Analysis Plan for Approximation in Complex Pricing Mechanisms

Jeffrey D. Michler, José G. Nuño-Ledesma, Peter Slade, and Steven Y. Wu

September 12, 2022

1 Introduction

Firms that produce multiple products have a wide variety of options in setting prices for those goods. Examples include digital media platforms like iTunes and Netflix, telecom companies like Rogers and Bell, restaurant chains, sports teams, museums, and theaters. The firms can sell their products at a single uniform price or set prices for each individual good. They can also sell their goods only as a complete bundle or bundle different sets of goods in a variety of ways to target specific types of consumers. While a wide variety of complex pricing mechanisms exist for firms to choose from, in reality firms implement very simple pricing strategies. A primary reason for this is the staggering complexity of the pricing problem as the number of products increases, which tests the bounds of rationality (Rubenstein, 1997).

Recent advances in economic theory have shown that simpler pricing strategies can closely approximate the more complex profit maximizing pricing mechanisms. But these simpler strategies can still be complicated and are often highly sensitive to the firm's knowledge of consumer demand for each good. This has restricted empirical research on the topic to case studies and limited the generalizability of these theories as useable strategies for firms. The question remains, how well do simpler pricing strategies approximate the profit maximizing mechanism?

The overall objective of this research is to build knowledge and understanding regarding the effectiveness of simpler pricing strategies in approximating more profitable, but more complex, pricing mechanisms. To accomplish this goal, we will use an economic lab experiment to answer the following research questions:

1. How robust is the profit maximizing pricing mechanism to gaps in a firm's knowledge regarding a) the demand for the goods and b) the correlation between consumers' value for the goods?
2. Among existing popular pricing strategies, which is the most robust to the absence of information regarding a and b?
3. What is the impact of different pricing strategies on both firm profits and social welfare?

Answers to these questions represent a significant advance in the application of what has till now been primarily a theoretical literature in economics. The experiment will generate concrete and specific advice to multi-product firms, allowing them to better understand the merits of feasible pricing strategies. Finally, by measuring the impact of different pricing strategies on social welfare, our research provides guidance to governments looking to regulate monopolistic behavior and curtail predatory pricing by firms.

2 Theoretical Framework

In this section we introduce the environment underlying the design of our experiment. The typical theoretical set-up is to assume a market with a monopolist firm which sells K distinct goods, and a mass of consumers N with independently-distributed reservation values for each good on sale: $(\rho_{i1}, ..., \rho_{iK})$. For simplicity, marginal costs are set equal to zero. Each consumer demands at most one copy of each good k. It is assumed that the firm knows the distribution of values for each good for every consumer, and chooses a pricing function $P(q_1, ..., q_K)$ which encompasses the choice of pricing strategy PS as well as the resulting J prices the firm must set for the chosen strategy.

In practice, we restrict this to four possible strategies:

1. **Component pricing**: the firm sets $(p_1^{cp}, ..., p_K^{cp})$ prices, one price for each good.
2. **Pure bundling**: the firm sets one price, p_K^{pb}, which requires the consumer to purchase all of the goods on sale.
3. **Mixed bundling**: the firm sets $(p_1^{mb}, ..., p_{2^K-1}^{mb})$ prices, one price for every possible combination of the k goods on sale.
4. **Bundle-size pricing**: the firm sets $(p_1^{bsp}, ..., p_K^{bsp})$ prices, one for each total quantity (or bundle size) the consumer buys, regardless of which specific goods are in that bundle.

When the firm offers component pricing, consumers simply make K distinct decisions: purchase good k if $\rho_{ik} \geq p_k^{cp}$. In the case of pure bundling, consumers make one decision: purchase the bundle if $\sum \rho_{ik} \geq p_K^{pb}$. In the case of mixed bundling, consumers search over the entire set of prices to choose the bundle B to maximize $\sum_{b \in B} \rho_{ib} - p_B^{mb}$. And finally for bundle-size pricing, consumers first identify the optimal goods to include in each size bundle, and then choose $|B|$ to maximize $\sum_{b \in B} \rho_{ib} - p_{|B|}^{bsp}$. Given how consumers optimize, and conditional on choosing a pricing strategy, firms simply maximize profits by maximizing sales.

