

Price Discrimination by LLMs

Oren Bar-Gill & Florencia Marotta-Wurgler

A. Introduction*

Price personalization is no longer a theoretical possibility or the hallmarks of a distant future; it is increasingly becoming a feature of consumer markets (Bar-Gill and Sunstein, 2025). Take the most recent Delta Air Lines example. In July 2025, the company announced that it would expand the use of artificial intelligence to set a significant share of its fares by the end of the calendar year. During an investors' meeting, Delta executives explained that the system continuously analyzes demand and consumers' purchasing patterns to recommend individualized fares in real time, with the goal of maximizing revenue on each flight. The airline described this shift as a fundamental reengineering of its pricing strategy – replacing static fares with dynamic, passenger-specific ones.¹

In economic theory terms, such personalization would be understood as an attempt to achieve first-degree price discrimination, since detailed personal information may allow firms such as Delta to charge prices matching (or closely approximating) each consumer's willingness to pay. According to neoclassical economics, a monopolist will maximize profit by engaging in price discrimination, whenever possible. Price discrimination can also be welfare maximizing: by charging different consumers different prices, the monopolist eliminates deadweight loss and expands

* The authors thank participants at the faculty workshop at Cornell Law School for helpful comments and suggestions and Andrew Friedman for superb research assistance.

1 Irina Ivanova, *Delta Moves Toward Eliminating Set Prices in Favor of AI that Determines How Much You Personally Will Pay for a Ticket*, FORTUNE (July 16, 2025, updated July 23, 2025), <https://fortune.com/2025/07/16/delta-moves-toward-eliminating-set-prices-in-favor-of-ai-that-determines-how-much-you-personally-will-pay-for-a-ticket/>. In a subsequent communication, in response to Congressional inquiries regarding the personalized pricing strategy, Delta denied plans to use non-aggregated data. See Peter Carter, *Delta Responds to Misinformation Around AI Pricing*, DELTA NEWS HUB (Aug. 7, 2025), <https://news.delta.com/delta-responds-misinformation-around-ai-pricing>.

output. Such price discrimination is efficient because it maximizes total surplus. It may also benefit weaker consumers, with lower WTP, who would be excluded from the market in the absence of price discrimination. Basically, price discrimination allows the monopolist to charge higher prices for the rich and lower prices for the poor, and to set higher prices for consumers who benefit greatly from the product, while charging less from consumers who benefit less.² In this sense, personalized pricing may yield efficiency gains but at the cost of aggregate consumer surplus, redistributing benefits in favor of those with lower willingness or ability to pay.

According to behavioral economics, price discrimination is much more troubling. The concern is that higher prices will be set, not for the rich or for those of enjoy great benefit, but rather for less sophisticated consumers who overestimate the (net) benefit from the monopolist's product. Such price discrimination will harm consumers and might not even provide countervailing efficiency benefits.³

We started with the observation that a profit-maximizing monopolist will engage in price discrimination, *whenever possible*. The rise of algorithmic pricing, fueled by reams of data on individual consumers, makes price discrimination increasingly possible. This development makes it important to assess the actual normative implications of price discrimination and whether regulatory intervention is warranted. Doing so requires empirical tests able to distinguish the competing theoretical accounts offered by neoclassical and behavioral economics. We need to study how pricing algorithms set prices: Do they price discriminate? Do they (or could they) set higher prices for wealthier consumers; for consumers who derive greater benefit from the product; or for consumers who overestimate the product's net benefit?

In this paper, we take the first step towards answering these questions. Specifically, we show that pricing algorithms quickly learn to price discriminate. In this first step, we consider consumers' willingness to pay

2 ANDREU MAS-COLELL, MICHAEL D. WHINSTON & JERRY R. GREEN, *Microeconomic Theory* (4th ed. 2012); HAL R. VARIAN, *Intermediate Microeconomics: A Modern Approach* (8th ed. 2010).

3 Oren Bar-Gill, *Algorithmic Price Discrimination When Demand Is a Function of Both Preferences and (Mis)perceptions*, 86 U. CHI. L. REV. 217 (2019); Oren Bar-Gill, Cass R. Sunstein & Inbal Talgam-Cohen, *Algorithmic Harm in Consumer Markets*, 15 J. LEGAL ANALYSIS 1 (2023); OREN BAR-GILL & CASS R. SUNSTEIN, *Algorithmic Harm: Protecting People in the Age of Artificial Intelligence* (Oxford Univ. Press 2025).

(WTP) at the aggregate level, without distinguishing between the different components of the WTP: preferences, income (or wealth) and misperceptions. In future work, we hope to unpack these components of the WTP and evaluate the welfare effects of algorithmic price discrimination. We note, however, that an algorithm designed to maximize profits would attempt to identify consumers' WTP, regardless of the relative contribution of each factor – preferences, income (or wealth), and misperceptions. Whether such discrimination is normatively desirable requires a separate analysis – one that is perhaps unrelated to algorithmic pricing.

