matching, and the transaction price (*i.e.,* the amount to be exchanged) is agreed upon by the agents *or* set by the system (c). Finally, if the transaction is mutually beneficial; it occurs. The initially higher-paying consumer pays the lower-paying consumer the transaction price, and receives the good (d).

To implement this system, we make decisions regarding the matching procedure and transaction price setting procedure for the resulting matches. We produce one optimal matching, and then test two transaction price-setting procedures: centralized and decentralized. As this system is intended to improve fairness, we consider four fairness targets for each transaction price-setting procedure. We produce the four fairness targets by examining two objectives (i.e., the value measured: mean and standard deviation in price), and scopes (i.e., who it is measured for: individuals and groups). We must also ensure that this system is financially sustainable. To do so, we explore how the system's revenue is affected by the number of agents in the collective, the cut the system takes from each transaction, and the dispersion (or range) of prices initially offered by the marketplace. We conclude by examining a case study on an empirical pricing distribution modeled in prior work.

We find that one method for setting the transaction price (decentralized negotiations between agents) and one fairness

[*]These authors contributed equally.

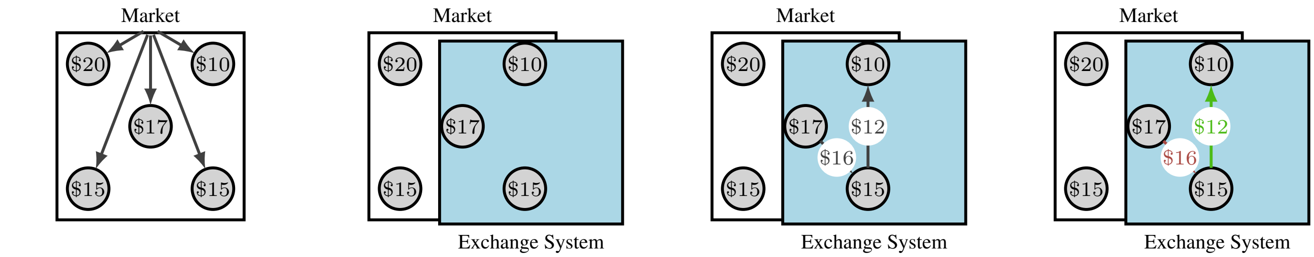

(a) Market assigns prices to each agent for a good

(b) Some agents decide to join collective

(c) System gives matching; prices are set

(d) Mutually beneficial transactions occur

Figure 1: Depiction of the relationship between the market (white square), the system (blue square), and consumers (gray circles). Consumers enter the market and are offered prices (a). Then, some of these consumers decide to join the collective, thereby choosing to participate in the exchange system (b). The system assigns a matching, and prices for these transactions are assigned or decided upon (c). Of these matched transactions, only those that are mutually beneficial to both agents occur (d).

target (minimizing the average net price paid by each agent) implemented together is fairness maximizing; we are able to decrease the average price paid by individuals and groups by 67%, when compared to a non-trading scenario (*i.e.,* original prices assigned by the market). We find that we can achieve fairness metrics simultaneously if we hold the objective (*i.e.,* the value we measure) constant while varying the scope (*i.e.,* who it is measured for). The same is not true while varying the objective and holding the scope constant. When considering the exchange system's revenue, the decentralized approach earns more. Our system acts as an effective counterweight to personalized pricing; we find that as dispersion of prices decreases, agents are no longer able to achieve the 67% welfare increase from higher dispersion scenarios. In low dispersion scenarios (*i.e.,* when limited personal pricing occurs in the market), the exchange system is not viable, failing to generate revenue even as the number of participants increases. Our contributions are as follows:

Exchange system design: We design an exchange system that could feasibly be implemented and financially sustain itself. Prior work in this area has found personalized pricing in different markets (Hannak et al. 2014; Aznar et al. 2018; Karan, Balepur, and Sundaram 2023), and some has even proposed fair pricing algorithms (Aurangzeb et al. 2021; Grari, Charpentier, and Detyniecki 2022; Xu, Qiao, and Wang 2023; Xu et al. 2022; Wang et al. 2016). However, none to our knowledge has proposed agent-driven solutions to unfair personalized pricing. In this work we close that gap — exploring design choices and fairness targets that will lead to a financially self-sustaining system that also maximizes welfare for consumers.

Opportunities from personalized pricing: We find that while personalized pricing can cause welfare loss to consumers, if taken advantage of, it provides a money-saving opportunity. While prior work notes that personalized pricing can help consumers by helping them afford certain goods (Dubé and Misra 2017), we design a system that helps all consumers. Surprisingly, we show that highly dispersed personalized pricing leads to higher welfare *and* higher revenue to the system, when compared to low dispersion. This finding emphasizes the strengths of our exchange system design: we do not require sellers to cooperate. In fact, when they don't, it can be beneficial to consumers and acts as a counterweight to extreme, unfair personalization.

2 Problem Statement

Consider a simple marketplace $\mathcal{M}$ that offers one type of good g to consumers (or agents) V. Each consumer $v \in V$ requires exactly one unit of good g; the supply of g is finite but sufficient to satisfy the population. Each agent $v \in V$ has a vector of attributes r_v. The agents 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. The 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 agents via a matching process $\mathcal{P}$. This matching process outputs a set of pairwise agent interactions $\mathcal{J}$. Let $j_{u \to v} \in \mathcal{J}$ be the (directed) interaction parameterized by agents u and v, where $p_u > p_v$. In an interaction $j_{u \to v}$, agent u will pay some m dollars to agent 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 agents by the matching process $\mathcal{P}$, but all interactions need not occur. An interaction is executed if both agents u and v benefit according to their utility functions f_u and f_v, which can include both positive and negative terms.

We frame the matching process $\mathcal{P}$ as a network problem. Our agents exist in a directed network $G = (V, E)$ where $(u, v) \in E$ if $p_u > p_v$. The matching process $\mathcal{P}$ produces interactions $\mathcal{J}$, the matched directed edges. Each agent v has at most k matched edges: its resource constraint. If an interaction $j_{u \to v}$ is executed, the edge (u, v) has been transacted on. We call u the buyer and v intermediary; one agent might be matched as both intermediary and 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 welfare: 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 size N of the collective, cut γ taken by the system, and dispersion δ of the pricing algorithm $\mathcal{A}$?

3 Related Work

Fair pricing: Unfair pricing caused by online behaviors has been found by several studies (Hannak et al. 2014; Aznar et al. 2018; Karan, Balepur, and Sundaram 2023). One solution for consumers is to redesign pricing algorithms for fairness. Seminal work on fair pricing (Heyman and Mellers 2008; Rotemberg 2011) examines what fair pricing is and how consumers may react to various pricing algorithms by firms. Kallus and Zhou (2020) examine how different markets and concerns (*e.g.,* information asymmetry) inform which fairness criteria to consider. Grari, Charpentier, and Detyniecki (2022) study fair pricing under adverse selection in the context of insurance, balancing actuarial risk with a demographic parity or equalized odds constraint. Xu et al. (2022) explore how imposing restrictions on the degree of personalized pricing (*i.e.,* the price of an object can't vary more than $x\%$) can be customized to balance the needs of both buyers and sellers. The goal of these works on fair pricing is to create fair algorithms that a seller can use while still achieving high profits.

