Essays on Economic Design: Theory and Estimation

HARVARD UNIVERSITY
Graduate School of Arts and Sciences

DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the

Department of Economics

have examined a dissertation entitled

"Essays on Economic Design: Theory and Estimation"

presented by **Olivia Reimann Hartzell**

candidate for the degree of Doctor of Philosophy and hereby certify that it is worthy of acceptance.

Signature Ariel Pakes

Typed name: Prof. Ariel Pakes

Signature ____________________

Typed name: Prof. Robin Lee

Signature ____________________

Typed name: Prof. Jerry Green

Date: July 24th, 2025

Essays on Economic Design: Theory and Estimation

A dissertation presented

by

Olivia Reimann Hartzell

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

July 2025

Dissertation Advisor: Professor Ariel Pakes Olivia Reimann Hartzell

Essays on Economic Design: Theory and Estimation

Abstract

This thesis studies optimal economic design choices, both theoretically and empirically, in environments where some market participants themselves make welfare-relevant design choices.

In the first two chapters, I study a platform's preferencing mechanism, or rule that assigns the placement of offers on a platform. I theoretically derive a platform's consumer-optimal preferencing rule in a setting in which firms choose prices in a locally envy-free profile, given the preferencing rule, which in special cases, correspond with equilibria in a Generalized Second Price auction. Additionally, I show that Reimers and Waldfogel (2023) empirical test to measure platform bias is valid even under prices endogenous to the preferencing rule in a locally envy-free profile. On Amazon, I measure the extent of platform bias using a large dataset from the platform during the years 2020-2022.

I then extend this study to an environment in which firm prices are determined in a non-myopic, dynamic setting, which depends on the platform's preferencing rule. These dynamic pricing strategies result in pricing patterns that closely resemble Edgeworth cycles. I provide a method to estimate the primitives that govern these pricing cycles, which allows me to assess the counterfactual welfare implications associated with various preferencing rules, using a large dataset from the Amazon platform from 2018-2022. The welfare effects of policy proposals in this environment, in particular those that eliminate self-preferencing, significantly differ from those under the static price competition environment.

The last chapter studies a class of complete information quantum games, where design choices in quantum versions of classical games, in particular the chosen basis and initial

state, determine the set of feasible outcome distributions in a Nash equilibrium in quantum strategies. I derive necessary and sufficient conditions for a Nash equilibrium in quantum strategies to Pareto improve upon classical correlated equilibria in certain classes of quantum games.

Contents

List of Tables

List of Figures

Acknowledgments

I am especially grateful to my advisors for their guidance. Ariel Pakes cultivated in me an approach of letting the data tell the story of my research and encouraged me to persevere when I wanted to give up. Robin Lee helped me distill my often-times unwieldy models to get to the root of the research question while never losing sight of the larger picture. Jerry Green encouraged and supported my unconventional ventures into unstudied theoretical territory that crossed departments and disciplines.

My research has greatly benefited from conversations and helpful comments over the years from Ed Glaeser, Yannai Gonczarowski, Myrto Kalouptsidi, Shengwu Li, Eric Maskin, and Elie Tamer.

Jeff Miron provided thoughtful advice in both academics and life. Brenda Piquet, fountain of wisdom, reminded me to climb my own ladder and no one else's.

I am particularly grateful to Andy Haupt for the invaluable lessons that I learned during our time together.

Many graduate students in the department have been supportive friends and colleagues over the past five years, including Sachet Bangia, Sage Belz, Lindsey Currier, Adriano Fernandes, Jeff Gortmaker, Tilman Graff, Stephanie Kestelman, Fanele Mashwama, Chris Walker, and Jennifer Zou.

In more recent years, Jacob Barandes generously welcomed me to the physics department and discipline. I am also grateful to Adam Brandenberger for sharing his wisdom on quantum games.

The bulk of this thesis would not have been completed without the generous contribution of an anonymous data provider, computational support from MIT CSAIL, and the help of my research assistant, Anton Tarasenko.

I attended graduate school in large part thanks to the encouragement of Dan O'Flaherty, Jose Montiel Olea, and John Asker. John Hartzell and Andrew C. Hartzell Jr. were instrumental in encouraging and supporting my return to academics after I retired from the stage.

I owe a great deal to my parents and three brothers, along with my extended and chosen family and friends for their love and support throughout the past five years and beyond. My dog Ruby quite literally carried the emotional burden of a PhD for not only myself, but also for the many graduate students in the department who have graciously cared for her.

Lastly, the past five years have afforded me the opportunity to witness a profound journey towards inner wisdom that has irrevocably shaped my intellectual and creative work. I am deeply grateful to all those who have guided me on the path.

To Mary Lee Hartzell, a true artist

Introduction

Mechanism design studies how an economic engineer can design a marketplace to optimize over welfare-relevant outcomes given market participants' incentives. This field has been highly successful insofar as using theoretical results to guide the implementation of real-world marketplaces that satisfy desirable properties such as efficiency and dominance incentive-compatibility.