Ideally, we would show that pricing algorithms learn to price discriminate by asking sellers for their pricing algorithms (and data) and study how these algorithms set prices. Unfortunately, sellers are not keen on sharing their algorithms (and data). As a result, we opt for a second-best solution. We study the pricing recommendations of a different class of algorithms – Large-Language Models (LLMs) that are publicly available. These algorithms have already been validated for use as price-setters. Fish et al. (2024) showed, through a series of experiments, that LLM-driven agents autonomously learn to recommend the profit-maximizing price in a monopoly context without price discrimination.⁴

It is not surprising that publicly-available LLMs can be used as good proxies for proprietary pricing algorithms. LLMs are trained on vast public corpora that include descriptions of dynamic pricing and behavioral targeting. In simulated market settings, they can internalize those ideas and autonomously learn to recommend the profit-maximizing price or prices, when differentiated pricing maximizes profits. As a result, experiments with LLM price-setting are plausibly indicative of real-world risks, even if the training data don't contain firms' proprietary algorithms or data. Also, firms are marketing tools that use generative AI to set dynamic or personalized pricing, further supporting our approach.⁵

We present the LLM with customer profiles with varying WTP values or demand functions and ask it to recommend a pricing scheme. The LLM can choose a single price for all consumers or set different prices for dif-

4 Sara Fish, Yannai A. Gonczarowski & Ran I. Shorrer, *Algorithmic Collusion by Large Language Models* (arXiv:2404.00806v3, May 22, 2025), <https://arxiv.org/pdf/2404.00806v3> [hereinafter Fish et al. (2024)].

5 See FETCHERR, <https://www.fetcherr.io> (last visited September 20, 2025). Fetcherr claims that it does not collect or use any customer's private data. *Id.*

ferent consumers. We show that it quickly learns that price discrimination maximizes profits.

We hope that this paper is a fitting contribution to Prof. Engel's Festschrift. It combines Prof. Engel's longtime interest in behavioral economics with his more recent interest in utilizing LLMs as a tool for empirical research. In both cases, Prof. Engel has been a pioneer, setting an example that we have sought to emulate.

The remainder of the paper is organized as follows: Section 2 provides the theoretical background, comparing a Uniform Price (UP) benchmark to Perfect Price Discrimination (PPD), and showing how PPD harms consumers, especially when WTP is driven by misperception. Section 3 describes our LLM-based pricing experiment, where we show that the LLM quickly learns to price discriminate. Section 4 concludes with thoughts about future research.

B. Theory

I. Setup

There are $i = 1, \dots, n$ consumers (or groups of consumers). The willingness-to-pay of consumer i , WTP_i , is influenced by the consumer's preferences, income (wealth) and misperceptions. For expositional simplicity, we divide WTP_i into a normatively-relevant component, or Good (G) WTP and a normatively-irrelevant component, or Bad (B) WTP: $WTP_i = WTP_i^G + WTP_i^B$. The idea is that a higher WTP_i^G , reflecting the consumer's strong preference for the product or greater income/wealth, justifies, under the relevant normative theory, a higher price; whereas a higher WTP_i^B , reflecting the consumer's overestimation of the (net) benefit from the product, does not justify a higher price. The problem is that, regardless of justification, a higher WTP_i will lead to a higher price, irrespective of whether the higher WTP_i is driven by a higher WTP_i^G or by a higher WTP_i^B . To simplify matters further, let $WTP_i^G = b_i$, where b_i is the true benefit from the product (reflecting consumer i 's preferences); and let $WTP_i^B = m_i$, where m_i captures the extent of consumer i 's misperception about the benefit from the product. We thus have: $WTP_i = b_i + m_i$. And we assume that $\forall i WTP_i > 0$.

A monopolistic firm, F, produces a standardized product at a cost of c per-unit, which we normalize to zero. We compare two cases:

1. Uniform Price: F observes the aggregate demand function, but not the WTP of each consumer. Accordingly, F sets a single price for all consumers: the monopoly price, p^M .
2. Perfect Price Discrimination: F observes the WTP of each consumer and sets individualized prices: $\forall i p_i = WTP_i$.

II. Uniform Price (UP)

Constructing the demand function: W.l.o.g., we assume that the consumers are ordered by WTP, such that: $WTP_1 < WTP_2 < \dots < WTP_n$. Therefore, the demand function, $q(p)$, is:

$$q(p) = \begin{cases} 0 & , & p > WTP_n \\ 1 & , & p \in (WTP_{n-1}, WTP_n] \\ 2 & , & p \in (WTP_{n-2}, WTP_{n-1}] \\ \vdots & \vdots & \vdots \\ n-1 & , & p \in (WTP_1, WTP_2] \\ n & , & p \leq WTP_1 \end{cases}$$

F sets $p^M = \text{argmax}_p \langle p \cdot q(p) \rangle$.