While this approach *could* achieve higher welfare 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 direct cooperation from the entity employing differential pricing. Some work has taken this approach — counterfactual fairness (Kusner et al. 2017) could provide a way for 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 (Karimi, Schölkopf, and Valera 2021). On the collective action side, Hardt et al. (2023) quantified the effect of a collective coordinating feature changes against an algorithmic system. In our work, collaboration and exchange happens not to change the direct outcome of an algorithmic system, but as a secondary layer to adjust final outcomes for individuals. Our work does not require designing a causal model or having individuals change their behaviors to get better prices. Rather, we focus on agents who are incentivized by monetary reward to share information.

Data marketplaces: Data marketplaces have grown in popularity as entities can directly monetize their own data by allowing purchasers easier access. Our exchange system can be framed as a data marketplace as well — agents share information with the system, and the system and agents themselves receive a monetary reward in return. Liang et al. (2018) survey pricing challenges and the structure of such data markets. Fernandez, Subramaniam, and Franklin (2020) examine existing data markets and their challenges in adoption. In particular, they discuss the importance of deciding the market type (internally versus externally-facing) and rule design. One key challenge is that in many data-trading markets it is hard for buyers and sellers to accurately value their data. Even if they could estimate a monetary value, the ultimate usefulness of the data is not known since it must be combined or processed for downstream applications (Fernandez, Subramaniam, and Franklin 2020).

In our scenario, the "data" being traded is pricing information, which may provide agents access to a lower price. Payoffs to our agents are better understood, but still difficult to value, as they never aware of the lowest possible price. Modeling this exchange is not easy, as it relies on inferences of user-specific behavior which may not be explicitly known. In this work we build off existing work on data marketplaces, defining a new sort of market where users receive direct monetary benefit from sharing.

Federated learning: In federated learning, data is indirectly "shared" via model updates and aggregated (*e.g.,* FedAvg (McMahan et al. 2017)). There is a very large existing body of work investigating fairness and trying to achieve it, on both group (Lyu et al. 2020; Ezzeldin et al. 2023; Zeng, Chen, and Lee 2021; Yu et al. 2020; Huang et al. 2020; Li et al. 2021; Du et al. 2021; Li et al. 2019) and individual (Li et al. 2023; Yue, Nouiehed, and Al Kontar 2023) 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 agents who act purely selfishly to achieve some notion of fairness. Donahue and Kleinberg (2021) too consider agent incentives; they 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 (Donahue and Kleinberg 2023) egalitarian and proportional fairness in the context of these model-sharing games. Salehi et al. (2012) develop a model-sharing architecture for agents' mental models but do not explicitly consider fairness.

Marketplace mechanisms: Shapley and Shubik (1971) and Roth and Sotomayor (1992) both examine the assignment game — a two sided matching with money, where they show properties of the core. In our work, notably, we're interested in situations where everyone has access to the same good at different prices, rather than valuing goods at different prices. Jagadeesan and Teytelboym (2021) look at how some markets give rise to universal pricing while others employ personalized personalized pricing. Babaioff et al.

Table 1: Our fairness definitions. All definitions rely on ω_u, the net cost to agent u. We use μ_g to represent the average net cost to agents u in group g.

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

(2021) and Branzei et al. (2024) design auctions to limit gains from post auction dealings, similar to the ones we introduce here.

4 Fairness

In **RQ 2** we ask whether some definitions of $\mathcal{F}$ are more feasible than others to achieve. Here we give those definitions as well as what we mean by feasibility. To develop our fairness definitions at the individual and group levels, we started by considering what the "ideal" scenario would be for all agents. We determined that the ideal fair outcome would result in all agents paying the *same, lowest* price for the good. In other words, we sought to minimize both average price and the standard deviation in prices paid by agents. 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 Section 5. We measure the "feasibility" of our fairness metrics by considering how close each mean and standard deviation come to the ideals.

Thus, we vary the objective and scope to get different definitions of fairness, *i.e., what* we measure, and *who* we measure it for. Our measured outcome is always the net cost incurred by agents from participating in system $\mathcal{S}$, assuming that each agent purchases exactly one unit of good g. For example, if agent u is offered price p_u by the market, but pays m dollars to agent v to buy good g for price p_v, then agent u's net cost is m. Agent v's profit from selling the ticket is $(m(1-\gamma) - p_v)$ and therefore their net cost for their own ticket is $p_v - (m(1-\gamma) - p_v)$. We denote ω_u as the net cost to agent u, where $\omega_u < 0$ implies monetary gain to u. We experiment with two objectives (mean and standard deviation) and two scopes (individual and group-level). We consider a group to be a collection of individuals who share a demographic attribute; *i.e.,* u and v are in the same group if $r_u = r_v$. We call this set of groups $\mathcal{G}$. This gives us 4 different fairness measures; we present them in Table 1.

We seek to minimize these values. Minimizing individual mean implies we want each agent to minimize the price they're paying (some may even earn money from $\mathcal{S}$). Minimizing individual standard deviation implies we want agents to benefit similarly from system $\mathcal{S}$; no one agent should make a large profit off of others, nor should any agent not profit while others do. 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 comprehensive. We also considered more standard fairness metrics, such as demographic parity, predicted parity, and equalized odds (Agarwal et al. 2018; Agarwal, Dudík, and Wu 2019; Bera et al. 2019; Zemel et al. 2013; Zafar et al. 2017; Kusner et al. 2017; Kleinberg et al. 2017; Dwork et al. 2012). However, we decided that for our application we could best capture fairness by designing metrics from our ideal scenario. We vary our fairness target over these definitions to test **RQ 2**.

5 The Model

In this section we detail our model, which consists of strategic agents V interacting in a market $\mathcal{M}$ alongside exchange system $\mathcal{S}$. We begin with a high-level overview of the model, followed by detailed descriptions of each relevant feature. We introduced considerable notation in Section 2; it is summarized in Table 3 in Appendix A.

Algorithm 1: Model overview

```
Data: V the set of agents, N the size of the collective
// PRICING ALGORITHM
for u ∈ V do
  | p_u ← getPrice(u)
// MATCHING PROCESS
D ← {} ;
G ← (V, E) such that directed edge (u, v) ∈ E
  exists iff p_u > p_v ;
for u ∈ V do
  | D(u) ← v from getMatching(G, k) ;
// EXCHANGE PROCESS
for u ∈ V do
  | v = D(u) ;
  | m ← getM(j_{u→v}) ;
  | if f_u(j_{u→v}) > 0 and f_v(j_{u→v}) > 0 then
  |   | agent u pays m dollars to v to get good g at
  |   |   price p_v ;
  | else
  |   | agent u pays p_u to the market M for good g
return G
```

Model overview

We begin with agents V who desire exactly one unit of good g on market $\mathcal{M}$. Each agent $v \in V$ is initialized with attribute vector r_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 agents. The transaction price m of each interaction is then set in either a centralized or decentralized fashion, as described in **RQ 1**. 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 agents make transaction decisions, any agent x who has not yet bought good g will do so at price p_x. Algorithm 1 above references sub-routines $getPrice$, $getMatching$ and $getM$, which we detail in this section, along with utility function f_u for agent u.

