Homogeneous Algorithms Can Reduce Competition in Personalized Pricing

Nathanael Jo[*]
Massachusetts Institute of Technology
nathanjo@mit.edu

Kathleen Creel[†]
Northeastern University
k.creel@northeastern.edu

Manish Raghavan[†]
Massachusetts Institute of Technology
mragh@mit.edu

Ashia Wilson[†]
Massachusetts Institute of Technology
ashia@mit.edu

Abstract

Firms' algorithm development practices are often *homogeneous*. Whether firms train algorithms on similar data, aim at similar benchmarks, or rely on similar pre-trained models, the result is correlated predictions. We model the impact of correlated algorithms on competition in the context of personalized pricing. Our analysis reveals that *(1)* higher correlation diminishes consumer welfare and *(2)* as consumers become more price sensitive, firms are increasingly incentivized to compromise on the accuracy of their predictions in exchange for coordination. We demonstrate our theoretical results in a stylized empirical study where two firms use personalized pricing algorithms to determine consumers' willingness to pay. Our results underscore the potential anti-competitive effects of algorithmic pricing and highlight the need for refined antitrust approaches in the era of digital markets.

1 Introduction

Competing firms increasingly use algorithms to price their goods and services. For example, dynamic pricing – where prices vary frequently, often in response to changes in supply or demand – is now widespread in e-commerce [54], airfare [29], and ridesharing apps [3]. In this work, we focus on *personalized pricing*, where prices are tailored to different consumers based on their willingness to pay. Companies have been observed to personalize prices in various domains, from travel websites that use browser information to charge US-based customers more [35] to a recent lawsuit against DoorDash alleging that the company is charging Apple users more in delivery fees [46].

The personalized pricing strategy where firms segment consumers and charge different prices to each segment is known as "third-degree price discrimination" [51, 43]. While third-degree price discrimination can be accomplished by simple rules (e.g., "charge higher prices to Mac users"), firms have recently turned to machine learning to segment consumers into smaller and more targeted categories [16]. One such example is offering a fixed discount to a subset of consumers [17], as has been documented in ridesharing competitors Uber and Lyft (see Figure 7).

In theory, using algorithms to better price-discriminate can be positive for consumers because those who are willing to pay less still get access to goods and services. Firms will inevitably make mistakes, giving some discounts to those who would be willing to pay more and overlooking others who would only buy with a discount. If these mistakes are not correlated, consumers can choose between a high

[*]Corresponding Author
[†]Ordered alphabetically

2nd Workshop on Regulatable ML at NeurIPS 2024.

and a low price. However, if pricing mistakes are correlated, firms are insulated from the competitive cost of their mistakes: the price-sensitive consumer to whom one firm failed to offer a discount is also not buying from a competitor [12].

Correlated predictions occur because firms often deploy algorithms with "shared components" such as algorithms trained on similar datasets, aimed at similar benchmarks, or based on similar pre-trained models [7]. In other words, model development practices can be *homogeneous*. At the extreme, firms might even adopt the same algorithm from a third party, creating an "algorithmic monoculture" with perfect predictive correlation across firms [26].

Contributions. Our work investigates how correlated algorithms impact competition in the context of personalized pricing. This is in line with, and in many ways an extension of, several recent cases brought by the Federal Trade Commission (FTC) involving multiple parties using the same pricing algorithm, which allegedly inflated hotel [47] and rent prices [48]. We build a game-theoretic model of competition between two firms who use algorithms to price-discriminate. We find that:

1. Consumers are always worse off when competing firms' algorithms are more correlated (Theorem 4.1).
2. As consumers become more price sensitive, firms increasingly prefer to correlate their algorithms (Theorem 4.2).
3. Firms are sometimes willing to give up non-trivial predictive performance in exchange for correlation with competitors (Theorem 5.3).

In all, algorithmic homogenization allows firms to sustain higher prices, potentially resulting in anti-competitive outcomes. Our work has implications for antitrust law, as governments seek to promote competition and prevent collusion via algorithms. We expand on this connection in Section 7.

2 Background and Related Work

Antitrust Law. The spirit of modern antitrust law is to promote competition. There is generally broad consensus that an open and free market economy – which at its core fosters competition – benefits consumers by lowering prices, spurring innovation, and increasing the quality of goods and services [41, 33, 20, 8]. In the United States, antitrust enforcement relies on three sets of federal laws: the Sherman Act, the Clayton Act, and the FTC Act, each prohibiting different actions that harm competition. In this work, we will focus our attention to the parts of the Sherman Act and the FTC Act that are intended to delineate which forms of collusion amongst competitors are illegal.

In general, agreement between competitors to fix prices or divide territory of operation is deemed *per se* illegal, meaning that no further inquiry is needed as to the action's effect on the market or the parties' intent in reaching such an agreement [44, 50]. Establishing additional violations of anti-trust law involves determining whether or not the alleged practice "unreasonably restricts trade" [49]. Absent proof of intent to form an agreement, firms are considered to engage in **tacit collusion**, which is generally not illegal. For example, firms might exhibit "parallel business behavior" by changing prices in response to market conditions or even consciously mirroring the public prices of their competitors without *intending* to reach supra-competitive prices [45, 28].

In evaluating the legality of parallel business behaviors, courts consider various "plus factors" that might tip the scales from tacit to illegal collusion. Evidence that firms are motivated to collude and take actions against their own economic self-interests are a common form of plus factor[27].

Homogeneity, Monoculture, and Model Multiplicity. Our work builds on recent work in machine learning on "algorithmic monoculture", namely the state of affairs in which "many decision-makers rely on the same algorithm" and in doing so correlate their behavior [26]. Existing literature focuses on how monoculture harms the welfare of those who are subject to correlated algorithmic errors or "homogeneous outcomes" [7, 24, 36]. Our work spotlights the harm to consumers that comes from higher prices in the context of personalized pricing.

Our work is also related to the concept of "model multiplicity", where arbitrarily many algorithms can achieve maximal accuracy but differ in other desiderata such as fairness, robustness, or interpretability. [6] showed that a model class' variance is directly proportional to its multiplicity. We argue that in the midst of competition, firms have a natural incentive to choose models with high correlation (perhaps from a model class with less variance), which results in higher overall prices.

Economic Models of Oligopoly Pricing. We consider competition under a duopoly, which has been extensively studied in the economics literature. The works most related to ours are game-theoretic models of duopolies under Bayesian uncertainty [32, 10, 52, 22, 40, 1]. Much of this literature considers whether firms have incentives to collude by sharing information with one another. Whether a model will suggest that firms are rewarded for sharing information depends on a variety of modeling choices including whether firms compete over production quantity [13] or price [5]. In our model, as in these information-sharing models, firms' information is parameterized by its performance and degree of correlation. This allows us to reason about strategic decisions firms may make regarding shared data, model components, or predictive algorithms.

Theoretical and Empirical Models of Personalized Pricing. A growing body of empirical, theoretical, and legal literature considers how personalized pricing interacts with concepts like competition and privacy [4, 17, 9, 18, 55, 11, 38]. Most related to our work are theoretical models of personalized pricing in the context of competition. Both Rhodes and Zhou [39] and Baik and Larson [2] consider models of competition in personalized pricing under first-degree price discrimination, when firms can acquire perfect information about consumer valuations. In contrast, our model is designed to provide insights when firms have imperfect but potentially correlated information.

3 Model

We consider a duopoly model where two firms sell identical goods. For each consumer, a firm decides whether to offer a default price H^r or a discounted price L^r.[3] Both firms incur similar unit costs C, leading to a per-unit profit of $H = H^r - C$ and $L = L^r - C$ when pricing high and low, respectively. We ignore consumers whose valuation V for the good is less than L^r, and we define θ to be the fraction of consumers with valuation at least H^r.[4] We will use τ_H and τ_L to refer to consumers with valuations at least or strictly less than H^r respectively.[5]

Consumer behavior. Consumers can purchase from either firm. Under perfect Bertrand competition, each consumer would simply choose to purchase the lower-priced good. The economics literature often relaxes the perfect competition assumption such that firms that price higher experience non-zero demand [see, e.g., 53]. This may be because firms have finite supply, meaning consumers are forced to purchase at a higher price when the low price goods sell out, or because some consumers are lazy and take the first price they encounter that is below their valuation. We parameterize this model as follows: When a consumer of type τ_H is offered a price L^r by one firm and H^r by the other, they purchase at price H^r with probability $\sigma \in [0, 0.5]$ and L^r with probability $1 - \sigma$. Thus, for larger values of σ, consumers are less price-sensitive.

We assume that consumers never pay a price above their valuation and always make a purchase as long as at least one firm offers a price below their valuation. Further, when a consumer is offered two identical prices, they choose a firm to purchase from uniformly at random. An intuitive example of this consumer behavior is riders choosing between ridesharing apps. When prices are the same, potential riders are largely indifferent between two rideshare services (i.e., they do not have brand loyalty). However, when prices differ and consumers are willing to pay the higher price, σ models the friction consumers face in comparing the two options. Perhaps some consumers check both apps to shop for the lowest price, but others choose one app at random and take the first price.

[3] This is consistent with pricing via "couponing" [e.g., 17], a strategy according to which firms target offers of fixed discounts (e.g., 20% off) to consumers.

[4] We note that H^r and L^r are exogenous to this model. One interpretation is that L^r is the equilibrium price under perfect competition. In our model, duopolists can extract more surplus from certain consumers via pricing H^r.

[5] While a more sophisticated model might seek to directly estimate consumer willingness to pay,organizations may in practice simplify continuous prediction problems into discrete ones [34] and collect data only at discrete price points [16].