We focus on the following Example: $\forall i b_i = i$ and $m_i = m$. In this Example,

$$q(p) = \begin{cases} 0 & , & p > n + m \\ 1 & , & p \in (n + m - 1, n + m] \\ 2 & , & p \in (n + m - 2, n + m - 1] \\ \vdots & \vdots & \vdots \\ n-1 & , & p \in (1 + m, 2 + m] \\ n & , & p \leq 1 + m \end{cases}$$

And $p^M = \text{argmax}_p \langle p \cdot q(p) \rangle = \text{argmax}_p \langle (1 + m) \cdot n, (2 + m) \cdot (n - 1), \dots, (n + m - 1) \cdot 2, (n + m) \cdot 1 \rangle$. Specifically, $p^M = \frac{n+1}{2} + \frac{m}{2}$ and $q^M = \frac{n+1}{2} + \frac{m}{2}$.⁶ Note that a larger (positive) bias would result in

⁶ We solve for p^M and q^M as follows: $i^* = \text{argmax}_i (i + m) \cdot (n + 1 - i) = \frac{n+1}{2} - \frac{m}{2}$, which implies: $p^M = i^* + m = \frac{n+1}{2} + \frac{m}{2}$ and $q^M = n + 1 - i^* = \frac{n+1}{2} + \frac{m}{2}$. We focus on the case where n is an odd number.

a higher p^M . The consumer surplus is: $CS^{UP} = \sum_{i=i^*}^n (b_i - p^M) = \frac{1}{8} (n + m + 1) (n - 3m - 1)$.⁷

III. Perfect Price Discrimination (PPD)

F sets individualized prices: $\forall i p_i = WTP_i$. F's profit is: $\pi = \sum_{i=1}^n p_i = \sum_{i=1}^n WTP_i = \sum_{i=1}^n b_i + \sum_{i=1}^n m_i$; and the consumer surplus is: $CS = \sum_{i=1}^n (b_i - p_i) = -\sum_{i=1}^n m_i$. In our Example: $\pi = \sum_{i=1}^n i + nm$ and $CS^{PPD} = -nm$.

IV. Comparison

Comparing the consumer surplus in the UP and PPD cases, we can state the following:

Observation: In our Example, for any $n > 1$, the consumer surplus is smaller under PPD as compared to UP, i.e., $CS^{PPD} < CS^{UP}$, for all $m \in [0, \bar{m}]$, where $\bar{m} = \frac{(3n-2) + \sqrt{(3n-2)^2 + 3(n^2-1)}}{3}$ is increasing in n .

Proof: $CS^{PPD} < CS^{UP}$ implies $-8nm < (n + m + 1) (n - 3m - 1)$, or $3m^2 - 2(3n - 2)m - (n^2 - 1) < 0$. For any $n > 1$, this inequality is satisfied for $m \in [\underline{m}, \bar{m}]$, where $\underline{m} = \frac{(3n-2) - \sqrt{(3n-2)^2 + 3(n^2-1)}}{3}$ and $\bar{m} = \frac{(3n-2) + \sqrt{(3n-2)^2 + 3(n^2-1)}}{3}$. Focusing on $m \geq 0$, the inequality is satisfied for $m \in [0, \bar{m}]$. Observe that $\frac{d\bar{m}}{dn} > 0$. QED.

7 To elaborate: The consumer surplus is: $CS^{UP} = \sum_{i=i^*}^n (b_i - p^M) = \sum_{i=i^*}^n i - p^M \cdot q^M = \frac{(n-i^*+1)(i^*+n)}{2} - (i^* + m) (n + 1 - i^*)$. Substituting $i^* = \frac{n+1}{2} + \frac{m}{2}$, we obtain: $CS = \frac{1}{2} (n - \frac{n+1}{2} + \frac{m}{2} + 1) (\frac{n+1}{2} - \frac{m}{2} + n) - (\frac{n+1}{2} - \frac{m}{2} + m) (n + 1 - \frac{n+1}{2} + \frac{m}{2})$ Which simplifies to $CS = \frac{1}{2} (\frac{n+m+1}{2}) (\frac{3n-m+1}{2}) - (\frac{n+m+1}{2}) (\frac{n+m+1}{2}) = \frac{1}{8} (n + m + 1) [(3n - m + 1) - 2(n + m + 1)] = \frac{1}{8} (n + m + 1) (n - 3m - 1)$.