Agent attributes

Our agents are associated with attributes and utility functions — in this section we elaborate further.

Consumer properties: We assign each agent v a vector of consumer properties r_v. We implement r_v as a scalar, but this can be trivially extended. 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 k: Each agent $v \in V$ has the same resource constraint k — the number of interactions for which agent v can serve as intermediary. Recall that all agents can serve as buyer for only one interaction.

Utility function: Each agent u has a utility function f_u. The form of this function is the same for each agent. Recall that in a given interaction $j_{u\to v}$, agent u is the buyer while agent v is the intermediary. Then, agent u will gain utility from any savings from buying the good from v at price m rather than from the market at price p_u. Agent v will get utility from any profit after buying good g for agent u at price p_v. We give the utility functions below:

$$f_u(j_{u\to v}) = p_u - m - \epsilon_{u_i}$$
$$f_v(j_{u\to v}) = m(1-\gamma) - p_v - \epsilon_{v_i}$$

The $(1-\gamma)$ term accounts for the system receiving γ proportion of the transaction amount m. ϵ_{u_i} and ϵ_{v_i} represent the disutility to u and v from spending time on this interaction. Each agent u is assigned a truncated Normal distribution $\mathcal{E}_u$, and ϵ_{u_i} is drawn from $\mathcal{E}_u$.

Pricing algorithm, $\mathcal{A}$

Here we describe how prices are assigned to agents in the market $\mathcal{M}$ ($getPrice$). Pricing algorithm $\mathcal{A}$ determines the price for the good g offered to each agent. To capture a wide range of pricing algorithm behavior, we define a notion of *dispersion* (δ) that captures the spread of prices outputted by algorithm $\mathcal{A}$. Mathematically, it is the range within which a large proportion of possible prices fall. We define a pricing algorithm $\mathcal{A}_\delta$ which assigns prices with dispersion δ. We construct each pricing algorithm $\mathcal{A}_\delta$ as follows:

- The range of feasible prices is in $(0, 1]$.
- $\mathcal{A}_\delta$ is biased based on some immutable attribute r_v of the consumer v. We partition the set of agents V into non-overlapping groups based on attribute r_v. We call this set of groups $\mathcal{G}$.
- Agent v's price p_v is drawn from a distribution $\mathcal{D}_{\mathcal{G}_v}$, where $\mathcal{G}_v$ is v's group. All agents from the same group have their price drawn from the same distribution.
- Dispersion δ represents the 2σ range of possible prices (*i.e.,* $\max_{g,h\in\mathcal{G}}((\mu_g - 2\sigma_g) - (\mu_h - 2\sigma_h))$).

Our construction of these pricing algorithms is intended to capture a range of different seller behaviors. We detail specific implementation in the results section. We also examine an empirical pricing algorithm; we use the pricing model presented in prior work (Karan, Balepur, and Sundaram 2023), which investigates differential prices for airline tickets. Details for this model can be found in Appendix C.

Exchange system, $\mathcal{S}$

The system $\mathcal{S}$ facilitates agent exchange of money for a good. Developing our system requires us to consider the design of two factors: the fairness target $\mathcal{F}$ and the choice of m, the money the buyer u pays to intermediary v. These modeling choices align with **RQ 2** and **RQ 1**; we have already described the design of the fairness targets. In this section we discuss the choice of m in detail.

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

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

$$\begin{aligned}
\text{minimize } & \frac{\sum_{u\in V}\omega_u}{|V|} \\
\text{s.t. } & \frac{p_v}{1-\gamma} \le m_{uv} \le p_u \quad \forall j_{u\to v}\in\mathcal{J} \\
& x_{uv}\in\{0,1\} \quad \forall j_{u\to v}\in\mathcal{J} \\
& \sum_v x_{uv} \le 1 \quad \forall u\in V \\
& \sum_u x_{uv} \le k \quad \forall v\in V \\
& \omega_u = \sum_{v\in V} x_{uv}m_{uv} + \sum_{v\in V}(1-x_{uv})p_u \\
& \quad - \sum_{v\in V} x_{vu}(m_{uv}(1-\gamma) - p_v) \quad \forall u\in V
\end{aligned} \tag{1}$$

Above is the linear program using μ_I 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_j for all $j_{u\to v}\in\mathcal{J}$. If no $j_{u\to v}$ exists in $\mathcal{J}$ for an agent u, then u will pay the original price it was assigned for good g, which is p_u.

Due to the presence of x_{uv} this is a mixed integer program, which is at least NP-hard (Arora and Barak 2009). When we optimize μ_I and $\mu_{\mathcal{G}}$ the objective is linear, while when we minimize σ_I or $\sigma_{\mathcal{G}}$, this is a Mixed Integer Quadratic Program, also NP-hard (Pia, Dey, and Molinaro 2017). 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 (Gurobi Optimization, LLC 2023) as our solver.

Decentralized (individual negotiation): Here, we keep the interactions (or matching) $\mathcal{J}$ given by the centralized process, but allow agents to negotiate individually for the transaction prices m_j. Rather than simulating bargaining between the agents, we make an assumption regarding the settled transaction price. In our implementation, we set each m_j equal to the Nash bargaining solution, which maximizes the product of the welfare gain (Osborne 1990). We choose this value because it is a likely outcome after individual negotiations, but other values could be used as well, such as the mean of prices.

We test four definitions for $\mathcal{F}$ and two m-setting processes; in total this gives eight methodologies. For convenience we denote a specific process and fairness optimization tuple as X^Y for $X \in \{\mu_I, \mu_\mathcal{G}, \sigma_I, \sigma_\mathcal{G}\}$ and $Y \in \{\text{C}, \text{D}\}$. For example, μ_I^C refers to optimizing μ_I via the centralized methodology. We notate μ^Y to refer to $\mu_I^Y, \mu_\mathcal{G}^Y$ methods and σ^Y to refer to $\sigma_I^Y, \sigma_\mathcal{G}^Y$ methods.

Theoretical claims

Given the model definition above, we make the following theoretical claims. First, we consider the bounds on mean agent welfare (Theorem 1) and standard deviation in agent welfare (Claim 1). Knowing these values allows us to determine the feasibility of our constructed fairness criteria. We also show that agents are always better off having participated in the system (Claim 2). We state them here and defer the proofs to Appendix B.

Theorem 1. *The mean net utility over all agents after trading is bounded below by* $p_{min}(1 + \frac{\gamma}{|V|(1-\gamma)})$.