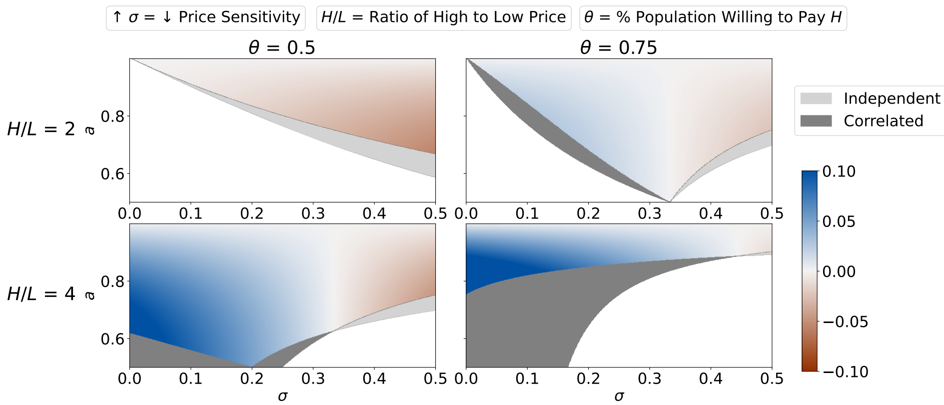

Figure 1: Regions where firms following the algorithm's recommendation is a Bayes Nash Equilibrium (BNE) for independent models only ($\rho = 0$, light gray), identical models only ($\rho = 1$, dark gray), and both (gradient). The gradient represents the difference in firm utility when $\rho = 1$ relative to $\rho = 0$; blue (red) signifies positive (negative) difference. Columns represent two values of $\theta \in \{0.5, 0.75\}$, while rows represent two values of $H/L \in \{2, 4\}$. The x-axis in each subfigure is σ and the y-axis is $a = a_1 = a_2$.

$[\tau_H]$ Firm 1	Firm 2: H	Firm 2: L
H	$(\frac{H}{2}, \frac{H}{2})$	$(\sigma H, (1-\sigma)L)$
L	$((1-\sigma)L, \sigma H)$	$(\frac{L}{2}, \frac{L}{2})$

$[\tau_L]$ Firm 1	Firm 2: H	Firm 2: L
H	$(0, 0)$	$(0, L)$
L	$(L, 0)$	$(\frac{L}{2}, \frac{L}{2})$

Table 1: Payoff matrices for both firms when the consumer is willing to pay the high price (τ_H, top) and low price (τ_L, bottom). Within each cell, we denote (Firm 1 payoff, Firm 2 payoff).

Firms' utility and information structure. Our consumer choice model yields the payoff matrices for the two firms for each consumer type shown in Table 1. Note that from firms' perspective, their utilities are with respect to unit profit as opposed to sale price. For ease of notation, we will drop the superscript and use H, L to refer to firms' pricing choices as well. We will denote $U_i(\cdot; \tau)$ as the utility/payoff for firm i for a given action profile for $\tau \in \{\tau_L, \tau_H\}$. For example, $U_1((H, L); \tau_H) = \sigma H$ and $U_2((H, L); \tau_H) = (1-\sigma)L$.

Firms do not have perfect information. Instead, we assume that when a consumer arrives with features x, each firm produces an algorithmic prediction $p_1(x), p_2(x) \in \{0, 1\}$ designed to segment users into types $\{\tau_L, \tau_H\}$. These algorithms are imperfect. For simplicity, we assume the algorithm has equal true positive and true negative rates, which we will denote a_1 for firm 1: $\mathbb{P}[p_1(x) = 1 \mid \tau_H] = \mathbb{P}[p_1(x) = 0 \mid \tau_L] = a_1$. We define the same quantity for firm 2 and will refer to a from hereon as the model's performance. We will drop x and simply refer to the algorithmic prediction as p_1, p_2.

An important feature of our model is that p_1 and p_2 need not be independent conditioned on user type. If, for example, both firms purchase data from a third party, their predictions may be correlated. In the extreme case of algorithmic monoculture, they may use the same model provided by a third party, meaning their predictions would be identical. We parameterize their correlation by $\rho \in [0, 1]$, where $\rho = 0$ implies independence ($p_1 \perp\!\!\!\perp p_2 \mid \tau$) and $\rho = 1$ implies maximal correlation.[6] When $a_1 = a_2$, $\rho = 1$ if and only if $p_1 = p_2$ deterministically. For now, we treat ρ as exogenous; we will consider strategic choices impacting ρ in Section 5. We assume all parameters are known to both firms.[7] In total, our model has five free parameters summarized in Table 2.

Equilibrium concept. A firm's strategy space is simple: for each segment given by the algorithm, set a price in $\{L, H\}$. Because all parameters are known, firms know the joint distribution on p_1, p_2, τ.

[6]Note that when $a_1 \neq a_2$, p_1 and p_2 cannot be perfectly correlated. See Appendix B for a formal definition of ρ.

[7]This assumption is especially common in oligopolies with few players that interact with each other frequently.

Parameter	Interpretation
$\theta \in [0,1]$	Frac. of consumers willing to pay H^r
$a_1, a_2 \in [0.5, 1]$	Model performance for firms 1 & 2
$\sigma \in [0, 0.5]$	Consumers' indifference to price
$\rho \in [0,1]$	Degree of model correlation

Table 2: List of free parameters in the model.

Thus, we assume that firms' strategies form a Bayes Nash Equilibrium (BNE). We do not require that firms price based on the algorithm's predictions; in fact, for some parts of the parameter space, firms may ignore the algorithm and either always or never offer the discount. We will focus on the region where both firms choose to follow their algorithms at equilibrium (i.e., price-discriminate), which is formally:

$$s^*(p) = \begin{cases} H, & \text{if } p = 1 \\ L, & \text{if } p = 0. \end{cases}$$

The strategy profile (s^*, s^*) (i.e., both firms price-discriminate) is an equilibrium if and only if the following conditions hold:

$$\mathop{\mathbb{E}}_{p_2,\tau}[U_1((H, s^*(p_2)); \tau) \mid p_1 = 1] \geq \mathop{\mathbb{E}}_{p_2,\tau}[U_1((L, s^*(p_2)); \tau) \mid p_1 = 1]$$
$$\mathop{\mathbb{E}}_{p_2,\tau}[U_1((L, s^*(p_2)); \tau) \mid p_1 = 0] \geq \mathop{\mathbb{E}}_{p_2,\tau}[U_1((H, s^*(p_2)); \tau) \mid p_1 = 0].$$

Analogous conditions must hold for firm 2. Intuitively, expected utility when both firms follow the algorithm's recommendation (both when $p_1 = 1$ and when $p_1 = 0$) must be higher than when one firm deviates.

4 Main Results

We find that *(1)* consumers are increasingly worse off as algorithms become more correlated; and *(2)* firms exhibit stronger preferences for correlation as consumers become more price sensitive.

Figure 1 shows regions where both firms following the algorithm's recommendation is a BNE for independent models only ($\rho = 0$, light gray), identical models only ($\rho = 1$, dark gray), and both (gradient), for various values of $\sigma, \theta, H/L$, and $a = a_1 = a_2$. Within the gradient region, we calculate the difference in firm utility between using perfectly correlated (i.e., identical) predictions ($\rho = 1$) and completely independent predictions ($\rho = 0$). Blue (red) indicates a positive difference, meaning that firms prefer using more correlated (more independent) predictions.

(1) Consumers are always worse off when pricing algorithms are correlated. When firms' pricing strategies are correlated (i.e., they price identically), consumers have less choice and must accept the given price or forgo the good. Conversely, when firms price independently, consumers often have the option to choose a lower price. We formalize this in Theorem 4.1.

Theorem 4.1. *Fix σ, a_1, a_2, θ, and H/L. For all ρ such that both firms following the algorithm's recommendation is a BNE, consumer welfare is decreasing in ρ.*

All proofs can be found in Appendix C. Our next few results describe when firms benefit by choosing correlated algorithms (and thereby harming consumers).

(2) Higher consumer price sensitivity leads to a stronger firm preference for correlation. As consumers become more price sensitive (i.e., σ decreases), firms increasingly prefer to use more correlated algorithms over more independent ones.

Theorem 4.2. *Suppose, for fixed θ, a_1, a_2, and H/L, both firms price-discriminating is a BNE when $\rho = \rho_A$ and $\rho = \rho_B$, with $\rho_B > \rho_A$. Assuming both $a_1, a_2 < 1$, firms prefer ρ_B when $\sigma < \frac{H\theta - L}{2\theta(H-L)}$ and otherwise prefer ρ_A.*

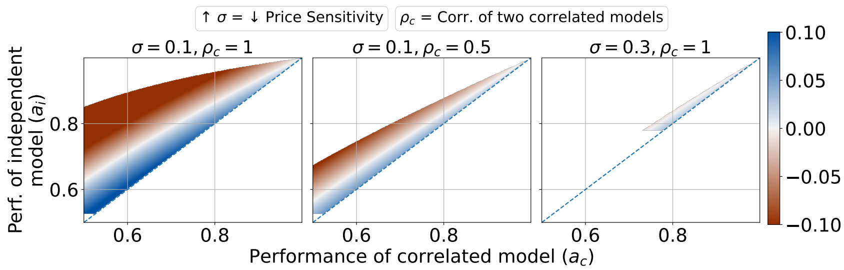

Figure 2: Regions where firms using both correlated models and independent models are Pure Nash Equilibra (first-stage game). An additional condition is that firms following the algorithm's recommendation must be a Bayes Nash Equilibrium (second-stage game). The x-axis is the performance of the correlated algorithm a_c, and the y-axis is the performance of the independent algorithm a_i. The gradient represents the difference in firm utility when $\rho = \rho_c$ (correlated) at performance a_c relative to the utility at $\rho = 0$ (independent) at performance a_i; blue (red) signifies positive (negative) difference. All subfigures show parameters for which firms have a preference for correlation at $a_c = a_i$ as per Theorem 4.1, with $H/L = 3, \theta = 0.75$.