This standard set-up of the firm and consumer problem relies on two key assumptions. First is that the monopolist has perfect information regarding the distribution of consumer preferences. Second is that both the firm and the consumers do not suffer from bounded rationality. The firm may posses perfect information but it may make mistakes in solving the pricing problem. This is not a trivial assumption since closed-form solutions do not exist for mixed bundling beyond the two good case. Further complicating the firm's problem is that consumers may be boundedly rational, making mistakes when faced with a long menu of choices. Thus, even if the firm has perfect information and can solve the pricing problem, consumers may make errors in solving their utility maximization problem, resulting in a realized demand that differs from true demand. In this experiment, we explore the assumptions of perfect information and unbounded rationality in the firm. Experiments run by Deutschmann et al. (2018) explore the implications of bounded rationality in the consumer.

3 Experimental Design

We design an experimental environment that mirrors the framework in Section 2 while relaxing two of they key assumptions regarding firms. In brief, a player in the lab experiment, denoted by i, will act as a firm trying to decide the prices to set for the various goods that they sell. The marginal cost for all goods is zero and disposal of unsold goods is costless. The "consumers" that the player is trying to sell to are a set of 10,000 computer buyers or bots that are pre-programmed to purchase goods or bundles of goods if the price is less than or equal to that bot's preset valuation. Firms are monopolistic, so each player is the sole seller in her market and does not compete against other players. We constrain the prices a player can set and the reservation values for the bots to be integers. The goal for the player is to set prices in order to maximize revenue (profit). Play takes place over a number of rounds, denoted by t. Payouts will be calculated as a fixed percentage of average revenue earned across three randomly selected rounds.

There are four parameters which we will experimentally manipulate in order to induce random variation in the markets that a given play i will face.

- The first parameter is the **pricing strategy**, denoted by b. Through the course of the game, all players will play each of the four pricing strategies and be asked to make decisions about how to price their goods using that strategy. The game will take place in 4 sub-games with 10 rounds in each sub-game. In a sub-game the player is assigned, at random, one of the four pricing strategies. They then use that strategy for all 10 rounds in the sub-game. After the sub-game, the player is then assigned, at random, to a second of the pricing strategies. This continues until the player has played with all 4 pricing strategies. During a set of practice rounds, each player will be provided with instructions and information about the characteristics of all 4 pricing schemes.

- The second parameter is the **number of goods** that the firm sells, denoted by n. In the experiment, there will be four possible goods that a firm could sell, denoted by color (blue, green, yellow, and red). Before the start of the game each player will be told the number of goods they are selling (2 through 4). For a given number of goods, all players selling that number of goods will be selling the same color goods.

- The third parameter is the **distribution** of the consumer valuations for each good, denoted by d. At the start of the game, a player will be randomly assigned to a market in which consumer valuations are either uniformly distributed or follow a beta distribution. Whether the player knows the distribution or not is one of our treatments (discussed below). For the uniform distribution, valuations will follow $U \sim (0, 100)$, which has a mean of 50 and a standard deviation of 29. For the beta distribution, valuations will follow $\beta \sim (5, 5)$. Scaling the beta distribution by 100 gives the distribution a mean of 50 and a standard deviation of 15. Most theoretical models of non-linear pricing assume consumer valuations follow either a uniform or Gaussian (normal) distribution. Since the support of a normal distribution is the entire real number line, we approximate the normal with a beta distribution, whose support is $x \in [0, 1]$. When scaled by 100, both the support and the mean of the uniform and beta distributions are equal, though the standard deviations still differ.

- The fourth parameter is the **correlation** between a given consumer's valuation for the goods, denoted by r. At the start of the game, a player will be randomly assigned to a market in which every consumer has the same correlation in their valuation for the goods on sale. Consumers will express either a negative correlation in their valuation for the goods or their valuations will be perfectly independent.[1] Regardless of the distribution, the value for all goods will be drawn from distributions that have a correlation coefficient of -0.25. As with the distribution of consumer valuations, a player's knowledge of the correlation in consumer valuations will be one of our treatments (discussed below).

With these four parameters, we can define four different markets which an individual player will be randomly assigned into.

1. Uniform, independent
2. Uniform, negative
3. Beta, independent
4. Beta, negative

We can also define three different production schedules which an individual player will be randomly assigned into.

1. Two goods (CP, PB, MB, BSP)
2. Three goods (CP, PB, MB, BSP)
3. Four goods (CP, PB, MB, BSP)

This gives us 12 market-schedule combinations.