Claim 1. *When* $\lambda = 0$ *and* $k = |V| - 1$*, the optimal standard deviation in price is* 0.

Claim 2. *This system satisfies Individual-Rationality (IR) regardless of how transaction price* m *is set.*

6 Results

We implement our model $\mathcal{S}$ with a collective size of $N = 100$. We sample $\mathcal{E}_u \sim$ truncated $N(\mu, 0.01)$ where $\mu \sim U(0, 0.02)$. We set $|\mathcal{G}| = 5$ and dispersion values $\delta = \{0.05, 0.25, 0.50, 0.75, 0.95\}$, which gives us five pricing algorithms $\{\mathcal{A}_{0.05}, \mathcal{A}_{0.25}, \mathcal{A}_{0.50}, \mathcal{A}_{0.75}, \mathcal{A}_{0.95}\}$. All pricing algorithms are constructed so prices range from $(0, 1]$. Each pricing algorithm has a set of Normal distributions $\mathcal{D}_\mathcal{G}$. For example, $\mathcal{A}_{0.95}$ involves five Normal distributions with means $\{0.1, 0.3, 0.5, 0.7, 0.9\}$ respectively and all having standard deviation $\frac{1}{30}$. We note that our construction changes the distribution of prices while ensuring that mean price across pricing algorithms is roughly the same (\$0.50) — this allows us to analyze the impact of dispersion directly.

RQ 1: Setting transaction price m

In **RQ 1** we ask: Given fairness metric $\mathcal{F}$ (Table 1), which of the following methods to set transaction price m will maximize welfare: individual negotiation or central allocation?

In Section 5 we described in detail two methods to set transaction price m for an interaction $j_{u \to v}$. In brief, the centralized setting solves for optimal edges and prices m, from which the agents can choose to accept or reject the given price. The decentralized option uses those solved edges from

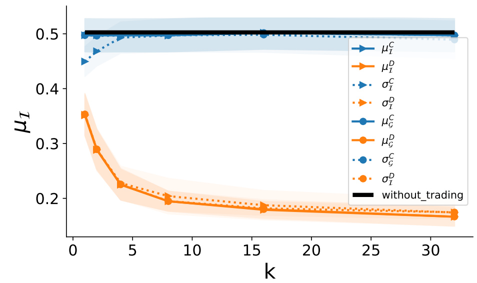

(a) Individual mean welfare, μ_I

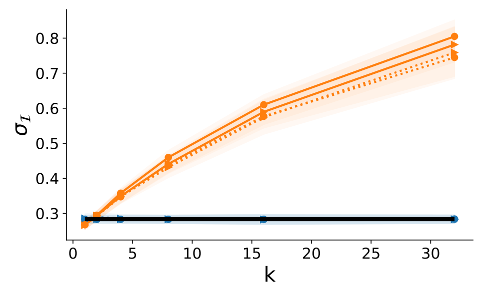

(b) Individual std. in welfare, σ_I

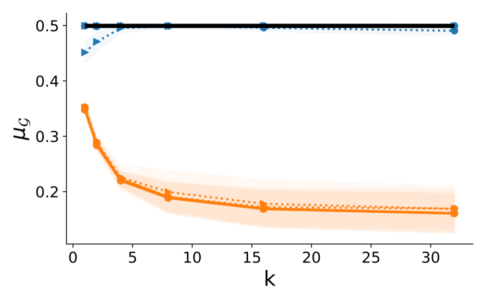

(c) Group mean welfare, $\mu_\mathcal{G}$

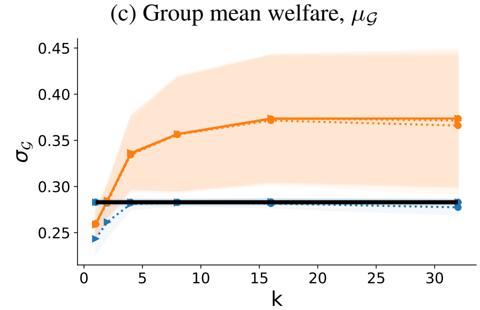

(d) Group std. in welfare, $\sigma_\mathcal{G}$

Figure 2: Realization of eight fairness definitions under different optimization procedures (four fairness targets, two m-setting procedures), $N = 100, \gamma = 0.4$ using pricing algorithm $\mathcal{A}_{0.95}$. We show centralized methods in blue, and decentralized in orange. We show variance bands for one standard deviation, and on the x-axis, we vary k. We see that μ_I and $\mu_\mathcal{G}$ can be effectively lowered by the decentralized methods targeting their respective metrics. σ_I and $\sigma_\mathcal{G}$ are much less feasible for all metrics; in particular σ_I for decentralized methods increases as k increases. This is because agents are still trading at the Nash bargaining solution. As k increases, those with the best price are able to profit more, increasing the variance.

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, under both centralized and decentralized m-setting. We note that μ^D methodologies are feasible while μ^C are not. This is in contrast to σ^C and σ^D methodologies, which see no change in performance. This is because prices set in σ^C scenarios are not transacted on; when agents decide for themselves, they can increase the spread of final prices.

		Centralized		Decentralized	
Metric	Pre Trade	Post Trade	% Change	Post Trade	% Change
μ_I	0.503	0.497	-1%	0.166	-67%
σ_I	0.284	0.283	0%	0.759	168%
μ_G	0.499	0.499	0%	0.160	-68%
σ_G	0.283	0.277	-2%	0.366	29%

the centralized case and then instead allows the paired agents to negotiate for the transaction price m. In our implementation, m values for these interactions are chosen according to the Nash bargaining solution — a likely equilibrium price.

To answer **RQ 1**, we run simulations on our system $\mathcal{S}$. We test a centralized and decentralized procedure for setting m, four fairness targets $\mathcal{F}$, six values for k (the agents' resource constraint), with $\gamma = 0.4$. Here we present $\mathcal{A}_{0.95}$ as the pricing algorithm (*i.e.,* prices are highly dispersed with $\delta = 0.95$). Each simulation is run 100 times; we present the average of these, as well as bands that show one standard deviation. In Figure 2, we plot all eight procedures against four fairness targets. We use blue to represent centralized m-setting while orange represents decentralized. When optimizing for μ_I (Figure 2a) and μ_G (Figure 2c), decentralized procedures that specifically optimize for said fairness metric outperform the centralized variants, achieving an average net cost of $0.17 (nearly a 66% reduction in price paid compared to the average price of $0.50). For σ_I (Figure 2b) we see that the decentralized procedures perform worse as as k increases. This occurs because while the edges are set to minimize the objectives, agents selfishly negotiate. As k increases, agents with better prices can complete more transactions, while those with worse prices can only get a better price once. For σ_G (Figure 2d) we see that decentralized and centralized methods which optimize for σ perform similarly. Notably, in all four settings, the centralized method varies little with k, as very few agents want to transact at the centrally set prices. This is due to the system $\mathcal{S}$ setting prices that are uninformed by agents' utility functions, so any profit to agents may not be sufficient for positive utility.