However, in certain practical economic settings, marketplace features may render both determining and implementing a welfare-optimal protocol more challenging. In this thesis, I study settings less-common in the canonical economic design literature, in which at least some of the market participants themselves make welfare-relevant design choices.

The first two chapters study marketplaces in which the designer, in this case a platform, also acts as a market participant, and therefore her incentives may not align with overall market efficiency.

The first chapter theoretically derives a consumer welfare-optimal design choice in an environment in which the platform is a participating producer. In particular, this chapter focuses on a design rule called *preferencing*, which determines the assignment of offers on a platform. Unlike in other environments such as auctions, in this setting, the theoretical result alone is not sufficient to practically implement the welfare-optimal rule – the designer must also precisely determine consumer preferences in order to implement the correctly-efficient preferencing rule. In this setting, these preferences, along with a methodology to estimate the existing preferencing rule, can also be used to determine *how far* the designer, in this the case the platform, diverges from the welfare-optimal rule. This procedure may be particularly useful for regulators as a first-pass welfare assessment in many real-world marketplaces.

In some settings, the same design choice may warrant the study of multiple induced models of market participant behavior. The second chapter studies solving for the optimal preferencing rule in the same market environment, but in which participants respond with dynamic behavior, as opposed to static behavior. Indeed, these participants may additionally fail to fully account for opponent behavior, as in the theoretical model studied in the first chapter. This chapter studies the effects of marketplace design choices on welfare when deriving theoretically-optimal rules may not be feasible, and assessing welfare implications requires careful estimation. In this setting, I compare the simulated counterfactual welfare effects of various design choices and determine the welfare loss associated with allowing for dynamic versus static behavior on the platform.

The final chapter studies settings in which *all* market participants act as designers in choosing certain parameters of their environment that dictate their feasible outcomes. In particular, I study normal form games of complete information in which players have access to quantum entanglement and show how design choices in these environments enable distributions over outcomes that may not coincide with classical Nash or correlated equilibria. In quantum settings, certain design choices may restrict the distributions over feasible outcomes and deviations. In this way, the quantum setup serves as a commitment device that players must jointly agree to opt into. In some very common economic settings, restricting the set of feasible outcomes *improves* their expected payoffs relative to classical Nash or correlated equilibria.

Chapter 1

Platform Preferencing and Price Competition I[1]

1.1 Introduction

In this chapter, we define a *preferencing* rule as an assignment rule that allocates featured or preferential status to an offer on a platform. The practice of preferencing, absent any effects on prices or foregone consumer surplus from limiting choices, may be beneficial to consumers who face search costs from choosing between multiple offers. However, recent regulatory concerns have emerged regarding such preferencing practices, as they may hinder price competition, restrict access to more preferred, non-featured offers, and adversely affect long-term market entry and competition from third-party sellers.

Policy proposals have gone so far as to suggest that certain preferencing practices be made downright unlawful. The American Innovation and Online Choice Act (U.S. Senate, 2021) proposes to prohibit practices that preference the "platform over those of another operator in a manner that would materially harm competition," while the recently implemented Digital Markets Act (European Union, 2022) forbids practices defined as "self-preferencing," or those "understood to be a more favourable treatment in ranking and related indexing

[1]Co-authored with Andreas Haupt

and crawling of first-party products and service than third-party offers."

Under this definition of self-preferencing, these policies propose that all else equal, a platform cannot allocate featured status to their own products over competitors. However, this definition fails to account for instances in which a platform's product is *fundamentally* preferred by consumers to non-platform offers, and the *extent* to which this preference should grant the platform featured status. In this paper, we will define *self-preferencing* as any preferencing rule that awards a higher probability of being awarded featured status to the platform itself, relative to a non-platform offer, than what is justified by consumer preferences. This definition provides a measurable means to assess whether a preferencing practice is disproportionately advantageous to a platform, which we will provide mathematical precision to in the sections that follow.

In order to study the effects of preferencing on competition and its potential harms to consumer and producer welfare more closely, we first theoretically study prices that are endogenous to a particular class of preferencing rules. We show the relationship between a particular instance of our theoretical setting and equilibrium outcomes in the well-known Generalized Second Price auction (Edelman et al., 2007; Varian, 2007). Given our model for endogenizing prices, we are then able to assess the effects of various preferencing rules on producer prices and resulting consumer welfare, and empirically measure the extent of self-preferencing.

We conduct the proposed self-preferencing test on the Amazon platform, which has recently faced considerable scrutiny over the potential harms of its preferencing practices. A lawsuit brought against Amazon by the Federal Trade Commission (United States Federal Trade Commission, 2023) accuses the platform of practices that prevent sellers from competing "on equal footing against Amazon's private label products." Additionally, the lawsuit alleges that Amazon "purposefully suppresses information about competing products to give its own private label products an artificial boost" and "degrades the quality of the shopping experience on Amazon by replacing helpful organic search results with biased 'widgets' that direct shoppers to purchase Amazon's private label products, . . . and display[s]

Amazon's private label products over other products sold on Amazon."