Subsection 3.1 provides details about the experiment interface and game play, subsection 3.2 describes the treatments, while subsection 3.3 discusses our approach for pre-registration.

[1] We do not test positive correlation, since theory does not exist which proves which pricing scheme is profit maximizing when goods are positively correlated.

3.1 Lab Experiment

We conduct our experiment in the Economic Science Laboratory (ESL) at the University of Arizona. Participants will be drawn from the undergraduate population the university. Given the size of the lab, we expect to have between 20 and 30 students in each session. Since students do not interact with each other but rather interact with a completely computerized market, we do not need to worry about having a certain number of students in each session. In each session, participants will show up and receive verbal and written instructions regarding the experiment. Participants will then take a quick quiz to measure comprehension of the instructions. They will also play four practice rounds in which they are walked through the specifics of each pricing strategy, asked to set prices, and shown what information will be provided at the end of each round.

Each participant is randomly assigned to one of the 12 market-schedule combinations. Participants complete 40 rounds of the game. At their computer terminals, participants will be told:

- Which pricing strategy they have been assigned to first (b),
- The number of goods they will sell (n),
- The number of buyers in the market ($10,000$),
- And the min and max valuation that buyers have for the goods ($\min = 0,\ \max = 100$).

The game then begins with each participant being given the opportunity to select what prices to set for their first randomly assigned schedule (sub-game). Rounds have no maximum time limit. After participants are satisfied with their selected prices, they will be asked to confirm their choice and then they will be shown a screen displaying a slide bar which corresponds to four different emojis that convey an emotional response. The emoji emotional response screen is displayed after every round between submitting prices and revealing performance. Subsection **??** provides more information on measuring emotional response.

After the emoji screen, participants are shown a results screen that displays the prices they chose, the resulting quantity sold of each good, the profit they earned, both per good and total, and what that translates to in real dollars. In subsequent rounds the information screen will also display the total profit and dollar equivalent from all previous rounds in the sub-game. Participants will then move on to the next round where they will be asked to use their previous experience in the "market" to revise their pricing decisions. At the end of 4 sub-games (40 rounds), participants will be shown the profits earned in each round and their payout based on three randomly selected rounds.

3.2 Treatments

Control participants will be provided with perfect information regarding both the distribution of consumer valuations and the correlation between values for each goods. This information will be communicated to participants in two ways. They will be told the parameters of the distribution along with its mean and standard deviation. They will also be shown a graph of the distributions probability density function. The information regarding the correlation will also be communicated to participants in two ways. Participants will be told whether the correlation is negative or independent along with the actual corresponding correlation coefficients. They will also be shown a scatter plot graph of the distribution of valuations.

In terms of treatments, participants will be randomly assigned to one of three treatments.

Perfect info	Unknown d
Unknown r	Unknown d & r

Unknown distribution (d) **treatment:** Participants randomly assigned to this treatment will be provided with no information regarding the underlying distribution of consumer valuations.

Unknown correlation (r) **treatment:** Participants randomly assigned to this treatment will be provided with no information regarding the correlation between consumers' valuation for the goods.

Unknown distribution (d) **and correlation** (r) **treatment:** Participants randomly assigned to this treatment will be provided with no information regarding both the distribution and the correlation.

3.3 Hybrid Pre-Registration Approach

Pre-analysis plans have become increasingly common in economics, particularly in experimental work. However, for work in which researchers have a large number of hypotheses on which they place a low weight ex ante, a strict pre-analysis plan may limit the ability to learn from data and exploratory analysis. Thus, in our work, we have elected to follow the strategy suggested in Anderson and Magruder (2017), which we feel allows us to increase credibility and limit the potential for false discovery, while also recognizing how we might learn from the data during analysis.

More specifically, we first submit a set of primary hypotheses, described below in Section 4, to the AEA registry before data collection begins. Second, once data collection is complete, we recruit a third party to split our sample before beginning any analysis work. The third party sets a seed without our knowledge and randomly generates a 35% sample of our data, stratified by treatment group and participants' experience with lab experiments. It is on this exploratory sample that we conduct all secondary analysis, exploring possible hypotheses and regression specifications. Third, having developed a satisfactory set of secondary hypotheses, we will update the pre-registration document with these new hypotheses. Finally, upon verification of that update, the third party will release the remainder of our data which will be used for completing the final confirmatory analysis.