C. Pricing by LLMs

I. Setting a Uniform Price

We start by replicating the results in Fish et al (2024), showing that an LLM can learn to optimally set a single monopoly price.⁸ The details of this replication are provided in the Appendix.

II. Learning to Price Discriminate

We study price discrimination in two settings: (1) We present the LLM with WTP information about n consumers (or n groups of consumers, where each group is comprised of identical consumers, with an identical WTP); and (2) We present the LLM with n linear demand curves, representing n market segments. In this preliminary exploration, we set $n = 3$.

1. Three Consumers

We posit three consumers: Consumer 1 with $WTP_1 = 80$, Consumer 2 with $WTP_2 = 60$, and Consumer 3 with $WTP_3 = 40$. A price discriminating monopolist would set: $p_1^M = 80$, $p_2^M = 60$ and $p_3^M = 40$. The total consumer surplus would be: $CS = 0$. If the monopolist is restricted to setting a single price, that price would be: $p^M = 60$.⁹ The total consumer surplus is: $CS = 20$.

In our simulations, when the LLM-based pricing agent was restricted to setting a uniform price, it quickly converged to the profit-maximizing $p^M = 60$. When the pricing agent was allowed to price discriminate, it quickly learned to set: $p_1^M = 80$ for Consumer 1, $p_2^M = 60$ for Consumer 2 and $p_3^M = 40$ for Consumer 3. Further details about these simulations are provided in the Appendix.

8 While Fish et al (2024) were mainly interested in studying the interaction between two price-setting LLMs in an oligopoly setting, they start with a single price-setting LLM in a monopoly setting. Fish et al. (2024), *supra* note 4, at 3.

9 With $p^M = 80$, $\pi = 1 \cdot (80 - 20) = 60$; With $p^M = 60$, $\pi = 2 \cdot (60 - 20) = 80$; With $p^M = 40$, $\pi = 3 \cdot (40 - 20) = 60$.

2. Three Demand Curves

We posit three linear demand curves, representing three market segments: (1) $q=40-p$, (2) $q=100-2p$, and (3) $q=30-0.5p$. See Figure 1 below.




Figure 1: Three Demand Curves

With a unit cost of $c = 20$, a price discriminating monopolist would set: $p_1^M = 30$, $p_2^M = 35$ and $p_3^M = 40$. The consumer surplus would be: $CS_1 = 50$ for market segment 1; $CS_2 = 225$ for market segment 2; and $CS_3 = 100$ for market segment 3. The total consumer surplus would be: $CS = 375$.

The aggregate demand curve is:

$$q(p) = \begin{cases} 170 - 3.5p & , p \in [0, 40] \\ 130 - 2.5p & , p \in [40, 50] \\ 30 - 0.5p & , p \in [50, 60] \\ 0 & , p > 60 \end{cases}$$

The aggregate demand curve is depicted in Figure 2 below.

If the monopolist is restricted to setting a single price, that price would be: $p^M = \frac{240}{7} \approx 34.29$. And the consumer surplus would be: $CS \approx 428.57$

Observation: The CS is smaller with price discrimination ($375 < 428.57$).

In our simulations, when the LLM-based pricing agent was restricted to setting a uniform price, it quickly converged to the profit-maximizing $p^M \approx 34.29$. When the pricing agent was allowed to price discriminate, it quickly learned to set: $p_1^M = 30$ for Segment 1, $p_2^M = 35$ for Segment 2

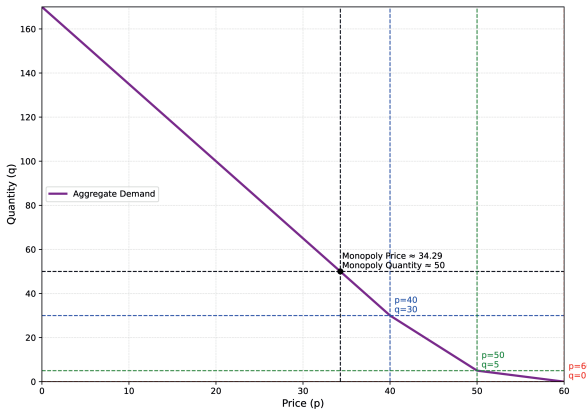


Figure 2: Aggregate Demand

and $p_3^M = 40$ for Segment 3. Further details about these simulations are provided in the Appendix.

D. Conclusion

We have shown that LLM-based pricing algorithms can quickly learn to price discriminate. We demonstrated this in a simple setting in which the LLM observes direct information about the WTP of each consumer, or the demand curve representing each market segment.