Intuitively, when consumers are more price sensitive, firms have a higher risk in pricing H because they may get undercut by their competitor and only attain a small percentage of the market. In these situations, firms prefer correlation because there is no risk of undercutting; both firms are guaranteed to get the same prediction and therefore price the same way. On the other hand, conditioned on pricing low, firms *prefer independence*: a firm would rather be undercutting its competitor than pricing identically. The balance between these two competing forces—a preference for correlation when pricing high and a preference for independence when pricing low—determine whether a firm prefers correlation overall.

The tension between these forces is mediated by σ, which determines the relative risk from being undercut. Indeed, in Figure 1 we observe that within the gradient region (where both independent and correlated models are equilibria), preference for correlation monotonically decreases (from blue to red) as σ increases. In the extreme case when $\sigma = 0.5$ and so consumers are completely price insensitive, firms always prefer independence. When a firm predicts $p_i = 1$ and prices H accordingly, there is zero risk in being undercut: the firm receives $\sigma H = 0.5H$ if $\tau = \tau_H$ and 0 otherwise, regardless of their opponent's price. However, when a firm predicts $p_i = 0$ and prices L, they would in fact prefer that their opponent prices H so that they guarantee a sale when the consumer's valuation is low ($\tau = \tau_L$).

5 First-Stage Game: Strategic Correlation

Before two firms compete in price, they must first decide which algorithms to deploy. In this section, we analyze strategic decisions that impact the degree of correlation ρ between firms' predictions.

5.1 Model

Two firms choose between two model development processes. For ease of exposition, we will illustrate an example where firms choose to either *(1)* collect their own training data or *(2)* buy training data from the same vendor. Purchasing data or using their own data results in a model with performance a_c and a_i, respectively. (For simplicity, we will assume that both models yield the same performance.) In this section, we will abuse notation slightly and not refer to p as the algorithm; instead, we will use a to refer to the algorithm with its associated performance.

When both firms buy data, their models are correlated at some $\rho = \rho_c > 0$. When a firm collects their own data, we make a simplifying assumption that their models make errors that are independent from the errors of their competitor's model (i.e., $\rho = \rho_0 = 0$); in practice, independent data may not lead to completely uncorrelated errors. We also note that any shared component in the model development process – not just data procurement – can lead to correlated outcomes. For instance, our experiments in Section 6 give firms a choice between two model classes that produce varying levels of correlation.

To summarize, firms are faced with the following payoff matrix:

		Firm 2	
		a_c	a_i
Firm 1	a_c	$E_{\rho_c}(s^*(a_c), s^*(a_c))$	$E_{\rho_0}(s^*(a_c), s^*(a_i))$
	a_i	$E_{\rho_0}(s^*(a_i), s^*(a_c))$	$E_{\rho_0}(s^*(a_i), s^*(a_i))$

where, for instance,

$$E_{\rho_c}(s^*(a_c), s^*(a_c)) = (E^1_{\rho_c}(s^*(a_c), s^*(a_c)), E^2_{\rho_c}(s^*(a_c), s^*(a_c))$$
$$= (\mathop{\mathbb{E}}_{a_c,a_c,\tau;\rho=\rho_c}[U_1((s^*(a_c), s^*(a_c))); \tau)], \mathop{\mathbb{E}}_{a_c,a_c,\tau;\rho=\rho_c}[U_2((s^*(a_c), s^*(a_c))); \tau)]).$$

We are interested in analyzing conditions under which equilibria exist. Two possible equilibria are *(1)* both firms choose algorithm a_c with correlation ρ_c, and *(2)* both firms choose algorithm a_i, resulting in independent outcomes $\rho = \rho_0 = 0$. From hereon, **we will refer to scenario (1) and (2) respectively as "correlated" and "independent"**, ignoring the fact that other actions can also lead to independent outcomes.

Formally, the following condition must hold for both firms correlating to be a Pure Nash Equilibrium (PNE):

$$E^1_{\rho_c}(s^*(a_c), s^*(a_c)) \geq E^1_{\rho_0}(s^*(a_i), s^*(a_c)) \text{ and } E^2_{\rho_c}(s^*(a_c), s^*(a_c)) \geq E^2_{\rho_0}(s^*(a_c), s^*(a_i))$$

and similarly for both firms choosing independence:

$$E^1_{\rho_0}(s^*(a_i), s^*(a_i)) \geq E^1_{\rho_0}(s^*(a_c), s^*(a_i)) \text{ and } E^2_{\rho_0}(s^*(a_i), s^*(a_i)) \geq E^2_{\rho_0}(s^*(a_i), s^*(a_c)).$$

As with the previous section, we focus our attention on the strategy s^* of price-discriminating. As such, an additional necessary condition for equilibrium is that in the downstream second-stage game, both firms following the algorithm's recommendation is a BNE.

5.2 Results

Our main result is that under certain conditions, there always exists some regime where correlating is preferred to independence, even when the correlated algorithm has worse performance. Critically, as mentioned before, the following theorems only apply for parameter spaces where firms play the strategy s^* of price-discriminating. We will also define $\sigma^*(\theta, R) = \frac{R\theta-1}{2\theta(R-1)}$, where $R = \frac{H}{L}$, as the maximum threshold on σ for firms to prefer correlation, as per Theorem 4.2.

Theorem 5.1. *When $a_i > a_c$, both firms choosing independence is always a PNE.*

We next establish the conditions under which both firms correlating are in equilibrium, which comes directly from Theorem 4.2.

Corollary 5.2 (Corollary to Thm 4.2)**.** *If, at $a_i = a_c$, firms have a strict preference for correlation ($\sigma < \sigma^*(\theta, R)$), then correlating is strictly a PNE, i.e.,*

$$\exists\, \delta > 0 \text{ s.t. } E^i_{\rho_c}(s^*(a_c), s^*(a_c)) \geq E^i_{\rho_0}(s^*(a_i), s^*(a_c)) + \delta, \forall i \in \{1, 2\}.$$

With Theorem 5.1 and Corollary 5.2 in hand, we can now state our final result.

Theorem 5.3. *Suppose, at $a_i = a_c$, that $\sigma < \sigma^*(\theta, R)$. Then $\exists\, a_i > a_c$ such that both correlation and independence are PNE **and** correlation is preferred over independence.*

In other words, Theorem 5.3 says that given a preference for correlation at $a_i = a_c$, there are settings where firms are willing to sacrifice accuracy to gain correlation with one another. We will demonstrate this effect in Section 6.

Figure 2 shows the various regions where both correlation and independence are PNE. All subfigures show model parameters for which firms have a preference for correlation at $a_i = a_c$. As expected,

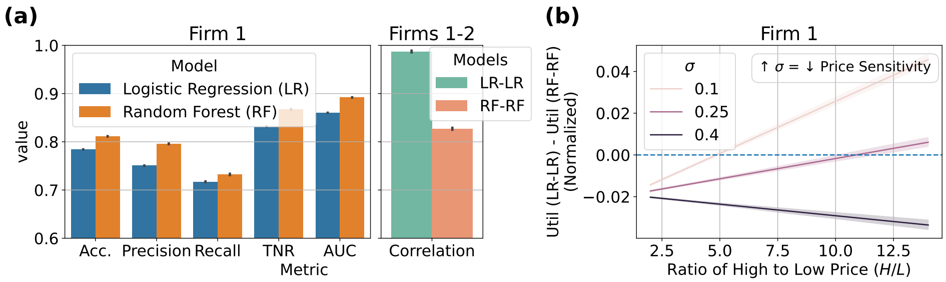

Figure 3: **(a)** [Left] Accuracy, precision, recall, true negative rate (TNR), and area under ROC curve (AUC) for Firm 1 deploying a logistic regression (LR) or random forest (RF) model. [Right] Correlation between both firms' models when they both use logistic regression (LR-LR) or both use random forests (RF-RF). Error bars indicate 95% confidence intervals over 15 seeds. **(b)** Utility when both firms use logistic regression models (LR-LR) subtracted by utility when both firms use random forests (RF-RF). Greater than 0 indicates a preference for correlation at the expense of predictive performance. x-axis varies the proportion of H (high price) to L (low price), and line colors indicate different values of σ, where a lighter color means higher consumer sensitivity to price. Shaded region indicates 95% confidence intervals over 15 seeds. **Results for Firm 2 are omitted because of symmetry** – since both firms face the same model options, their results are identically distributed.