4 Empirical Strategy: Pre-registered Hypotheses

The goal of the experiment is to determine if the simplicity of CP, PB, or BSP, though sub-optimal in terms of profits, allow the strategies to outperform MB in the absence of perfect information. We will measure performance of each mechanism according to six criteria, which are our primary outcomes of interest.

4.1 Robustness of Pricing Strategies Given Perfect Information

We judge the performance of each mechanism by six criteria. A robust mechanism is one that not only is a short distance from the maximum achievable profits but is adaptable (profitable across environments), performs best in worst case scenarios (maximizes the minimum profits achieved), is easy to learn (quick convergence of prices across rounds), and reduces variability in profit (performs well across all individuals).

All of our pre-registered hypotheses relate only to the control group that has perfect information. The following regressions are will use data from the control group only.

4.1.1 Profitability

Profitability is the simplest criteria on which to judge the performance of each pricing strategy. The outcome simply asks which strategy, b, actually delivers the most money for the participant i. We measure profitability as:

$$Y_{bi}^{\Pi} = \sum_{t=1}^{10} \pi_{bit}^{att}, \tag{1}$$

where π_{bit}^{att} is the profit attained by individual i in round t using pricing strategy b. This is summed over all rounds to determine the total amount of profit earned by the individual for each pricing strategy in the experiment.

To determine which pricing strategy generates the most profit, we estimate:

$$Y_{bi}^{\Pi} = \alpha + \beta_1 PB_i + \beta_2 BSP_i + \beta_3 MB_i + \gamma_g + \theta_i + \epsilon_{bi}. \tag{2}$$

Here we control for the ordering in which an individual saw each pricing strategy (θ_i) in order to account for the learning that may occur from playing easier strategies earlier. We also control for the number of goods being sold (γ_g) so that we are comparing someone selling two goods via PB to someone selling two goods via CP. This also allows us to control for the fact that people using BSP to sell four goods will make more money than people using BSP to sell two goods, simply because they are selling more items. We cluster the standard errors by participant.

This regression estimates the average treatment effect of each pricing strategy on the profit attained by individuals in the experiment. It answers the question: on average, which pricing strategy allows individuals selling the same goods to generate the most profits, regardless of market. Based on existing models of non-linear pricing mechanisms, hypothesis 1 is as follows.

Hypothesis 1 *All three complex pricing strategies will increase profits relative to CP* ($\beta_1 > 0;\ \beta_2 > 0;\ \beta_3 > 0$). *Furthermore, BSP will outperform PB while MB will outperform both* ($\beta_3 > \beta_2 > \beta_1$).

4.1.2 Adaptability

We define adaptability as the most profitable strategy across rounds and markets. The criteria seeks to determine which pricing mechanism performs best across all sources of variation introduced by the experiment. We measure adaptability as simply the attained profit by each individual in each round:

$$Y_{it}^{APT} = \pi_{it}^{att}. \tag{3}$$

We then estimate a regression controlling for all possible sources of variation:

$$Y_{it}^{APT} = \alpha + \beta_1 PB_i + \beta_2 BSP_i + \beta_3 MB_i + \tau_t + \gamma_g + \theta_i + \delta_d + \rho_r + (\delta_d \times \rho_r) + \epsilon_{it}. \tag{4}$$

As before, γ_g controls for the number of goods and θ_i controls for the order in which an individual saw each pricing strategy. We also control for the distribution of demand (δ_d), the correlation of values between goods (ρ_r), and their interaction. Finally we include round fixed effects (τ_t) and cluster errors at the participant-level. The resulting regression allows us to compare performance of one pricing scheme to another within the exact same environment.

Given our controls, the regression estimates the average treatment effect of each pricing strategy on the profit attained by individuals in each round, in each market, for each number and type of goods in the experiment. It answers the question: when comparing one pricing strategy to another in the each of our settings, which one, on average, performs the best. Based on existing models of non-linear pricing mechanisms, hypothesis 2 states:

Hypothesis 2 *All three complex pricing strategies will be more adaptable than CP* ($\beta_1 > 0;\ \beta_2 > 0;\ \beta_3 > 0$). *Furthermore, BSP will be more adaptable than PB while MB will be more adaptable than both* ($\beta_3 > \beta_2 > \beta_1$).