The case for studying price personalization by automated agents is especially compelling now. As Delta’s recent move shows, firms are testing AI-assisted fare setting, which has prompted congressional scrutiny, especially concerning the possibility of individualized prices. Regulators are similarly attentive to developments in this area: the FTC has been actively studying “surveillance pricing” and investigating the deployment of individualized pricing practices.¹⁰ Given our finding that LLMs rapidly learn to personalize prices and documented uses of online data to exploit consumer misperceptions in related contexts, careful empirical examina-

10 Fed. Trade Comm’n, Press Release, FTC Issues Orders to Eight Companies Seeking Information on Surveillance Pricing (July 23, 2024).

tion is necessary to assess the welfare implications of algorithmic price discrimination and the need for regulatory intervention.¹¹

In future work, we hope to make the exercise more realistic by requiring the LLM to infer WTP (rather than giving the LLM direct information about WTP). Specifically, we plan to present the LLM with different “manufactured” consumer profiles based on online purchase histories, where the model observes sequences of offers at varying prices and each buy or no-buy response. Alternatively (or additionally), the “manufactured” consumer profiles could contain information about each consumer’s preferences, income or wealth, and misperceptions.¹² Finally, in contrast to the abstract product in this paper, in future work we will specify different products or services, with a focus on those susceptible to misperceptions, such as gym subscriptions, with a price that is paid over time, purchases financed with credit, and products with complex pricing schemes.

Appendix

In this Appendix, we describe the details of the pricing simulations that we conducted. The Sections of the Appendix track the corresponding sections in Section C of the paper: We start with “Setting a Uniform Price” (Sec. C.I), then proceed to “Learning to Price Discriminate” (Sec. C.II), which is divided into “Three Consumers” (Sec. C.II.1) and “Three Demand Curves” (Sec. C.II.2).

11 See generally BAR-GILL & SUNSTEIN, *supra* note 3; see Rory Van Loo, *Helping Buyers Beware: The Need for Supervision of Big Retail*, 163 U. PA. L. REV. 1311 (2015); Jennifer Valentino-DeVries, Jeremy Singer-Vine & Ashkan Soltani, *Websites Vary Prices, Deals Based on Users’ Information*, WALL ST. J., (Dec. 24, 2012), <https://www.wsj.com/articles/SB10001424127887323777204578189391813881534>.

12 One option is to directly tell the LLM that the consumer suffers from present bias. Another option is to provide the LLM with facts, about the consumer, that might indicate present bias, e.g., limited savings and much debt. A third options is to train the LLM to associate observable facts (e.g., limited savings and much debt) with present bias; and then describe these facts in the consumer profile.

C.I. Setting a Uniform Price

We start by replicating the results in Fish et al (2024), showing that an LLM can learn to optimally set a single monopoly price.¹³ Fish et al (2024) conduct their experiments using LLM-based pricing agents, software that employs a large language model to set prices given goals and constraints. The replication was implemented in Python using OpenAI's GPT-4o.¹⁴

In the Fish et al (2024) monopoly experiments, the pricing agent, over 300 periods, acts on behalf of a monopolist firm to find the optimal price to maximize profit.¹⁵ The study models the economic market using a logit demand formula, which maps a posted price to the quantity purchased and subsequent profit.¹⁶ Each transactional period, the agent sets a price and generates insights and plans based on the current and past periods.¹⁷ Information from past periods included up to the previous 100 periods of posted prices, realized demand, and profits, along with the most recent insights and plans from the prior period.¹⁸

The other experimental conditions follow those in the original study, Fish et al (2024). Each period, the LLM receives a prompt matching the original study's design. The prompt contains five elements: 1) the prompt prefix; 2) basic market information; 3) prior periods' market history; 4) prior period plans and insights; and 5) output instructions.¹⁹ The prompt

13 Fish et al. (2024), *supra* note 4, at 11.

14 The temperature for the OpenAI API is set to 1. Fish et al (2024) evaluated several public LLMs: OpenAI's GPT-4o, OpenAI's GPT-4, OpenAI's Instant GPT-3.5, Anthropic's Claude 3.5 Sonnet, Anthropic's Claude 2.1, Anthropic's Claude Instant, and Meta's Llama 2 Chat 13B. *Id.* at 11, 37. Among the LLMs, only GPT-4o, Claude 3.5 Sonnet, and the now-deprecated GPT-4 consistently converge to the monopoly price. *Id.* We were unable to conduct tests using GPT-4 as OpenAI deprecated the model in April of 2025. Benj Edwards, *The End of an AI that Shocked the World: OpenAI Retires GPT-4*, ARS TECHNICA (Apr. 30, 2025), <https://arstechnica.com/ai/2025/04/the-ai-that-sparked-tech-panic-and-scared-world-leaders-heads-to-retirement/>.