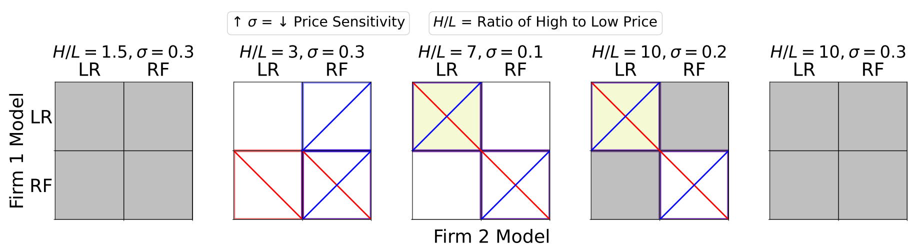

Figure 4: Best response matrices for the two firms where the action space is to deploy a logistic regression (LR) or random forest (RF) model, over five select model parameters. Best response for Firms 1 and 2 are highlighted in blue and red, respectively. Nash equilibria exist when both blue and red are highlighted in the same box (e.g., (LR, LR) in the middle subfigure). When both (LR, LR) and (RF, RF) are equilibria, a yellow square indicates higher firm utility between the two. Grey boxes are "invalid" regions because following the algorithm would not have been a BNE in the downstream game where firms compete on prices. These results use the average firm utility over 15 seeds.

all subfigures have a region at $a_i > a_c$ where correlation is still preferred to independence despite having a lower performance (blue gradient region). It seems that higher price sensitivity and a higher correlation option tend to increase the valid region of ϵ. For example, when $\theta = 0.75, H/L = 3, \sigma = 0.1$, and $\rho_c = 1$, firms would rather correlate at a performance of $a_c = 0.6$ than have a much more informative independent model of $a_i = 0.72$.

6 Empirical Study

We now demonstrate our theoretical results in a stylized game between two firms who are predicting income based on demographic attributes. We use ACSIncome data [15], which contains US Census data from 2018. The task is to predict whether or not a person's income is greater than $50,000.

6.1 Setup: Different Model Classes

Our experiment involves two firms. Each firm chooses between a better performing (random forests) and worse performing (logistic regression) model. However, logistic regression – despite having worse performance – has lower variance, meaning that it is likely to be more correlated when the opposing firm also chooses the same model class. We will show empirically that firms may prefer to sacrifice predictive performance in exchange for correlation, which leads to lower consumer welfare.

Both firms train and test on Census data in California. The test set is 30% of the data ($n = 58,700$) and is fixed across both firms. We randomly split half of the remaining 70% as the training set for Firm 1, and the other half for Firm 2, each having 35% of the entire data to train ($n = 68,482$). We repeat the training data splits over 15 random seeds.

Figure 3(a) shows the performance for both firms when training one logistic regression (LR) model and one random forest (RF) model, as well as the correlation between the two firms when they both employ the same model class. Note that RF outperform the LR across many performance metrics: accuracy, precision, recall, true negative rate (TNR), and area under ROC curve (AUC). This is by design – our goal is to simulate a scenario where firms have a choice between a more correlated model with worse performance (LR) and a better performing but less correlated model (RF). See Appendix D.1 for additional details. This model multiplicity setup is one of many ways firms can correlate their models. Refer to Appendix D.2, where firms can independently choose to procure the same third-party data to correlate their outcomes.

6.2 Results

Preference for Correlation. Figure 3(b) shows the difference in utility when both firms use logistic regression models (LR-LR) subtracted by the utility when both firms use random forests (RF-RF), over multiple values of H/L and σ. A positive difference indicates a preference for correlation. We observe that both firms tend to prefer correlation more when σ is low (consumers are more price sensitive) and the ratio between H and L prices is large. As per Theorem 4.2, correlation is most beneficial to firms when there is a high risk of being undercut by the opponent; therefore firms would rather have certainty about the other firm's actions than a better performing model.

When Correlation is in Equilibrium. We next model firms' *choice* of algorithms in a first-stage game. Figure 4 shows best response matrices for both firms when given the option to deploy a logistic regression (LR) or a random forest (RF), over various values of σ and H/L. Cells with a blue and red cross indicate a Pure Nash Equilibrium (PNE) for that action profile. In the extremes, when H/L is too low or too high, firms will never choose to follow their personalized pricing algorithms to begin with (grey cells) because always pricing low or high will give a higher expected utility. When H/L is moderate, less correlation (RF, RF) is always a PNE as per Theorem 5.1. More correlation (LR, LR) is a PNE under the condition outlined in Corollary 5.2. Finally, when both (LR, LR) and (RF, RF) are PNE, the difference in performance between LR and RF are small enough such that (LR, LR) is higher in utility (yellow cell) than (RF, RF) as per Theorem 5.3.

7 Discussion

Correlation is a mechanism to sustain higher prices. When firms make up a duopoly, using more correlated algorithms allows firms to better price discriminate and reduce competition, which increases prices. When algorithms are not correlated, firms naturally attempt to undercut their opponent in order to extract more surplus and the high price equilibrium cannot be sustained. This undercutting will continue to lower prices until firms reach a new equilibrium.

Firms prefer lower variance models under competition. Lower variance models have less predictive multiplicity [6], and thus predictive errors are more correlated. Our empirical study suggests that in the midst of competition, firms are pushed to adopt simpler models (i.e., higher bias, lower variance) on the margins.

Correlation of any form may induce collusive outcomes. Models can become correlated in a variety of ways, such as using the same pre-trained model or training on similar data. In Appendix D.2, we demonstrate that firms may tend to buy the same third-party data that degrades their models' performance in exchange for higher correlation, resulting in collusive outcomes.

Publicly signalling choice of model may prompt collusion. Expanding on [21], [28, 454] suggests that adopting a pricing algorithm that transparently "broadcasts" its intentions can be an invitation for other firms to collude on price. Currently there is little precedent for broadcasting intention by publicly adopting an inferior model. However, our model suggests that this may be a risk. Assume that both correlated models (LR, LR) and independent models (RF, RF) are equilibria (e.g., middle subfigure of Figure 2) and assume further that firms currently use (RF, RF). In order for firms to reach the collusive outcome of (LR, LR) without explicit communication or agreement, one firm would need to switch to LR at the expense of their own utility (because they are leaving an equilibrium state) in the hopes that their opponent will follow suit. This action could be treated as a form of signalling. Sending a (costly) signal by adopting an inferior model suggests that a firm took an action that was against its own (short-term) economic self-interest but eventually led to collusive outcomes. Antitrust law may need to determine whether public announcement of model choice can be anti-competitive in the same way that public announcement of intent to price high can be anti-competitive.

8 Conclusion

Taken together, these results suggest that firms will sometimes prefer a less accurate personalized pricing algorithm when doing so allows them to better correlate their behavior with their competitors, per Theorem 5.3. This behavior reduces consumer welfare, per Theorem 4.1. Furthermore, as firms are more likely to prefer correlated algorithms when consumers are price sensitive per Theorem 4.2, the consumers most likely to suffer are those to whom the price matters most. Our results add to the growing body of scholarship suggesting that the ease of collusion that algorithmic price-setting facilitates may support a revision of traditional anti-trust standards [30, 19, 28].

References

[1] Cédric Argenton, Radosveta Ivanova-Stenzel, and Wieland Müller. Cournot meets bayes-nash: A discontinuity in behavior in finitely repeated duopoly games. *Games and Economic Behavior*, 144:1–12, 2024.

[2] Alicia Baik, Simon Anderson and Nathan Larson. Price discrimination in the information age: Prices, poaching, and privacy with personalized targeted discounts. *The Review of Economic Studies*, 90(5):2085–2115, 2023.

[3] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. Dynamic pricing in ridesharing platforms. *ACM SIGecom Exchanges*, 15(1):65–70, 2016.

[4] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–957, 2015.

[5] Joseph Bertrand. Théorie des richesses: revue de 'théories mathématiques de la richesse sociale' de l. walras and 'recherches sur les principes mathématiques de la théorie des richesses' de a. cournot. *Journal des Savants*, 1883.

[6] Emily Black, Manish Raghavan, and Solon Barocas. Model multiplicity: Opportunities, concerns, and solutions. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 850–863, 2022.

[7] Rishi Bommasani, Kathleen A Creel, Ananya Kumar, Dan Jurafsky, and Percy S Liang. Picking on the same person: Does algorithmic monoculture lead to outcome homogenization? In *Advances in Neural Information Processing Systems*, volume 35, pages 3663–3678, 2022.

[8] Heather Boushey and Helen Knudsen. The importance of competition for the american economy, 2021. URL `https://www.whitehouse.gov/cea/written-materials/2021/07/09/the-importance-of-competition-for-the-american-economy/`.

[9] Zhijun Chen, Chongwoo Choe, and Noriaki Matsushima. Competitive personalized pricing. *Management Science*, 66(9):4003–4023, 2020.

[10] Richard N Clarke. Duopolists don't wish to share information. *Economics Letters*, 11(1-2): 33–36, 1983.

[11] Stefano Colombo, Luigi Filippini, and Aldo Pignataro. Information sharing, personalized pricing, and collusion. *Information Economics and Policy*, 66:101032, 2024.

[12] CEA Council of Economic Advisors. Big data and differential pricing, 2015.

[13] Antoine Augustin Cournot. *Mathematical principles of the theory of wealth*. Paris: L. Hachette, 1838.

[14] Shidi Deng, Maximilian Schiffer, and Martin Bichler. Algorithmic collusion in dynamic pricing with deep reinforcement learning, 2024.

[15] Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring adult: New datasets for fair machine learning. *Advances in neural information processing systems*, 34:6478–6490, 2021.

[16] Jean-Pierre Dubé and Sanjog Misra. Personalized pricing and consumer welfare. *Journal of Political Economy*, 131(1):131–189, 2023.

[17] Jean-Pierre Dubé, Zheng Fang, Nathan Fong, and Xueming Luo. Competitive price targeting with smartphone coupons. *Marketing Science*, 36(6):944–975, 2017.

[18] Adam N Elmachtoub, Vishal Gupta, and Michael L Hamilton. The value of personalized pricing. *Management Science*, 67(10):6055–6070, 2021.