4.1.3 Effectiveness

Effectiveness is defined as the share of the maximum achievable profit obtained by a pricing strategy. We measure effectiveness as:

$$Y_i^{EFF} = \frac{\sum_{t=1}^{40} \pi_{it}^{att}}{40 * \pi_i^{max}} \tag{5}$$

This gives us one value for each participant - the share of total profits they obtained by summing across all rounds. Our regression is:

$$Y_i^{EFF} = \alpha + \beta_1 PB_i + \beta_2 BSP_i + \beta_3 MB_i + \epsilon_i. \tag{6}$$

By including maximum achievable profits on the left hand side of the equation we effectively control for market, type of goods, and the number of goods. This means there is no need to control for them directly on the right hand side of the equation.

The regression estimates the average treatment effect of each pricing strategy on the share of possible profits attained by participants, controlling for the correlation in valuation between goods, the distribution of consumer valuations, or the number and type of goods on sale. It answers the question: on average, which pricing strategy allows individuals to obtain profits closest to the maximum achievable profit for the market they are facing. Hypothesis 3 states.

Hypothesis 3 *All three complex pricing strategies will increase share of profits attained relative to CP ($\beta_1 > 0;\ \beta_2 > 0;\ \beta_3 > 0$). Furthermore, BSP will outperform PB while MB will outperform both ($\beta_3 > \beta_2 > \beta_1$).*

4.1.4 Convergence

The convergence criteria seeks to assess the most profitable strategy across rounds. This can be visualized as the area between the maximum obtainable profit and profit attained by a participant (see Figure 1). Mathematically, we measure convergence as:

$$Y_{bi}^{CNV} = \sum_{t=1}^{10} \left(\pi_{bi}^{max} - \pi_{bit}^{att} \right). \tag{7}$$

This gives us a single value for each participant for each pricing strategy - the sum of the distance between max profit and achieved profit for each pricing strategy. We then estimate:

$$Y_{bi}^{CNV} = \alpha + \beta_1 PB_i + \beta_2 BSP_i + \beta_3 MB_i + \gamma_g + \kappa_j + \epsilon_{bi}. \tag{8}$$

Again, we control for the number of goods sold (γ_g) and the order in which an individual saw each pricing strategy (θ_i) but estimate the effect of each pricing strategy regardless of market (distribution and correlation). Including fixed effects for the number of goods allows us to control for the fact that MB with four goods is inherently more difficult than MB with two goods and thus we would only want to compare MB with four goods to PB with four goods, not PB with two goods. We cluster the standard errors by participant.

The regression estimates the average treatment effect of each pricing strategy on the profit attained by individuals across all rounds in the experiment. It answers the question: on average, which pricing strategy allows individuals selling the same goods to converge quickest on the

Figure 1: Convergence of pricing strategies across rounds

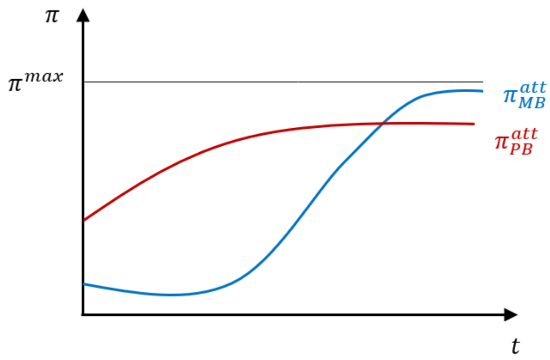

maximum profit, regardless of market. Based on the number of prices each strategy requires, hypothesis 4 is as follows.

Hypothesis 4 *As PB requires setting only one price, it will converge faster than the other three pricing strategies* ($\beta_1 < 0$; $\beta_1 < \beta_2$; $\beta_1 < \beta_3$). *Furthermore, since BSP requires setting the same number of prices as CP, we expect both to converge at the same rate* ($\beta_2 = 0$). *Finally, since MB requires setting more prices than the other three pricing strategies, we expect MB to converge slower than CP, PB, and BSP* ($\beta_3 > 0$; $\beta_3 > \beta_1$; $\beta_3 > \beta_2$).