Allowing agents to negotiate after determining the matched edges is key for transactions to occur. As a reminder, when centrally deciding the m value, the system $\mathcal{S}$ is unaware of individual disutilities which can cause agents to reject an interaction at price m that was set for them. Conversely, a rejection when m is set to be the Nash bargaining solution is much less likely.

RQ 2: Fairness definitions

In **RQ 2** we ask: In the context of our exchange system $\mathcal{S}$, are some fairness conditions (*i.e.,* definitions of $\mathcal{F}$, described in Section 4) more feasible to achieve than others? In particular, how does varying $\mathcal{F}$ in scope and objective affect the ability of the system to achieve the ideal 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. Notably, we see that while when specifically optimizing for μ_I and μ_G, decentralized m-setting is able to effectively reduce the average price paid. On the other hand, achieving low standard deviation (Figures 2b and 2d) is more difficult under both centralized and decentralized m-setting. 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 procedures either do no better or significantly worse than the initial starting point.

Here, we also examine trade-offs — can we achieve multiple definitions of fairness simultaneously? Figure 3a depicts how each methodology achieves a combination of individual mean welfare μ_I and group mean welfare μ_G. Every methodology either achieves both in tandem or achieves neither (the same holds for σ_I and σ_G). However, achieving low mean and standard deviation in welfare simultaneously is difficult (Figure 3b). One can try to achieve low mean, but will sacrifice low standard deviation, and vice versa. Our results suggest that in practice, the designer will need to choose which objective to optimize for.

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

In **RQ 3** we ask: How does revenue to exchange system $\mathcal{S}$ change with the number of consumers N, the cut taken by the system γ, and dispersion δ of pricing algorithm $\mathcal{A}$?

For exchange system $\mathcal{S}$ to earn reasonable revenue, two things must happen. First, there must be a mechanism to recover some cost from transactions; and second, agents must have an incentive to participate even with the system takes a cut. We compute the revenue generated by each method: two m-setting methodologies for four fairness objectives (Figure 4) to determine which earns the system the most revenue. In Figure 4a we show, for different γ values, how much the system is able to earn from each approach. We see that the μ^D methodologies result in highest revenue for the system. In contrast, σ^D methodologies earn some revenue, while centralized methods fail to earn. For decentralized negotiation methods, as gamma increases, revenue increases until some maximum (approximately $\gamma = 0.8$), where agents refuse to trade, and revenue decreases. At this point we see the system earning $\approx$ $22. Given that the original seller would have made $50 when agents were not trading, this represents a substantial amount of revenue going to the ex-

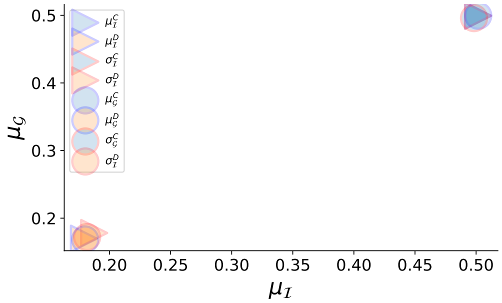

(a) Individual vs group mean welfare, $\mu_{\mathcal{I}}$ vs $\mu_{\mathcal{G}}$

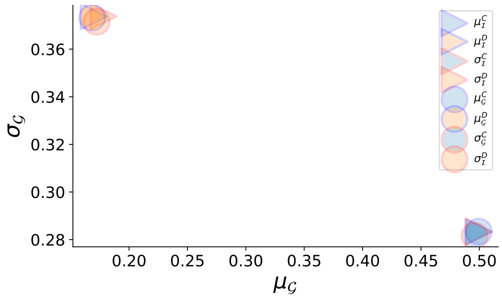

(b) Group mean welfare vs s.d., $\mu_{\mathcal{G}}$ vs $\sigma_{\mathcal{G}}$

Figure 3: Fairness trade-offs in different settings. Ideally, average welfare loss and standard deviation are both low. Achieving low $\mu_{\mathcal{I}}$ and $\mu_{\mathcal{G}}$ simultaneously is possible ($\sigma_{\mathcal{I}}$ and $\sigma_{\mathcal{G}}$ behavior similarly). Conversely, achieving low mean and standard deviation in conjunction are not possible, in either the individual or group case.

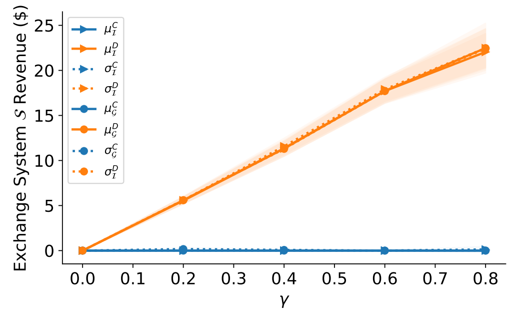

(a) System revenue for eight methodologies

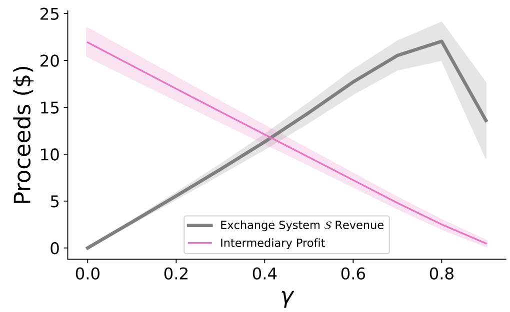

(b) Agent and system proceeds under μ_I^D

Figure 4: In (a) we show how system revenue changes with respect to γ under high dispersion $\delta = 0.95$ and $k = 16$. μ^D methods at least match σ methods, which in turn outperform centralized m-setting methods. In (b) we show for μ_I^D, agent earnings from trades and system revenue. For sufficiently high γ, system revenue falls as agents find fees too high.

change. In Figure 4b we show system revenue *and* agent intermediary profits specifically for μ_I^D. For sufficiently high γ values, system revenue decreases as the system takes too much and agents find it unprofitable to trade. Intermediary agents on the other hand consistently lose profits as γ increases — their profit maximization occurs at $\gamma = 0$.

To study the impact of dispersion and N on revenue, we again focus on μ_I^D, which is fairness and system revenue-maximizing. We examine the importance of dispersion in Figure 5, 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 5a 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 \$0.16 versus \$0.50 in the lowest dispersion case — a near 64% reduction. Figure 5b 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). This is because when dispersion is sufficiently high, agents are willing to trade. If dispersion increases (say, from $\delta = 0.75$ to 0.95), it results in a decrease in the Nash bargaining solution, meaning that the system earns slightly less. Nevertheless, the revenues are comparable.

We can also use the collective size N and 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 would be able to sustain this system, for $\delta = 0.5$ to $\delta = 0.95$ the revenue scales linearly up to $N = 500$. In order to justify developing such a system, the market prices need to exhibit sufficiently high dispersion — otherwise there is no revenue to be made regardless of the collective size. 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. 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 even when sellers employ extreme personalization.