[19] Ariel Ezrachi and Maurice E. Stucke. Sustainable and unchallenged algorithmic tacit collusion. *Northwestern Journal of Technology & Intellectual Property*, (217):255–58, 2020.

[20] FTC. Competition counts: How consumers win when businesses compete. Technical report, FTC, 2015.

[21] FTC. Ftc: Protecting america's consumers. August 2016. URL `https://www.ftc.gov/news-events/blogs/competition-matters/2016/08/compliance-check-collaborators-who-also-compete`.

[22] Esther Gal-Or. Information sharing in oligopoly. *Econometrica: Journal of the Econometric Society*, pages 329–343, 1985.

[23] Jason D Hartline, Sheng Long, and Chenhao Zhang. Regulation of algorithmic collusion. In *Proceedings of the Symposium on Computer Science and Law*, pages 98–108, 2024.

[24] Shomik Jain, Vinith Suriyakumar, and Ashia Wilson. Algorithmic pluralism: A structural approach towards equal opportunity. *arXiv preprint arXiv:2305.08157*, 2023.

[25] Timo Klein. Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics*, 52(3):538–558, 2021.

[26] Jon Kleinberg and Manish Raghavan. Algorithmic Monoculture and Social Welfare. *Proceedings of the National Academy of Sciences*, 118(22):e2018340118, 2021.

[27] William E Kovacic, Robert C Marshall, Leslie M Marx, and Halbert L White. Plus factors and agreement in antitrust law. *Mich. L. Rev.*, 110:393, 2011.

[28] Aneesa Mazumdar. Algorithmic collusion: Reviving section 5 of the ftc act. *Columbia Law Review*, 122(2):449–488, March 2022. URL `https://www.jstor.org/stable/27114356`.

[29] R Preston McAfee and Vera Te Velde. Dynamic pricing in the airline industry. *Handbook on economics and information systems*, 1:527–67, 2006.

[30] Salil K. Mehra. Antitrust and the robo-seller: Competition in the time of algorithms. *Minnesota Law Review*, (100), 2016.

[31] Renato Nazzini and James Henderson. Overcoming the current knowledge gap of algorithmic" collusion" and the role of computational antitrust. *Stan. Computational Antitrust*, 4:1, 2024.

[32] William Novshek and Hugo Sonnenschein. Fulfilled expectations cournot duopoly with information acquisition and release. *The Bell Journal of Economics*, pages 214–218, 1982.

[33] Vilfredo Pareto. *Manuel d'économie politique*, volume 38. Giard & Brière, 1909.

[34] Samir Passi and Solon Barocas. Problem formulation and fairness. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 39–48, 2019.

[35] Katie Pederson. How companies use personal data to charge different people different prices for the same product. *CBC*, 2017. URL `https://www.cbc.ca/news/business/marketplace-online-prices-profiles-1.4414240`.

[36] Kenny Peng and Nikhil Garg. Monoculture in matching markets. *arXiv preprint arXiv:2312.09841*, 2023.

[37] Clemens Possnig. *Reinforcement learning and collusion*. Department of Economics, University of Waterloo, 2023.

[38] Liying Qiu, Yan Huang, Param Vir Singh, and Kannan Srinivasan. Does personalization in product rankings facilitate or mitigate algorithmic pricing collusion? *SSRN https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4132555*, 2023.

[39] Andrew Rhodes and Jidong Zhou. Personalized pricing and competition. *American Economic Review*, forthcoming.

[40] Dan Sasaki. The value of information in oligopoly with demand uncertainty. *Journal of Economics*, 73:1–23, 2001.

[41] Adam Smith. *An inquiry into the nature and causes of the wealth of nations*. London: printed for W. Strahan; and T. Cadell, 1776., 1776.

[42] Maurice E Stucke and Ariel Ezrachi. Antitrust, algorithmic pricing and tacit collusion. In *Research Handbook on the Law of Artificial Intelligence*, pages 624–648. Edward Elgar Publishing, 2018.

[43] Jacques-Francois Thisse and Xavier Vives. On the strategic choice of spatial price policy. *The American Economic Review*, pages 122–137, 1988.

[44] Seventh Circuit United States Court of Appeals. United states v. socony-vacuum oil co., 310 u.s. 150, 1940.

[45] Seventh Circuit United States Court of Appeals. In re text messaging antitrust litigation, 782 f.3d 867, 2015.

[46] District of Maryland United States District Court. Hecox et al. v. doordash, inc., 2023.

[47] District of Nevada United States District Court. Gibson v. mgm resorts int'l, 2023 u.s. dist. lexis 190442, 2023.

[48] Nashville Division United States District Court, Middle District of Tennessee. In re realpage, inc., rental software antitrust litig. (no. ii), 2023 u.s. dist. lexis 230200, 2023.

[49] United States Supreme Court. Standard oil co. of new jersey v. united states, 221 u.s. 1, 1910.

[50] United States Supreme Court. United states v. topco assocs., inc., 405 u.s. 596, 1972.

[51] Hal R Varian. Price discrimination. *Handbook of industrial organization*, 1:597–654, 1989.

[52] Xavier Vives. Duopoly information equilibrium: Cournot and bertrand. *Journal of economic theory*, 34(1):71–94, 1984.

[53] Xavier Vives. Rationing rules and bertrand-edgeworth equilibria in large markets. *Economics Letters*, 21(2):113–116, 1986.

[54] Robert M Weiss and Ajay K Mehrotra. Online dynamic pricing: Efficiency, equity and the future of e-commerce. *Va. JL & Tech.*, 6:1, 2001.

[55] Ramsi A Woodcock. Personalized pricing as monopolization. *Conn. L. Rev.*, 51:311, 2019.

A Other Related Work

Algorithmic (Tacit) Collusion. In general, legal scholars consider three broad mechanisms for algorithmic collusion. First, an algorithm can act as a tool that aids humans in explicitly sustaining cartel-like behavior. Second, an algorithm can be a hub that coordinates actions, or be the sole algorithm used among competitors. Third, highly sophisticated algorithms can learn each other's behaviors and collectively achieve supra-competitive prices without explicit communication. From the first to third category, the likelihood that the behavior is illegal decreases or, at best, the action becomes more likely to fall into a contested grey area [42]. This is because humans become less involved in achieving collusion, making it harder to prove an intent or conspiracy to agree to fix prices. Recent legal scholarship has raised concerns about the potential for algorithms to facilitate tacit collusion, which falls somewhere between the second and third category. Various works have proposed legal and legislative pathways to expand the powers of regulatory agencies [28] or methods to more effectively screen and audit for tacit collusion [31, 23].

However, the mechanisms for algorithmic tacit collusion have not been extensively studied. Several theoretical papers have found collusive outcomes under the third mechanism for algorithmic collusion, where sufficiently sophisticated reinforcement learning models interact and compete in prices over time [25, 37, 14]. Our work formally analyzes a novel mechanism for tacit collusion based on the widespread correlation of algorithmic decisions.

B Model (Continued)

B.1 Correlation Parameter

We parameterize correlation between two models p_1 and p_2 with $\rho \in [0, 1]$, see Table 3 for the joint distribution on $\mathbb{P}[p_1, p_2, \tau]$. Note that in the case where $\rho = 1$ and $a_1 = a_2$, we are modeling a scenario where both firms are using the same algorithm (i.e., monoculture).

τ	p_1	p_2	$\mathbb{P}[p_1, p_2, \tau]$
τ_H	1	1	$\theta[a_1a_2 + \rho(\min(a_1, a_2) - a_1a_2)]$
τ_H	1	0	$\theta[a_1(1 - a_2) - \rho(\min(a_1, a_2) - a_1a_2)]$
τ_H	0	1	$\theta[(1 - a_1)a_2 - \rho(\min(a_1, a_2) - a_1a_2)]$
τ_H	0	0	$\theta[1 - a_1 - a_2 + a_1a_2 + \rho(\min(a_1, a_2) - a_1a_2)]$
τ_L	1	1	$(1 - \theta)[1 - a_1 - a_2 + a_1a_2 + \rho(\min(a_1, a_2) - a_1a_2)]$
τ_L	1	0	$(1 - \theta)[(1 - a_1)a_2 - \rho(\min(a_1, a_2) - a_1a_2)]$
τ_L	0	1	$(1 - \theta)[a_1(1 - a_2) - \rho(\min(a_1, a_2) - a_1a_2)]$
τ_L	0	0	$(1 - \theta)[a_1a_2 + \rho(\min(a_1, a_2) - a_1a_2)]$

Table 3: Joint distribution $\mathbb{P}[p_1, p_2, \tau]$.

C Proofs

C.1 Consumer Welfare and Proof of Theorem 4.1

Before proving the theorem, we will introduce some additional notation. Let $W((\cdot); \tau)$ denote consumer welfare under the action profile $(\cdot)$ and demand state τ. We define consumer welfare as the consumer valuation of the good subtracted by the cost of the good. As such, we define two additional variables V_L, V_H to be consumers' expected valuation under τ_L and τ_H, respectively. Let $\delta_L = V_L - L^r$ and similarly for $\delta_H = V_H - H^r$. By definition, $\delta_L, \delta_H \geq 0$ – otherwise, consumers will not purchase the good. Consumer welfare under the various actions and demand states can be summarized in Table 4.

τ_H

Firm 1 \ Firm 2	H	L
H	δ_H	$\delta_H + (1-\sigma)(H^r - L^r)$
L	$\delta_H + (1-\sigma)(H^r - L^r)$	$\delta_H + (H^r - L^r)$

τ_L

Firm 1 \ Firm 2	H	L
H	0	δ_L
L	δ_L	δ_L

Table 4: Consumer welfare under all action possibilities and both demand states (τ_H, τ_L).