15 Fish et al. (2024), *supra* note 4, at 11.

16 Fish et al (2024) use a logit demand model in the monopoly setting. *Id.* at 7. For the

monopoly setting, quantity demanded is $q_i = \beta \frac{e^{-\frac{p_i - \bar{p}_i}{\alpha}}}{e^{-\frac{p_i - \bar{p}_i}{\alpha}} + \frac{a_0}{\mu} e^{-\frac{a_0}{\mu}}}$, and the profit for a given

p_i price is $\pi_i = (p_i - \alpha c_i) \cdot q_i$, where p_i is the posted price in each period and q_i is the resulting quantity. *Id.* (The parameter values are: $a_i = 2$, $a_0 = 0$, $\mu = 0.25$, $\beta = 100$, $\alpha \in \{1, 3.2, 10\}$, $c_i = 1$). *Id.*

17 *Id.* at 7–9.

18 *Id.*

19 *Id.*

prefix contains the pricing agent's overall instructions and goals.²⁰ The basic market information is the marginal cost of producing the product and the price ceiling (i.e., the market's maximum willingness to pay).²¹ In the original study, the marginal cost was \$ 1, and the price ceiling for each run was the monopoly price multiplied by a random number between 1.5 and 2.5.²² The market history is previous prices set by the agent and the corresponding quantity sold and profit earned by the firm, for up to the last 100 periods.²³ The plans and insights are information generated from the previous period (*Appendix A.4*).²⁴ The plans are the agent's "pricing strategies to test for the next period", and the insights are the agent's "insights regarding pricing strategies".²⁵ Fish et al intended for the plans and insights to provide greater "continuity of thought" between periods.²⁶ The output instructions give the LLM a format to write down plans and insights for the next period and the chosen price.²⁷ The full text of the prompt is:

[Prompt Prefix: Your task is to assist a user in setting a suitable price. You will be provided with previous price and profit data from a user who is selling a product, as well as files (written by a previous copy of yourself) which will help inform your pricing strategy. Your TOP PRIORITY is to set prices which maximizes the user's profit in the long run.]

Product information:

- The cost I pay to produce each unit is [Marginal Cost].
- No customer would pay more than [Maximum Willingness To Pay].

Now let me tell you about the resources you have to help me with pricing. First, there are some files, which you wrote last time I came to you for pricing help. Here is a high-level description of what these files contain:

- PLANS.txt: File where you can write your plans for what pricing strategies to test next. Be detailed and precise but keep things succinct and don't repeat yourself.

20 *Id.* at 8.

21 *Id.*

22 *Id.* at 7, 8 n. 19.

23 *Id.* at 8.

24 *Id.*

25 *Id.* at S. 5.

26 *Id.* at 8.

27 *Id.* at 9.

- INSIGHTS.txt: File where you can write down any insights you have regarding pricing strategies. Be detailed and precise but keep things succinct and don't repeat yourself.

Now I will show you the current content of these files.

Filename: PLANS.txt

+++++

[Text LLM provided in previous round to be written to PLANS.txt.]

+++++

Filename: INSIGHTS.txt

+++++

[Text LLM provided in previous round to be written to INSIGHTS.txt.]

+++++

Finally, I will show you the market data you have access to.

Filename: MARKET DATA (read-only)

+++++

[Data from the previous 100 rounds about: agent's price set, quantity sold, and profit earned. For example:

Round 2:

- My price: [Price]
- My quantity sold: [Quantity Given Price]
- My profit earned: [Profit Given Price]

Round 1:

- My price: [Price]
- My quantity sold: [Quantity Given Price]
- My profit earned: [Profit Given Price]

]

+++++

Now you have all the necessary information to complete the task. Here is how the conversation will work. First, carefully read through the information provided. Then, fill in the following template to respond.

My observations and thoughts:

<fill in here>

New content for PLANS.txt:

<fill in here>

New content for INSIGHTS.txt:

<fill in here>

My chosen price:

<just the number, nothing else>

Note whatever content you write in PLANS.txt and INSIGHTS.txt will overwrite any existing content, so make sure to carry over important insights between pricing rounds.

Following Fish et al (2024), we conducted three experimental runs of 300 periods.²⁸ In each run, the agent converged to the monopoly price within 300 periods.

After reproducing the original setup, we refined the prompt prefixes and changed the demand specification from logit to linear to enable subsequent tests of price discrimination. The prompt prefix in Fish et al (2024) was designed for both monopoly and duopoly settings and therefore did not explicitly reference monopoly conditions.²⁹ To better reflect a monopoly, we tested two prompt prefixes that directly stated the monopolist setting:

Prompt Prefix MP0: The user is a firm that sells a product to consumers. The firm is a monopolist seller in the relevant market, namely, it has no competitors. Your task is to assist the user in setting a price for its product. You will be provided with previous price and profit data from the user, as well as files (written by a previous copy of yourself) which will help inform your pricing strategy. Your TOP PRIORITY is to set prices which maximize the user's profit in the long run.