7 Discussion

Analysis of an empirical pricing distribution: Here we show our system $\mathcal{S}$ on a real (rather than simulated) price distribution. We use prices as modeled in prior work (Karan, Balepur, and Sundaram 2023); find details regarding implementation in Appendix D. Using this pricing distribu-

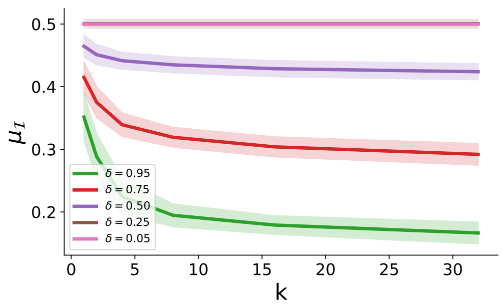

(a) Effect of k on μ_I under different δ

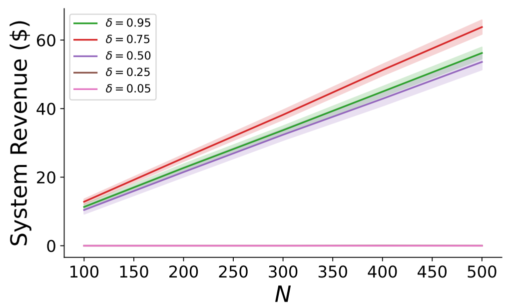

(b) Effect of N on revenue under different δ

Figure 5: Impact of price dispersion on outcomes. Because lower prices are more common in high dispersion scenarios, as k increases, more agents are able to access this lower price (a). In (b) we see that if the dispersion is too low for a given γ, trading will decrease. Assuming a viable γ value, system revenue increases with N, and $\delta = 0.75$ leads to the highest revenue. At $\delta > 0.75$ nearly all agents have incentive to trade; increasing δ lowers the Nash bargaining solution, resulting in slightly lower system revenue.

tion (which includes nine non-overlapping groups), we run a simulation with $N = 100$ agents. Figure 6 shows that even though the range in prices is small, we still see improvement. Specifically, the range is \$6.15, and the maximum price is \$275.82; this gives a dispersion of ≈ 0.02 after normalizing prices to $(0, 1]$. 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 6a). We set γ values to be much smaller than in the simulated examples since the system $\mathcal{S}$ takes γ proportion of the higher transaction prices. Even with low γ values, revenues are reasonable, peaking at $\gamma = 0.01$ (Figure 6b) and about \$100 in revenue for $N = 100$. These results show that it could be feasible to implement our system in a market such as this one.

What is fair?: We consider various definitions of fairness in this work, outlined in Table 1. Ideally, one would like to maximize all of these 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. The results in Figure 3a suggest that simultaneously achieving group and individual mean welfare objectives is feasible. However, Figure 3b suggests that 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. We make no claim on which one is ideal — that is context dependent. However, if finances are of concern, Figure 4a suggests that focusing on μ is more likely to result in a self-sustaining system. Other notions of fairness (*e.g* counterfactual) were not explored here but could be more or less feasible and produce different revenues.

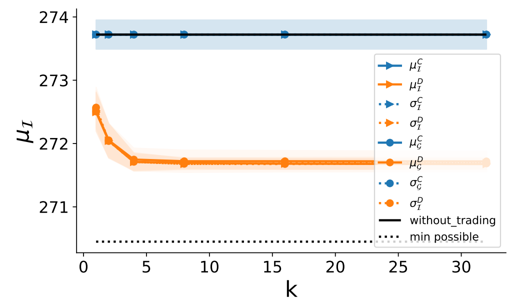

(a) Welfare improvement for flights, ($\gamma = 0.005$)

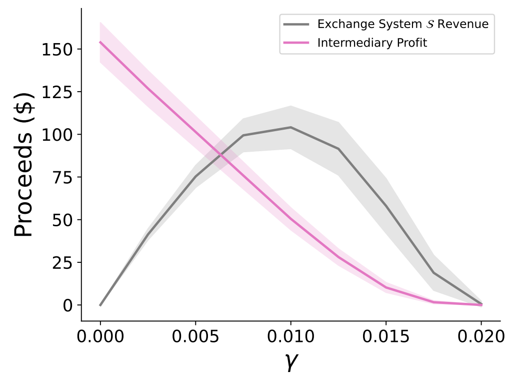

(b) γ vs system proceeds

Figure 6: Exchange system on an empirical pricing distribution. In (a) we see again that μ^D performs best, and is 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.

The closeness of μ_I and μ_G: In our implementation, μ_I and μ_G values track very closely in many results. 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 groups are roughly equal in size, μ_I and μ_G are algebraically quite close; if group sizes were exactly the same, μ_I and μ_G would be identical. Purposefully varying group sizes could dramatically change the two measures; we leave this as a topic for future exploration.

Collusion of system and market: We assumed that the system is fairness-minded and independent from the market. However, when two entities are profit-seeking, it is possible that owners of the system could forgo fairness in exchange for higher revenue. If the system and market collude, the market could produce prices that when passed through the colluding exchange system $\mathcal{S}$, result in worse outcomes for consumers, but higher revenue for the system and market. Considering our finding that higher dispersion can increase revenue and welfare, existence of such a strategy is very 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 type of good g. We assume that agents V all browse the market $\mathcal{M}$ around the same time for good g, so the system $\mathcal{S}$ matches these agents together. We believe this closeness in browsing time is important so all agents in V can acquire good g in a timely manner. It is possible, however, that multiple agents are browsing the same marketplace for multiple goods. This means that multiple matchings and transactions would occur in parallel, increasing the profit but also the cost of the system $\mathcal{S}$ 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 this 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 agent *must* buy the good. However, one rationale for personalized pricing is that some users are less willing to pay for a good than others (Dubé and Misra 2017). Incorporating this requires a more detailed utility function for purchasing the good in the first place, 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.

Response by pricing algorithm: Importantly, we assume that the market's pricing algorithm does not respond to this new 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. However, 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 — this has been suggested by Kosmopoulou, Liu, and Shuai (2016) This dynamic between the pricing algorithm and agents is a rich area for future exploration.

Collective Formation: Here we allowed agents to join the collective for free but pay a fee per transaction. Another potential structure is to charge agents a flat fee as soon as they choose to join the collective, and then allow them to exchange without additional transaction cost. However, this structure comes with the challenge (both practically and in modeling) of convincing users to join a collective 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 the actions of the agents, *e.g.,* the format of the utility functions, the probability of joining the collective, the demographic attributes, and the resource constraints. To truly test these assumptions, future human studies would be ideal. This would allow us to test our design principles on boundedly-rational humans who may not respond as we originally modeled them. Other design principles could be tested in such a setup as well (*e.g.,* wording of messaging to consumers) to incentivize them to make trades.