Proof. We will denote expected consumer welfare for a given value of ρ as

$$\mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau)].$$

Our goal is to show that

$$\frac{d}{d\rho} \mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau)] < 0.$$

Our approach will be to show that increasing ρ increases the likelihood that $p_1 = p_2$, which in turn reduces consumer welfare. First, observe that

$$\begin{aligned}\mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau)] &= \mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 = p_2] \mathop{\Pr}_{p_1,p_2,\tau;\rho}[p_1 = p_2] \\ &+ \mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 \neq p_2] \mathop{\Pr}_{p_1,p_2,\tau;\rho}[p_1 \neq p_2].\end{aligned}$$

We will show that

$$\mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 = p_2]$$

and

$$\mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 \neq p_2]$$

do not depend on ρ.

$$\begin{aligned}\mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 = p_2] &= \mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 = p_2, \tau = \tau_H] \mathop{\Pr}_{p_1,p_2,\tau;\rho}[\tau_H \mid p_1 = p_2] \\ &+ \mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 = p_2, \tau = \tau_L] \mathop{\Pr}_{p_1,p_2,\tau;\rho}[\tau_L \mid p_1 = p_2] \\ &= W((s^*(p_1), s^*(p_2)); \tau_H) \mathop{\Pr}_{p_1,p_2,\tau;\rho}[\tau_H \mid p_1 = p_2] \\ &+ W((s^*(p_1), s^*(p_2)); \tau_L) \mathop{\Pr}_{p_1,p_2,\tau;\rho}[\tau_L \mid p_1 = p_2]\end{aligned}$$

Note that by definition, $\Pr_{p_1,p_2;\rho}[p_1 = p_2 \mid \tau = \tau_H] = \Pr_{p_1,p_2;\rho}[p_1 = p_2 \mid \tau = \tau_L] = \Pr_{p_1,p_2,\tau;\rho}[p_1 = p_2]$. Therefore,

$$\begin{aligned}\mathop{\Pr}_{p_1,p_2,\tau;\rho}[\tau = \tau_L \mid p_1 = p_2] &= \frac{\Pr_{\tau;\rho}[\tau = \tau_L] \Pr_{p_1,p_2;\rho}[p_1 = p_2 \mid \tau = \tau_L]}{\Pr_{p_1,p_2;\rho}[p_1 = p_2]} = 1 - \theta \\ \mathop{\Pr}_{p_1,p_2,\tau;\rho}[\tau = \tau_H \mid p_1 = p_2] &= \frac{\Pr_{\tau;\rho}[\tau = \tau_H] \Pr_{p_1,p_2;\rho}[p_1 = p_2 \mid \tau = \tau_H]}{\Pr_{p_1,p_2;\rho}[p_1 = p_2]} = \theta\end{aligned}$$

This implies

$$\mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 = p_2] = \theta W((s^*(p_1), s^*(p_2)); \tau_H) + (1-\theta) W((s^*(p_1), s^*(p_2)); \tau_L).$$

A similar argument shows that $\mathbb{E}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 \neq p_2]$ does not depend on ρ. Next, we will show that that

$$\mathop{\mathbb{E}}_{p_1,p_2,\tau}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 = p_2] \leq \mathop{\mathbb{E}}_{p_1,p_2,\tau}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 \neq p_2], \quad (1)$$

meaning that consumers have higher expected utility when offered different prices. Because τ is independent of the event $p_1 = p_2$, we can analyze each $\tau \in \{\tau_L, \tau_H\}$ separately. For τ_H,

$$\begin{aligned}
\mathop{\mathbb{E}}_{p_1,p_2}[W((s^*(p_1), s^*(p_2)); \tau_H) \mid p_1 = p_2, \tau = \tau_H] - \delta_H &= \Pr_{p_1,p_2}[p_1 = p_2 = 1 \mid p_1 = p_2, \tau = \tau_H](W((H,H); \tau_H) - \delta_H) \\
&+ \Pr_{p_1,p_2}[p_1 = p_2 = 0 \mid p_1 = p_2, \tau = \tau_H](W((L,L); \tau_H) - \delta_H) \\
&= \Pr_{p_1,p_2}[p_1 = p_2 = 1 \mid p_1 = p_2, \tau = \tau_H] \cdot 0 \\
&+ \Pr_{p_1,p_2}[p_1 = p_2 = 0 \mid p_1 = p_2, \tau = \tau_H](H^r - L^r) \\
&= \frac{1 - a_1 - a_2 + (1-\rho)a_1 a_2 + \rho \min(a_1, a_2)}{1 - a_1 - a_2 + 2(1-\rho)a_1 a_2 + 2\rho \min(a_1, a_2)}(H^r - L^r) \\
&\le \frac{1}{2}(H^r - L^r)
\end{aligned}$$

because a_1 and a_2 are both at least 0.5. Similarly,

$$\begin{aligned}
\mathop{\mathbb{E}}_{p_1,p_2}[W((s^*(p_1), s^*(p_2)); \tau_H) \mid p_1 \neq p_2, \tau = \tau_H] - \delta_H &= (1-\sigma)(H^r - L^r) \\
&\ge \frac{1}{2}(H^r - L^r)
\end{aligned}$$

since $\sigma \le 0.5$, meaning

$$\mathop{\mathbb{E}}_{p_1,p_2}[W((s^*(p_1), s^*(p_2)); \tau_H) \mid p_1 = p_2, \tau = \tau_H] \le \mathop{\mathbb{E}}_{p_1,p_2}[W((s^*(p_1), s^*(p_2)); \tau_H) \mid p_1 \neq p_2, \tau = \tau_H], \tag{2}$$

Next, observe that

$$\mathop{\mathbb{E}}_{p_1,p_2}[W((s^*(p_1), s^*(p_2)); \tau_L) \mid p_1 = p_2, \tau = \tau_L] \le \mathop{\mathbb{E}}_{p_1,p_2}[W((s^*(p_1), s^*(p_2)); \tau_L) \mid p_1 \neq p_2, \tau = \tau_L] \tag{3}$$

simply because the left hand side is at most δ_L and the right hand side is deterministically δ_L. Combining (2) and (3) and using the fact that τ is independent of the event $p_1 = p_2$ proves (1). As a result,

$$\begin{aligned}
\frac{d}{d\rho} \mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho}[W((s^*(p_1), s^*(p_2)); \tau)] &= \frac{d}{d\rho}\Bigg(\mathop{\mathbb{E}}_{p_1,p_2,\tau}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 = p_2] \Pr_{p_1,p_2,\tau;\rho}[p_1 = p_2] \\
&+ \mathop{\mathbb{E}}_{p_1,p_2,\tau}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 \neq p_2] \left(1 - \Pr_{p_1,p_2,\tau;\rho}[p_1 = p_2]\right)\Bigg) \\
&= \Bigg(\mathop{\mathbb{E}}_{p_1,p_2,\tau}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 = p_2] \\
&- \mathop{\mathbb{E}}_{p_1,p_2,\tau}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 \neq p_2]\Bigg) \frac{d}{d\rho} \Pr_{p_1,p_2,\tau;\rho}[p_1 = p_2] \\
&= \Bigg(\mathop{\mathbb{E}}_{p_1,p_2,\tau}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 = p_2] \\
&- \mathop{\mathbb{E}}_{p_1,p_2,\tau}[W((s^*(p_1), s^*(p_2)); \tau) \mid p_1 \neq p_2]\Bigg) \\
&\cdot \frac{d}{d\rho} 1 - a_1 - a_2 + (1-\rho)a_1 a_2 + \rho \min(a_1, a_2) \\
&\le 0,
\end{aligned}$$

where the last line follows by (1) and using the fact that

$$\frac{d}{d\rho} 1 - a_1 - a_2 + (1-\rho)a_1 a_2 + \rho \min(a_1, a_2) \ge 0.$$

This inequality is strict as long as $a_1, a_2 < 1$ (otherwise ρ has no impact on the joint distribution of p_1, p_2, τ). □

C.2 Proof of Theorem 4.2

Proof. The following condition for Firm 1 must hold for them to prefer prefer $\rho = \rho_B$ over $\rho = \rho_A$:

$$\mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho=\rho_B}[U_1((s^*(p_1), s^*(p_2)); \tau)] > \mathop{\mathbb{E}}_{p_1,p_2,\tau;\rho=\rho_A}[U_1((s^*(p_1), s^*(p_2)); \tau)].$$