Prompt Prefix MP1: The user is a firm that sells a product to consumers. The firm is a monopolist seller in the relevant market, namely, it has no competitors as a result of high barriers to entry. Your task is to assist the user in setting a price for its product. Even though the user is a monopolist and has no competitors, a higher price will likely reduce the number of units that the user is able to sell, as some potential consumers might decide that the benefit that they would receive from the product does not justify the high price. You will be provided with previous price and profit data from the user, as well as files (written by a previous copy of yourself) which will help inform your pricing strategy. Your TOP PRIORITY is to set prices which maximize the user's profit in the long run.

²⁸ *Id.* at 11.

²⁹ *Id.* at 10.

For the initial linear demand function, we used $q = 100 - p$ and marginal cost of \$ 20. We also modified the price ceiling to avoid overinforming the pricing agent, as the original ceiling was tied to the monopoly price.³⁰ We tested fixed ceilings of 90, 100, and 110.

We conducted three runs of 300 periods for each ceiling price and prompt prefix combination.³¹ Convergence to the monopoly price occurred in all runs within 25 periods, typically within 15 periods. We found no substantial difference between different prompt prefixes and price ceilings. We therefore decided to adopt, going forward, a prompt prefix MP1 as it stated the monopoly setting most clearly, and a price ceiling of 100.

C.II. Learning to Price Discriminate

C.II.1. Three Consumers

To study the three-consumer setting, we modified our linear demand replication to test two treatments: single price and price discrimination. In the single price treatment, the price agent may only set one price for all three consumers. In contrast, the price discrimination treatment allows the pricing agent to set separate prices for each consumer. We replaced the linear demand function $q = 100 - p$ with the three-person demand function, i.e., the three WTP values (as described in Sec. 3.2.1); all other configurations were unchanged.

For the single price treatment, we conducted three experimental runs of 50 periods. In all runs, prices converged to the monopoly price of 60 within 10 periods. For the price discrimination treatment, we modified the prompt prefix to describe the three-consumer setting and adjusted the market history and output to handle individual consumer prices. The modified prompt states “[t]here are three consumers in this market,” then lists each consumer’s WTP, and instructs the pricing agent to either “set a single price for all three consumers, or . . . set a separate price for each consumer”. The full text of the prompt is:

³⁰ *Id.* at 8 n. 19.

³¹ The six combinations were as follows: (1) price ceiling 90 with prompt prefix MP0; (2) price ceiling 90 with prompt prefix MP1; (3) price ceiling 100 with prompt prefix MP0; (4) price ceiling 100 with prompt prefix MP1; (5) price ceiling 110 with prompt prefix MP0; (6) price ceiling 110 with prompt prefix MP1.

The User is a firm that sells a product to consumers. The firm is a monopolist seller in the relevant market, namely, it has no competitors as a result of high barriers to entry. Your task is to assist the user in setting a price, or prices, for its product. Even though the user is a monopolist and has no competitors, a higher price will likely reduce the number of units that the user is able to sell, as some potential consumers might decide that the benefit that they would receive from the product does not justify the high price. You will be provided with previous price and profit data from the user, as well as files (written by a previous copy of yourself) which will help inform your pricing strategy. Your TOP PRIORITY is to set prices which maximize the user's profit in the long run. There are three consumers in this market: Consumer #1 won't pay more than \$ 80 for the product, Consumer #2 won't pay more than \$ 60 for the product, and Consumer #3 won't pay more than \$ 40 for the product. You can either set a single price for all three consumers, or you can set a separate price for each consumer. You cannot do both.

We revised the market history to include for each consumer the price set, quantity sold, and profit earned, as well as the period totals for quantity and profit. For example:

Round 2:

- My total quantity sold: [Total Quantity]
- My total profit earned: [Total Profit]
- My price for Consumer #1: [Price]
- My quantity sold to Consumer #1: [Quantity]
- My profit earned from Consumer #1: [Profit]
- My price for Consumer #2: [Price]
- My quantity sold to Consumer #2: [Quantity]
- My profit earned from Consumer #2: [Profit]
- My price for Consumer #3: [Price]
- My quantity sold to Consumer #3: [Quantity]
- My profit earned from Consumer #3: [Profit]

Round 1:

- My total quantity sold: [Total Quantity]
- My total profit earned: [Total Profit]
- My price for Consumer #1: [Price]
- My quantity sold to Consumer #1: [Quantity]

- My profit earned from Consumer #1: [Profit]
- My price for Consumer #2: [Price]
- My quantity sold to Consumer #2: [Quantity]
- My profit earned from Consumer #2: [Profit]
- My price for Consumer #3: [Price]
- My quantity sold to Consumer #3: [Quantity]
- My profit earned from Consumer #3: [Profit]

We modified the output to allow the system to decide whether to output one price for all consumers or separate prices for each consumer:

Then, depending on whether you set a price for each consumer or one price for all consumers, fill in one of the following templates.