4.1.5 Maximin

The maximin criteria is designed to determine which strategy performs best at limiting losses. To determine which pricing strategy has the largest minimum profits, we first find the minimum profit earned by each participant with each pricing strategy:

$$Y_{bi}^{MXN} = \min_{\pi_{bi}} \pi_{bit}^{att}. \tag{9}$$

Since participants have zero marginal cost in producing goods and free disposal of unsold goods, the absolute minimum profit a participant could earn in any given round is zero. We then estimate:

$$Y_{bi}^{MXN} = \alpha + \beta_1 PB_i + \beta_2 BSP_i + \beta_3 MB_i + \gamma_g + \kappa_j + \delta_d + \rho_r + (\delta_d \times \rho_r) + \epsilon_{bi}. \tag{10}$$

As before, γ_g controls for the number of goods and θ_i controls for the order in which an individual saw each pricing strategy but also the distribution of demand (δ_d), the correlation of values between goods (ρ_r), and their interaction. This allows us to compare the minimum values earned by each pricing strategy within markets and within number and type of goods. Standard errors are clustered by participant.

The regression estimates the average treatment effect of each pricing strategy on producing the largest minimum profits obtained by individuals within a given market conditional on selling the same number and types of goods. It answers the question: on average, which pricing strategy results in largest minimum profit, controlling for all other factors. We form our null hypothesis 5 by ranking pricing strategies based on the number of options available for consumers to purchase. The intuition is that since PB is an all-or-nothing decision for consumers, it is likely that the bundle is set too high for any consumer to find value it purchasing it. Conversely, MB, which sets many prices, is likely that a consumer will find value in purchasing some set or sub-set of goods.

Hypothesis 5 *MB will earn the largest minimum profit compared to the other three pricing strategies* ($\beta_3 > 0;\ \beta_3 > \beta_1;\ \beta_3 > \beta_2$)*. Since BSP offers the same number of options as CP, we expect both to earn the same minimum* ($\beta_2 = 0$)*. Finally, since PB offers only one option, it will earn the smallest minimum profit* ($\beta_1 < 0;\ \beta_1 < \beta_2;\ \beta_1 < \beta_3$)*.*

4.1.6 Variability

Our final criteria seeks to determine which pricing strategy minimizes the variance of profits earned across individuals. We measure variability as the second moment of each participant's distribution of achieved profits for each pricing strategy:

$$Y_{bi}^{VAR} = \mathrm{Var}\left[\pi_{bit}^{att}\right] = \mathrm{E}\left[\left(\pi_{bit}^{att} - \mu_{bi}\right)^2\right], \tag{11}$$

where μ_{bi} is the average profit attained by an individual for a pricing strategy. We then estimate:

$$Y_{bi}^{VAR} = \alpha + \beta_1 PB_i + \beta_2 BSP_i + \beta_3 MB_i + \kappa_j + \epsilon_{bi}. \tag{12}$$

Here we do not control for the order in which a participant saw a strategy and not the market, number of goods, or type of goods. This is because we are interested in comparing variance in profit across all settings and not within each setting.

The regression estimates the average treatment effect of each pricing strategy on the variability in profits earned by participants. It answers the question: which pricing strategy, on average, has the smallest variability in attained profit, regardless of the specifics on time, market, and the number and type of goods. Similar to the maximin criteria, we rank pricing strategies based on the number of options available to consumers to purchase. The intuition is that with more option, a participant will end up in fewer all-or-nothing setting, which will reduce the variance in attained profits.

Hypothesis 6 *MB will reduce variance in profits compared to the other three pricing strategies* ($\beta_3 < 0;\ \beta_3 < \beta_1;\ \beta_3 < \beta_2$)*. Since BSP offers the same number of options as CP, we expect profits from both to have the same variance* ($\beta_2 = 0$)*. Finally, since PB offers only one option, it will have the largest variance in profits* ($\beta_1 > 0;\ \beta_1 > \beta_2;\ \beta_1 > \beta_3$)*.*

4.2 Experimental Treatments

The goal of our experimental treatments is to understand the importance of perfect information versus bounded rationality in solving the complex pricing problems. The standard theoretical set-up of the firm and consumer problem relies on two key assumptions. First is that the monopolist has perfect information regarding the distribution of consumer preferences. Second is that both the firm and the consumers do not suffer from bounded rationality. By using human participants in the experiment, we allow for individuals to be less than fully rational, relaxing a key element in the second assumption. In all likelihood, many of the student participants will be unable to perfectly solve the complex pricing problems posed to them in the experiment. To understand the importance of the first assumption, regarding perfect information, our three treatments remove information about the demand space participants face. The use of human subjects in combination with the experimental treatment will allow us to compare performance in the theoretical ideal relative to relaxing each of the two assumptions. Results from our treatments will allow us to understand which information is more important (distribution or correlation) and how performance suffers from loss of this information compared to simply being boundedly rational.