For ease of notation, let $A = \min(a_1, a_2) - a_1 a_2$ and let $\Delta_\rho = \rho_B - \rho_A > 0$. We will see that the probabilities cancel out when subtracting $\rho = \rho_B$ to $\rho = \rho_A$, leaving only the A and Δ_ρ terms:

$$\sum_{p_1\in\{0,1\}} \sum_{p_2\in\{0,1\}} \sum_{\tau\in\{\tau_H,\tau_L\}} U_1[(s^*(p_1), s^*(p_2); \tau] \left[\mathbb{P}[p_1, p_2, \tau; \rho = \rho_B] - \mathbb{P}[p_1, p_2, \tau; \rho = \rho_A]\right] > 0$$

$$\frac{H\theta A\Delta_\rho}{2} - H\sigma\theta A\Delta_\rho + \frac{L(1-\theta)A\Delta_\rho}{2} - L(1-\theta)A\Delta_\rho - L\theta(1-\sigma)A\Delta_\rho + \frac{L\theta A\Delta_\rho}{2} > 0$$

$$\frac{1}{2}A\Delta_\rho[H\theta(1-2\sigma) + L(2\sigma\theta - 1)] > 0$$

We can derive the same exact inequality for firm 2. When $\min(a_1, a_2) - a_1 a_2 \neq 0$ (or, when both $a_1, a_2 < 1$), we get

$$\sigma < \frac{H\theta - L}{2\theta(H-L)}.$$

We can further show that lower σ monotonically increases preference for correlation:

$$\frac{\partial}{\partial\sigma} A\Delta_\rho[H\theta(1-2\sigma) + L(2\sigma\theta - 1) = 2A\Delta_\rho\theta(L-H),$$

which is always negative because $L < H$ by definition. □

C.3 Proof of Theorem 5.1

Proof. The main intuition behind this proof is that an algorithm with performance a_i can simulate an algorithm with lower performance a_c. Recall that we define s^* to be the optimal strategy of following the algorithm. Let $s^{\sim *}$ be the strategy of doing the opposite of the algorithm's recommendations. We define s' to be the following strategy:

$$s'(a_i) = \begin{cases} s^*(a_i), & \text{w.p. } q \\ s^{\sim *}(a_i), & \text{w.p. } 1-q, \end{cases}$$

where $q = \frac{a_c + a_i - 1}{2a_i - 1}$. The strategy $s'(a_i)$ is equivalent in expectation to $s^*(a_c)$ in terms of firm utility. To see this, **we will prove that the conditional distribution** $\mathbb{P}[\tau|s'(a_i)]$ **is equivalent to** $\mathbb{P}[\tau|s^*(a_c)]$:

$$\mathbb{P}[\tau = \tau_H|s'(a_i) = 1] = \mathbb{P}[\tau = \tau_H|s^*(a_c) = 1]$$
$$\frac{\mathbb{P}[s'(a_i) = 1|\tau = \tau_H]\mathbb{P}[\tau_H]}{\mathbb{P}[s'(a_i) = 1|\tau = \tau_H]\mathbb{P}[\tau_H] + \mathbb{P}[s'(a_i) = 1|\tau = \tau_L]\mathbb{P}[\tau_L]} = \frac{\mathbb{P}[s^*(a_c) = 1|\tau = \tau_H]\mathbb{P}[\tau_H]}{\mathbb{P}[s^*(a_c) = 1|\tau = \tau_H]\mathbb{P}[\tau_H] + \mathbb{P}[s^*(a_c) = 1|\tau = \tau_L]\mathbb{P}[\tau_L]}$$

and

$$\mathbb{P}[\tau = \tau_H|s'(a_i) = 0] = \mathbb{P}[\tau = \tau_H|s^*(a_c) = 0]$$
$$\frac{\mathbb{P}[s'(a_i) = 0|\tau = \tau_H]\mathbb{P}[\tau_H]}{\mathbb{P}[s'(a_i) = 0|\tau = \tau_H]\mathbb{P}[\tau_H] + \mathbb{P}[s'(a_i) = 0|\tau = \tau_L]\mathbb{P}[\tau_L]} = \frac{\mathbb{P}[s^*(a_c) = 0|\tau = \tau_H]\mathbb{P}[\tau_H]}{\mathbb{P}[s^*(a_c) = 0|\tau = \tau_H]\mathbb{P}[\tau_H] + \mathbb{P}[s^*(a_c) = 0|\tau = \tau_L]\mathbb{P}[\tau_L]}.$$

Based on the Bayes' Rule expansion above, it suffices to prove the following equivalences:

$$\mathbb{P}[s'(a_i) = 1|\tau = \tau_H] = \mathbb{P}[s^*(a_c) = 1|\tau = \tau_H] \tag{4}$$

$$\mathbb{P}[s'(a_i) = 1|\tau = \tau_L] = \mathbb{P}[s^*(a_c) = 1|\tau = \tau_L] \tag{5}$$

Proof of (4):

$$\begin{aligned}
\mathbb{P}[s'(a_i) = 1|\tau = \tau_H] &= q\mathbb{P}[s^*(a_c) = 1|\tau = \tau_H] + (1-q)\mathbb{P}[s^{\sim *}(a_c) = 1|\tau = \tau_H] \\
&= q\mathbb{P}[s^*(a_c) = 1|\tau = \tau_H] + (1-q)\mathbb{P}[s^*(a_c) = 0|\tau = \tau_H] \\
&= qa_i + (1-q)(1-a_i) = \frac{a_c + a_i - 1}{2a_i - 1}a_i + \frac{a_i - a_c}{2a_i - 1}(1-a_i) = \frac{a_c(2a_i - 1)}{2a_i - 1} = a_c \\
&= \mathbb{P}[s^*(a_c) = 1|\tau = \tau_H]
\end{aligned}$$

Proof of (5):

$$\begin{aligned}
\mathbb{P}[s'(a_i) = 1|\tau = \tau_L] &= q\mathbb{P}[s^*(a_c) = 1|\tau = \tau_L] + (1-q)\mathbb{P}[s^{\sim *}(a_c) = 1|\tau = \tau_L] \\
&= q\mathbb{P}[s^*(a_c) = 1|\tau = \tau_L] + (1-q)\mathbb{P}[s^*(a_c) = 0|\tau = \tau_L] \\
&= (1-q)a_i + q(1-a_i) = \frac{a_i - a_c}{2a_i - 1}a_i + \frac{a_c + a_i - 1}{2a_i - 1}(1-a_i) = \frac{(2a_i - 1)(1-a_c)}{2a_i - 1} = 1 - a_c \\
&= \mathbb{P}[s^*(a_c) = 1|\tau = \tau_L]
\end{aligned}$$

Note that the space of possible accuracies is $a \geq 0.5$ for an algorithm to be useful. When $a_c = 0.5$, $a_i > 0.5$ by assumption of the Theorem and therefore q is never undefined. Then,

$$E^1_{\rho_0}[(s^*(a_i), s^*(a_i))] \geq E^1_{\rho_0}[(s'(a_i), s^*(a_i))] = E^1_{\rho_0}[(s^*(a_c), s^*(a_i))],$$

and similarly for Firm 2.

□

C.4 Proof of Corollary 5.2

Proof. We will show that the condition for a strict preference for correlation (in the second-stage game) is equivalent to correlation being strictly in equilibrium (in the first-stage game). We first start with the preference correlation in the proof for Theorem 4.2:

$$\mathop{\mathbb{E}}_{p_1, p_2, \tau; \rho = \rho_B}[U_1((s^*(p_1), s^*(p_2)); \tau)] > \mathop{\mathbb{E}}_{p_1, p_2, \tau; \rho = \rho_A}[U_1((s^*(p_1), s^*(p_2)); \tau)].$$

Since this condition is for any $\rho_A < \rho_B$, we will let $\rho_B = \rho_c$ and $\rho_A = 0$. Further, we will change the p notation to a_c and a_i where relevant, since $a_i = a_c$ by assumption.

$$\mathop{\mathbb{E}}_{a_c, a_c, \tau; \rho = \rho_c}[U_1((s^*(a_c), s^*(a_c)); \tau)] > \mathop{\mathbb{E}}_{a_i, a_c, \tau; \rho = 0}[U_1((s^*(a_i), s^*(a_c)); \tau)],$$

which is equivalent to the condition that both firms using correlated models is in equilibrium; this strict inequality implies a strict equilibrium. Symmetric argument applies for Firm 2. □

C.5 Proof of Theorem 5.3

Proof. First, both firms using independent algorithms is always a PNE when $a_i > a_c$ as per Theorem 5.1.

We will next state what is needed to prove the theorem. When firms have a preference for correlation at $a_i = a_c$, both firms using correlated algorithms should be a PNE when $a_i = a_c + \epsilon$, for small enough ϵ:

$$\exists\ \epsilon > 0 \text{ s.t. } E^1_{\rho_c, s^*}(a_c, a_c) > E^1_{\rho_0, s^*}(a_i + \epsilon, a_c) \tag{6}$$

On top of that, firms also prefer correlated algorithms over independence at $a_i = a_c + \epsilon$, for small enough ϵ:

$$\exists\, \epsilon > 0 \text{ s.t. } E^1_{\rho_c, s^*}(a_c, a_c) > E^1_{\rho_0, s^*}(a_i + \epsilon, a_i + \epsilon). \tag{7}$$

The proofs for (6) and (7) come from Corollary 5.2, which states that when firms strictly prefer correlation at $a_i = a_c$, correlating is δ-strictly a PNE:

$$\exists\, \delta > 0 \text{ s.t. } E^1_{\rho_c, s^*}(a_c, a_c) \geq E^1_{\rho_0, s^*}(a_i, a_c) + \delta,$$

We define the following shorthand:

A	$E^1_{\rho_c, s^*}(a_c, a_c)$
B	$E^1_{\rho_0, s^*}(a_i, a_c)$
$C(\epsilon)$	$E^1_{\rho_0, s^*}(a_i + \epsilon, a_c)$
$D(\epsilon)$	$E^1_{\rho_0, s^*}(a_i + \epsilon, a_i + \epsilon)$

Put another way, Corollary 5.2 states that

$$\exists\, \delta > 0 \text{ s.t. } A \geq B + \delta \tag{8a}$$
$$A \geq C(\epsilon) + \delta \tag{8b}$$
$$A \geq D(\epsilon) + \delta \tag{8c}$$

at $\epsilon = 0$ and $a_i = a_c$ because $B = C(\epsilon = 0) = D(\epsilon = 0)$. Since $C(\epsilon)$ and $D(\epsilon)$ are continuous in ϵ, (6) and (7) are true by (8b) and (8c).