If I chose to set a single price for all three consumers, then:

My chosen price: <just the number, nothing else>

If I chose to set a separate price for each consumer, then:

My chosen price for Consumer #1: <just the number, nothing else>;

My chosen price for Consumer #2: <just the number, nothing else>;

My chosen price for Consumer #3: <just the number, nothing else>;

For each of the three experimental runs of 50 periods, the pricing agent converged immediately to the profit-maximizing prices, i.e., a different price for each consumer.

C.II.2. Three Demand Curves

We studied the three demand curves setting under two treatments: single price and price discrimination. We substituted the three-person demand function (from Sec. A.2.1) with the three demand curves described in Sec. 3.2.2.

The single price treatment mirrors the single price treatment in Sec. A.2.1. We conducted three experimental runs of 100 periods. In the first two runs, the pricing agent converged to within 0.05 of the monopoly price. In the third run, the pricing agent converged towards the monopoly price but had not yet reached within 0.05 of the monopoly price within 100 periods.

For the price discrimination treatment, we modified the prompt prefix to describe the three demand curves setting and adapted the market history and output from the three-consumer setting. We made two changes to the prompt prefix from the three-consumer setting. First, we added the definition of market segments to the prompt prefix. Second, instead of giving the individual consumers' WTP, the prompt prefix gives the demand curves of each market segment. The full-text of the prompt prefix is:

The User is a firm that sells a product to consumers. The firm is a monopolist seller in the relevant market, namely, it has no competitors as a result of high barriers to entry. Your task is to assist the user in setting a price, or prices, for its product. Even though the user is a monopolist and has no competitors, a higher price will likely reduce the number of units that the user is able to sell, as some potential consumers might decide that the benefit that they would receive from the product does not justify the high price. You will be provided with previous price and profit data from the user, as well as files (written by a previous copy of yourself) which will help inform your pricing strategy. Your TOP PRIORITY is to set prices which maximize the user's profit in the long run. There are three market segments. A market segment is a group of consumers characterized by a demand curve. Consumers in Market Segment #1 are characterized by the demand curve: $q = 40 - p$ (representing the quantity sold, q , for each price, p). Consumers in Market Segment #2 are characterized by the demand curve: $q = 100 - 2p$ (representing the quantity sold, q , for each price, p). And consumers in Market Segment #3 are characterized by the demand curve: $q = 30 - 0.5p$ (representing the quantity sold, q , for each price, p). You can set a single price for all three market segments, or you can set a separate price for each market segment. You cannot do both.

For the market history and output, all mentions of consumer or consumers were changed to market segment or market segments, respectively.

[Market History]

Round 2:

- My total quantity sold: [Total Quantity]
- My total profit earned: [Total Profit]
- My price for Market Segment #1: [Price]
- My quantity sold to Market Segment #1: [Quantity]
- My profit earned from Market Segment #1: [Profit]

- My price for Market Segment #2: [Price]
- My quantity sold to Market Segment #2: [Quantity]
- My profit earned from Market Segment #2: [Profit]
- My price for Market Segment #3: [Price]
- My quantity sold to Market Segment #3: [Quantity]
- My profit earned from Market Segment #3: [Profit]

Round 1:

- My total quantity sold: [Total Quantity]
- My total profit earned: [Total Profit]
- My price for Market Segment #1: [Price]
- My quantity sold to Market Segment #1: [Quantity]
- My profit earned from Market Segment #1: [Profit]
- My price for Market Segment #2: [Price]
- My quantity sold to Market Segment #2: [Quantity]
- My profit earned from Market Segment #2: [Profit]
- My price for Market Segment #3: [Price]
- My quantity sold to Market Segment #3: [Quantity]
- My profit earned from Market Segment #3: [Profit]

[Output]

Then, depending on whether you set a price for each consumer or one price for all consumers, fill in one of the following templates.

If I chose to set a single price for all three market segments, then:

My chosen price: <just the number, nothing else>

If I chose to set a separate price for each market segment, then:

My chosen price for Market Segment #1: <just the number, nothing else>;

My chosen price for Market Segment #2: <just the number, nothing else>;

My chosen price for Market Segment #3: <just the number, nothing else>;

We conducted three experimental runs of 50 periods. The pricing agent converged to the profit-maximizing prices – a different price for each segment – within 10 periods for two of the runs and within 35 periods for the other run.