We expect performance across all criteria to fall when we remove information. We expect performance to fall by the same degree regardless of whether we remove information about distribution or correlation. We expect performance to be the worst when we have removed information about both distribution and correlation.

We will test these hypotheses using all six of our criterion as defined above. For each criteria, we will test the interaction between the pricing scheme and the treatment. For a generic criteria, we estimate:

$$Y = \alpha + \beta_1 PB + \beta_2 BSP + \beta_3 MB + \beta_4 (PB \times T) + \beta_5 (BSP \times T) + \beta_6 (MB \times T) + \beta_7 T + \epsilon, \quad (13)$$

where T is an indicator for the treatment being tested. This leads us to our final hypothesis:

Hypothesis 7 *A lack of information will result in lower profits* ($\beta_7 < 0$). *Furthermore, a lack of information will reduce performance of each pricing strategy relative to its performance in the control group* ($\beta_1 > \beta_4;\ \beta_2 > \beta_5;\ \beta_3 > \beta_6$). *Finally, a lack of information will have a larger negative effect for pricing strategies that require setting more prices* ($\beta_4 > \beta_7;\ \beta_5 = \beta_7;\ \beta_6 < \beta_7$).

5 Secondary Outcomes and Supplementary Data

For our secondary outcomes we have pre-specified variable definition but not estimation strategies. This is because there are not strong theoretical predictions that allow us to form strong priors for hypothesis testing. Rather, estimation strategies for the remaining outcomes will be settled on during the pre-specified exploratory stage of analysis. We will then update the pre-analysis plan and confirm these hypotheses on the remaining (non-exploratory) sample.

5.1 Decision-making Process

The experimental software allows us to track two key indicators of player behavior during the game:

1. Response time (time spent in each round);
2. Values entered into the price fields.

For each player, we can reconstruct the options tried for each price and the total time spent completing this process.

One avenue we think worth exploring with this data is how many times participants enter "provisional" prices before settling on their final price. We measure the difficulty of a task by 1) how many provisional prices a player entered and 2) the distribution of those provisional prices. We can also measure difficultly by the length of time each participant spent on setting prices, both in each round and over the entire experiment.

5.2 Measuring Emotional Responses

In addition to participants' choices in each round, we measured their "emotional response" over the course of the game. After entering their prices but before learning their profit, we elicit participants' current emotional state. To do this, we created a slide bar that corresponds to four emojis. Figure 2 shows an example of the table as seen by participants. The goal is to assess how people react to information and performance and determine if this correlates with past and future performance.

Figure 2: Emoji interface example

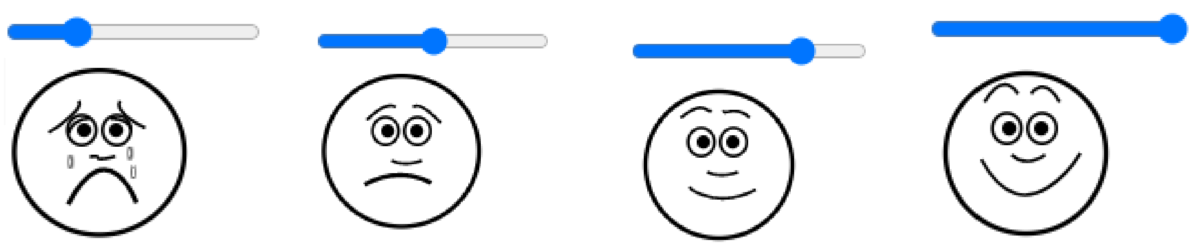

References

Anderson, M. L. and J. Magruder (2017). Split-sample strategies for avoiding false discoveries. NBER Working Paper No. 23544.

Deutschmann, J. W., J. D. Michler, and E. Tjernström (2018). Complex pricing and financial stress. AEA RCT Registry. April 19. https://doi.org/10.1257/rct.2904-1.0.

Rubenstein, A. (1997). *Modeling Bounded Rationality.* Cambridge: M.I.T. Press.