9 Conclusion

In this work we introduced a system that takes advantage of personalized pricing to improve fairness. We examined the effect of price dispersion and explored two different transaction price-setting procedures paired with four fairness targets. These modeling choices set the transaction prices from which the system and agents can profit. We showed that our system's revenue is higher when prices are more dispersed, and that agents are able to achieve lower prices (up to 67% improvement compared to baseline). We demonstrated that a decentralized negotiation approach is better able achieve most notions of fairness compared to a centralized approach. We also showed that fairness targets with the same objective (*i.e.,* the value measured), but different scopes (*i.e.,* who it is measured for) were able to be achieved simultaneously. However, fairness targets with the same scope but different objectives were difficult to achieve in conjunction. While a designer could choose to focus on any one of these targets, from a financial sustainability perspective, minimizing mean cost paid by an individual or group 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 regulations on the marketplace. Further directions include varied collective formation approaches as well as evaluating this system on humans. Our results are theoretical evidence that such a system could improve fairness for consumers while sustaining itself financially and providing a useful check against extreme personalized pricing.

References

Agarwal, A.; Beygelzimer, A.; Dudík, M.; Langford, J.; and Wallach, H. 2018. A reductions approach to fair classification. In *International conference on machine learning*, 60–69. PMLR.

Agarwal, A.; Dudík, M.; and Wu, Z. S. 2019. Fair regression: Quantitative definitions and reduction-based algorithms. In *International Conference on Machine Learning*, 120–129. PMLR.

Arora, S.; and Barak, B. 2009. *Computational complexity: a modern approach.* Cambridge University Press.

Aurangzeb, K.; Aslam, S.; Mohsin, S. M.; and Alhussein, M. 2021. A fair pricing mechanism in smart grids for low energy consumption users. *IEEE Access*, 9: 22035–22044.

Aznar, J. P.; Sayeras, J. M.; Segarra, G.; and Claveria, J. 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): 16–28.

Babaioff, M.; Immorlica, N.; Li, Y.; and Lucier, B. 2021. Making auctions robust to aftermarkets. *arXiv preprint arXiv:2107.05853*.

Bera, S. K.; Chakrabarty, D.; Flores, N. J.; and Negahbani, M. 2019. Fair Algorithms for Clustering. 1–21.

Branzei, S.; Derakhshan, M.; Golrezaei, N.; and Han, Y. 2024. Learning and Collusion in Multi-unit Auctions. *Advances in Neural Information Processing Systems*, 36.

Donahue, K.; and Kleinberg, J. 2021. Model-sharing games: Analyzing federated learning under voluntary participation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 5303–5311.

Donahue, K.; and Kleinberg, J. 2023. Fairness in Model-Sharing Games. In *Proceedings of the ACM Web Conference 2023*, WWW '23, 3775–3783. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394161.

Du, W.; Xu, D.; Wu, X.; and Tong, H. 2021. Fairness-aware agnostic federated learning. In *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*, 181–189. SIAM.

Dubé, J.-P.; and Misra, S. 2017. Personalized Pricing and Consumer Welfare. Working Paper 23775, National Bureau of Economic Research.

Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. *ITCS 2012 - Innovations in Theoretical Computer Science Conference*, 214–226.

Ezzeldin, Y. H.; Yan, S.; He, C.; Ferrara, E.; and Avestimehr, A. S. 2023. Fairfed: Enabling group fairness in federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 7494–7502.

Fernandez, R. C.; Subramaniam, P.; and Franklin, M. J. 2020. Data Market Platforms: Trading Data Assets to Solve Data Problems. *Proc. VLDB Endow.*, 13(12): 1933–1947.

Gans, J. S.; and King, S. P. 2007. Perfect price discrimination with costless arbitrage. *International Journal of Industrial Organization*, 25(3): 431–440.

Grari, V.; Charpentier, A.; and Detyniecki, M. 2022. A fair pricing model via adversarial learning. *arXiv preprint arXiv:2202.12008*.

Gurobi Optimization, LLC. 2023. Gurobi Optimizer Reference Manual.

Hannak, A.; Soeller, G.; Lazer, D.; Mislove, A.; and Wilson, C. 2014. Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 conference on internet measurement conference*, 305–318.

Hardt, M.; Mazumdar, E.; Mendler-Dünner, C.; and Zrnic, T. 2023. Algorithmic Collective Action in Machine Learning. arXiv:2302.04262.

Heyman, J. E.; and Mellers, B. A. 2008. Perceptions of fair pricing. *Handbook of consumer psychology*.

Hogan, K. 2018. Consumer Experience in the Retail Renaissance. Technical report, Deloitte Digital, SalseForce.

Huang, W.; Li, T.; Wang, D.; Du, S.; and Zhang, J. 2020. Fairness and accuracy in federated learning. *arXiv preprint arXiv:2012.10069*.

Jagadeesan, R.; and Teytelboym, A. 2021. Matching and prices. Technical report, Working paper, Stanford Univ.

Kallus, N.; and Zhou, A. 2020. Fairness, Welfare, and Equity in Personalized Pricing. arXiv:2012.11066.

Karan, A.; Balepur, N.; and Sundaram, H. 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*, FAccT '23, 717–735. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701924.

Karimi, A.-H.; Schölkopf, B.; and Valera, I. 2021. Algorithmic Recourse: From Counterfactual Explanations to Interventions. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, 353–362. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383097.

Kleinberg, J.; Mullainathan, S.; Raghavan, M.; and Chouldechova, A. 2017. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data*, 67(2): 1–23.

Kosmopoulou, G.; Liu, Q.; and Shuai, J. 2016. Customer poaching and coupon trading. *Journal of Economics*, 118: 219–238.

Kusner, M.; Loftus, J.; Russell, C.; and Silva, R. 2017. Counterfactual fairness. *Advances in Neural Information Processing Systems*, 2017-Decem(Nips): 4067–4077.

Li, J.; Zhu, T.; Ren, W.; and Raymond, K.-K. 2023. Improve individual fairness in federated learning via adversarial training. *Computers & Security*, 132: 103336.

Li, T.; Hu, S.; Beirami, A.; and Smith, V. 2021. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning*, 6357–6368. PMLR.

Li, T.; Sanjabi, M.; Beirami, A.; and Smith, V. 2019. Fair resource allocation in federated learning. *arXiv preprint arXiv:1905.10497*.

Liang, F.; Yu, W.; An, D.; Yang, Q.; Fu, X.; and Zhao, W. 2018. A Survey on Big Data Market: Pricing, Trading and Protection. *IEEE Access*, 6: 15132–15154.

Lyu, L.; Xu, X.; Wang, Q.; and Yu, H. 2020. Collaborative fairness in federated learning. *Federated Learning: Privacy and Incentive*, 189–204.

McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and Arcas, B. A. y. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Singh, A.; and Zhu, J., eds., *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, 1273–1282. PMLR.

Osborne, M. J. 1990. Bargaining and markets.

Pia, A. D.; Dey, S. S.; and Molinaro, M. 2017. Mixed-integer quadratic programming is in NP. *Mathematical Programming*, 162: 225–240.