□

D Experiments (continued)

D.1 Additional Details for Model Multiplicity Setup

We chose the following model hyperparameters to simulate a higher performance for random forests compared to logistic regression:

Model	Hyperparameters
Logistic Regression	$\ell 1$-penalty, saga solver
Random Forest	· # trees = 9 · min # samples in each leaf = 7 · weight: 1.2x for negative class

All unspecified hyperparameters use the default values set by scikit-learn.

D.2 Additional Experiments: Data Procurement

D.2.1 Setup

Our experiment involves two firms who may independently choose to correlate with each others' models by using overlapping datasets. Firm 1 trains on Census data in Texas while Firm 2 trains on Census data in Florida. They both have the option to purchase supplementary data **of worse quality** from a third-party, which in this case is Census data from California whose labels have been perturbed 25% of the time. In doing so, we are giving firms the choice of correlating their models at the expense of predictive accuracy.

In order to smoothly interpolate between independence and correlation, we define a parameter γ; for instance, Firm 1 can use the training data $(1-\gamma)$ TX $+\gamma$ CA, and similarly for Firm 2. If both firms use $\gamma = 0$, there is no overlap in training data and their resulting models will be the least correlated.

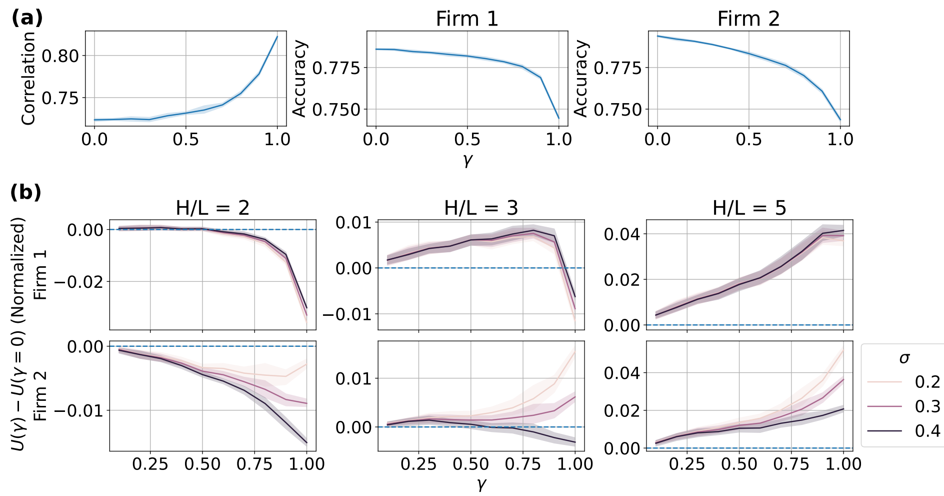

Figure 5: **(a)** [Left] Correlation between both firms' models in the empirical study across various values of γ. $\gamma = 0$ (1) corresponds to no overlap (full overlap) in training data. [Middle and Right] Accuracy of Firm 1 and 2's models over various values of γ. Error bars are 95% confidence intervals over 15 seeds. **(b)** Difference in utility between γ at the x-axis and $\gamma = 0$ (no overlap in training data) for the empirical study, over various values of H/L and σ. Top and bottom rows correspond to Firm 1 and 2's utilities, respectively. Shaded regions indicate 95% confidence intervals over 15 seeds.

Conversely, when both firms use $\gamma = 1$, their training data is identical and their models will be the most correlated.

We randomly sample $n = 200,000$ datapoints from TX, FL, and CA in order to standardize the effect of γ. We then further sample γ proportion of each dataset to ensure that all training data used have exactly n observations. We run this experiment over 15 random seeds, and over $\gamma \in [0, 1]$ in 0.1 increments. Both firms train the same model class (random forests) and have the same test data: Census data from Illinois.

D.2.2 Results: Second-Stage Game

Figure 5(a) shows the predictive accuracy for both firms and the correlation between both firms as γ varies. As expected, accuracy monotonically decreases and correlation monotonically increases as γ increases since firms use more and more of the same lower-quality training data. We observe a significant decrease in accuracy for both firms when $\gamma = 1$, presumably because both models no longer receive the more predictive signal from their original training data.

Figure 5(b) shows the difference in utility between γ at the x-axis and $\gamma = 0$ (independent datasets). When this difference is above 0 (blue dashed line), firms have a preference for correlation at that γ value. We observe such a preference for correlation when consumers are more price sensitive (lower σ) and when the ratio between the H and L prices is larger, as per Theorem 4.2. Firms prefer correlation even when accuracy marginally decreases (subfigure (a)); this happens particularly when there is a high risk of being undercut, making correlating especially beneficial even at the expense of predictive accuracy. However, firms no longer prefer correlation when the trade-off in accuracy is too high (e.g., Firm 1 in $\gamma = 1$). We note that firms are asymmetric: because their models' accuracies differ at various γ, they do not always prefer correlation in the same way, but the general trends remain.

D.2.3 Results: First-Stage Game

We also model firms' decision to correlate at a particular γ. In particular, Figure 6 shows the best response matrices for both firms in choosing various values of γ, over various model parameters $(H/L, \sigma)$. Cells with a red and blue cross indicate a Pure Nash Equilibrium. In general. higher correlation (γ) is only in equilibrium for higher H/L and lower σ, which reflect the same trends as the second-stage game. For example, When $H/L = 6, \sigma = 0.1$, the sole equilibrium exists at

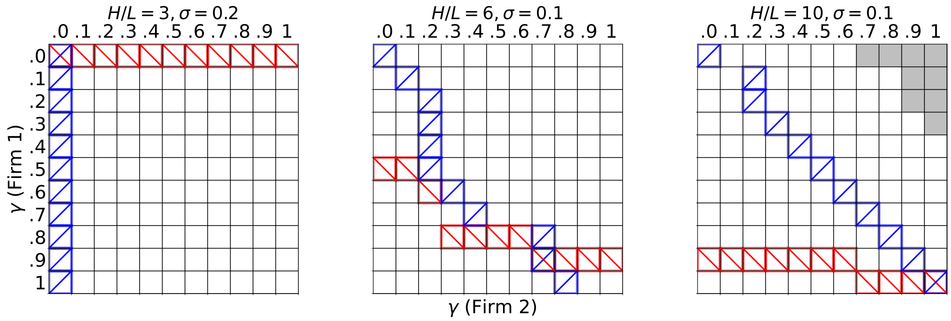

Figure 6: Best response matrices for the two firms in the empirical study, over three select model parameters. $\gamma = 0$ means no overlap in training data (least correlated) while $\gamma = 1$ indicates identical training data (most correlated). Best response for Firms 1 and 2 are highlighted in blue and red, respectively. Nash equilibria exist when both blue and red are highlighted in the same box (e.g., $(0, 0)$ in the left subfigure). Grey boxes are "invalid" regions because following the algorithm would not have been a BNE in the downstream game where firms compete in prices. These results use the average firm utility over X seeds.

$(0.9, 0.7)$. When H/L increases to 10, equilibrium is at $(1, 1)$. We note, however, that in extreme H/L values, certain regions are "invalid" in the sense that firms would not follow the algorithm in the downstream second-stage game (grey cells).

E Additional Results

E.1 Firms choose to correlate, even when algorithms are uninformative

Figure 1 displays regions where firms following the algorithm's recommendation is a BNE for independent models only (light gray) and correlated models only (dark gray). When $a = 0.5$, independent models are never in equlibrium because the algorithms are as good as random. However, when $a = 0.5$ and models are correlated, firms may still choose to follow the algorithm. This region is more likely to be in low σ regimes – where there is the highest risk in being undercut by one's opponent – and therefore there is value in coordinating actions despite the model having no predictive power.

F Miscellaneous

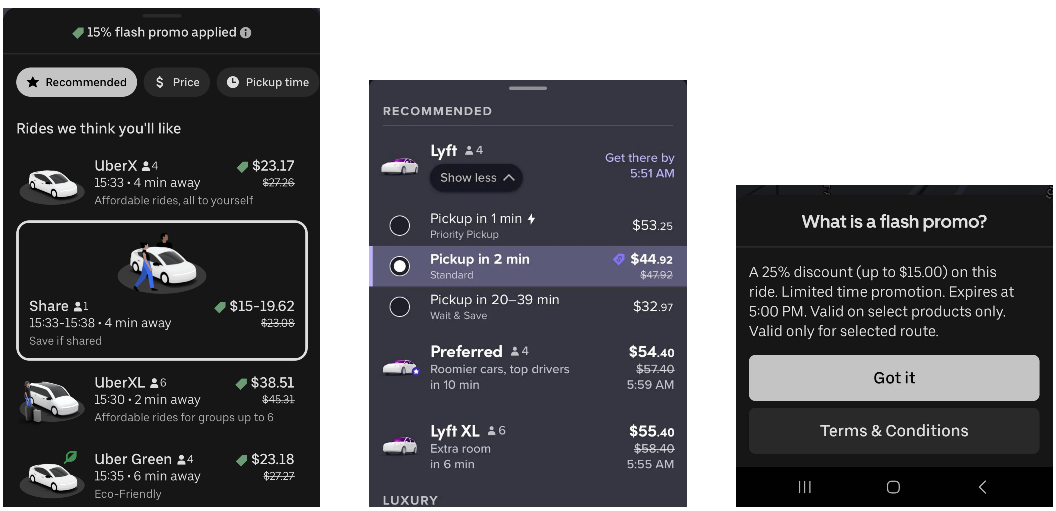

Figure 7: Examples of discounts offer to potential riders on Uber (left) and Lyft (middle). Rightmost figure explains Uber's "flash promo" offering.