Rotemberg, J. J. 2011. Fair pricing. *Journal of the European Economic Association*, 9(5): 952–981.

Roth, A. E.; and Sotomayor, M. 1992. Two-sided matching. *Handbook of game theory with economic applications*, 1: 485–541.

Salehi, S.; Taghiyareh, F.; Saffar, M.; and Badie, K. 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.

Shapley, L. S.; and Shubik, M. 1971. The assignment game I: The core. *International Journal of game theory*, 1(1): 111–130.

Wang, Q.; Sun, Q.; Ren, K.; and Jia, X. 2016. THEMIS: Collusion-resistant and fair pricing spectrum auction under dynamic supply. *IEEE Transactions on Mobile Computing*, 16(7): 2051–2064.

Xu, J.; Qiao, D.; and Wang, Y.-X. 2023. Doubly Fair Dynamic Pricing. In *International Conference on Artificial Intelligence and Statistics*, 9941–9975. PMLR.

Xu, R.; Zhang, X.; Cui, P.; Li, B.; Shen, Z.; and Xu, J. 2022. Regulatory instruments for fair personalized pricing. In *Proceedings of the ACM Web Conference 2022*, 4–15.

Yu, H.; Liu, Z.; Liu, Y.; Chen, T.; Cong, M.; Weng, X.; Niyato, D.; and Yang, Q. 2020. A fairness-aware incentive scheme for federated learning. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 393–399.

Yue, X.; Nouiehed, M.; and Al Kontar, R. 2023. GIFAIR-FL: A Framework for Group and Individual Fairness in Federated Learning. *INFORMS Journal on Data Science*, 2(1): 10–23.

Zafar, M. B.; Valera, I.; Rodriguez, M. G.; and Gummadi, K. P. 2017. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. *26th International World Wide Web Conference, WWW 2017*, 1171–1180.

Zemel, R.; Wu, Y.; Swersky, K.; Pitassi, T.; and Dwork, C. 2013. Learning fair representations. *30th International Conference on Machine Learning, ICML 2013*, 28(PART 2): 1362–1370.

Zeng, Y.; Chen, H.; and Lee, K. 2021. Improving fairness via federated learning. *arXiv preprint arXiv:2110.15545*.

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
$f_v(j_{u \to v})$	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 Proofs

Here we show proofs of our theorems and claims.

Theorem 1: *The mean net utility over all agents after trading is bounded below by $p_{min}(1 + \frac{\gamma}{|V|(1-\gamma)})$.*

Proof. The total gain from this transaction is $p_u - m_{uv} + m_{uv}(1-\gamma) - p_v = p_u - p_v - \gamma m_{uv}$. We can sum across all transaction to get the total gain. The total net price is the sum of original agent prices minus the gains from trade. In other words, $\sum_{i \in V} p_i - \sum_{(u,v) \in \mathcal{J}} p_u - p_v - \gamma m_{uv}$. We note that in $\sum_{i \in V} p_i - \sum_{(u,v) \in \mathcal{J}} p_u$, the second term at most one p_i, as we only allow for buyers to trade once. This leaves us only with nodes that never bought from an intermediary. We can write this as $\sum_{i \in V \setminus \mathcal{J}} p_i + \sum_{(u,v) \in \mathcal{J}} p_v + m_{uv}\gamma$. The first two sums together give us the price of the nodes that didn't transact plus the sum of the prices of the nodes that did. This covers all nodes, and we know that these prices can't exceed the min price. Hence we can lower bound this by $\geq Np_{min} + \sum_{(u,v) \in \mathcal{J}} \gamma m_{uv}$. We also know that for a transaction to occur it must be that $\frac{p_v}{1-\gamma} \leq m_{uv}$. Hence we can further upper bound by $\geq Np_{min} + \sum_{(u,v) \in \mathcal{J}} \gamma \frac{p_v}{1-\gamma} \geq Np_{min} + \gamma|\mathcal{J}| \frac{p_{min}}{1-\gamma}$. Computing the average we have $\frac{Np_{min} + |\mathcal{J}| p_{min} \frac{\gamma}{1-\gamma}}{N} = p_{min} + \frac{|\mathcal{J}|}{N} p_{min} \frac{\gamma}{1-\gamma} \geq p_{min}(1 + \frac{\gamma}{|V|(1-\gamma)})$. □

Claim 1: *When $\lambda = 0$ and $k = |V| - 1$, the optimal standard deviation in price is* 0.

Proof. If $\lambda = 0$ then if all agents trade with the agent with the lowest price at $m = p_{min}$ then all agents will receive p_{min}, while the lowest price agent doesn't earn any positive benefit. The net price to all agents is p_{min} and hence the standard deviation will be 0. □

Claim 2: *This system satisfies Individual-Rationality (IR) regardless how transaction price m is set.*

Proof. Regardless how m is set, buyers only transact on m if $f_u(j_{u \to v}) = p_u - m - \epsilon_{u_i} \geq 0$ where $\epsilon_{u_i} \geq 0$ is assumed. If m is small enough then the agent will choose to accept, if m is too large, the agent will reject and instead pay their original price p_u, hence the agent is no worse off. A similar process holds for the intermediaries. If both agents have positive utility they will trade and both receive a benefit. If one agent would not receive a benefit, that agent will reject the trade and both agents will be just paying their original price. Hence they are no worse off in participating in this system. □

C Pricing algorithm details

As described in Section 5, we construct a family of pricing distributions $\mathcal{A}_\delta$ parameterized by δ, where δ represents the 2σ range of possible prices. We consider $\delta = \{0.05, 0.25, 0.5, 0.75, 0.95\}$ with $|G| = 5$. For a given pricing algorithm, we have $\mu_1, \mu_2, ...\mu_5$ 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 $N(\mu_g, \sigma)$. Appendix C details for all five pricing algorithms 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 ($0.50), which allows us to specifically focus on understanding the role of dispersion within our proposed system.

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

δ	μ_1	μ_2	μ_3	μ_4	μ_5	σ
0.95	0.1	0.3	0.5	0.7	0.9	$\frac{3}{90}$
0.75	0.2	0.35	0.5	0.65	0.8	$\frac{3}{90}$
0.5	0.3	0.4	0.5	0.6	0.7	$\frac{2}{90}$
0.25	0.4	0.45	0.5	0.55	0.6	$\frac{1}{90}$
0.05	0.5	0.5	0.5	0.5	0.5	$\frac{1}{90}$

D Flight Pricing Simulation

Flight Price Model

We use a pricing model from prior work to produce our empirical pricing distribution (Karan, Balepur, and Sundaram 2023). 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 the work, they 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].

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

From these mean values we created nine non-overlapping groups of agents, 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.

Agent Utilities

The utility structure of the agents is the same as described in Section 5. For the individual disutility each agent u is assigned a truncated Normal distribution $\mathcal{E}_u$ with varying means (draw from a $U[0, 1]$) but the same standard deviation (0.5) where the numbers are chosen with respect to the prices from the flight price model.