In this chapter, we focus most of our study of self-preferencing on Amazon product pages, whose preferencing mechanism is governed by an assignment rule which allocates a single offer on a product page to the "Buy Box" also known as the "Buy Now" box. When a consumer arrives on a given product page on the platform, they purchase directly from the individual offer assigned featured status under the Buy Box when it clicks on the "Buy Now" button. In order to view the other available, non-featured offers on the product page, the consumer must click on an additional button at the bottom of the page to expand information on these remaining offers. Indeed, products that are awarded the Buy Box are reported to face considerable advantages. It is estimated that about 80 % of sales on Amazon occur in the Buy Box, and that purchase rates on mobile devices are even higher (Feedvisor, 2021). An example of a product page with an offer assigned to the Buy Box is shown in Figure 1.1, with all corresponding non-Buy Box offers shown in Figure 1.2. In this example, the Buy Box is assigned to Amazon, and is not the lowest price offer available on the product page.

Previous work has suggested that Amazon engages in self-preferencing when awarding the Buy Box, relative to other, third-party offers that may vary based on fulfillment method, or Prime status (Raval, 2022). Offers on Amazon may be fulfilled in one of two ways, either Fulfilled by Amazon (FBA), or Fulfilled by Merchant (FBM). Offers that are FBA, may be sold by a non-Amazon seller, but packed and shipped by Amazon for a fee charged by the platform, which typically depends on the size and weight of the order. All offers sold by Amazon are FBA. FBM offers are packed and shipped by the third-party merchants themselves, or by some other non-Amazon fulfillment service. Additionally, a third-party offer may be Prime eligible, meaning it is eligible for free two-day shipping to consumers who are enrolled in Amazon's Prime membership program. All FBA products are Prime eligible, while some FBM products may be.

The empirical test that we conduct allows us to measure the extent to which third-party offers should be preferenced relative to each other and Amazon offers, if the platform

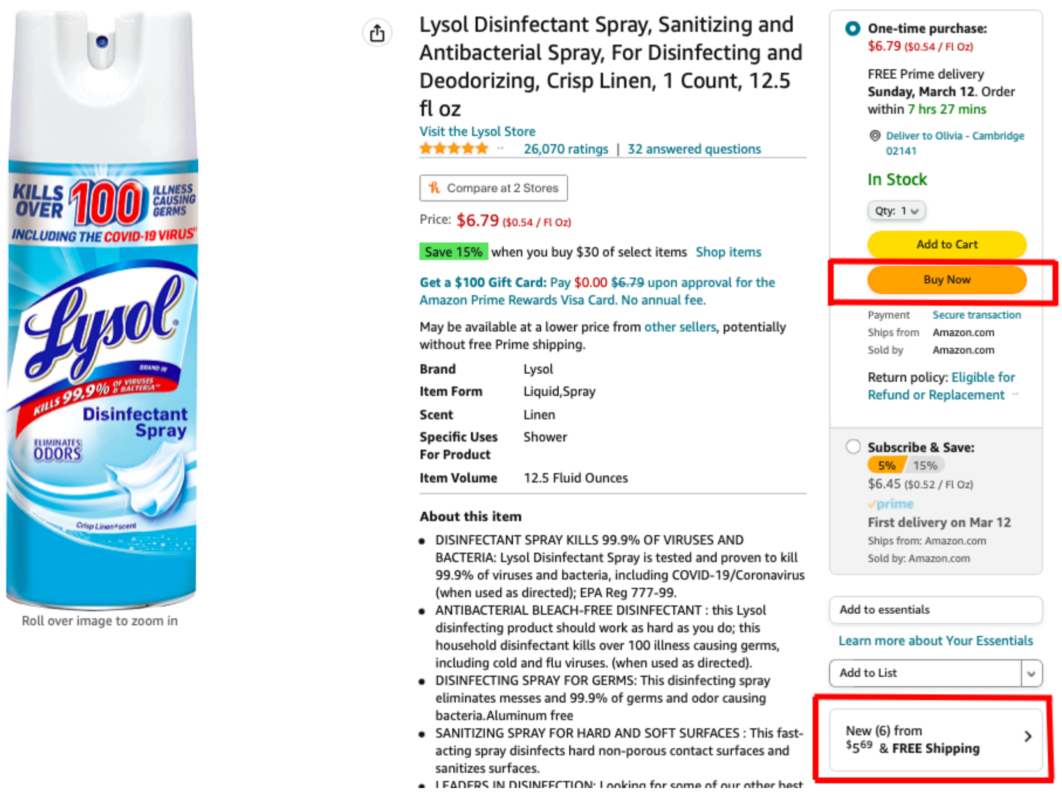

Figure 1.1: *Amazon product page with assigned Buy Box.*

awarded the Buy Box according to the rule that would maximize consumer surplus under our theoretical model, and measure the extent to which Amazon's current rule deviates from this rule. Additionally, our data allow us to provide evidence for the extent to which keyword search rankings are optimized for purchases and consumer surplus.

1.1.1 Related Literature

Several papers have theoretically studied the various trade-offs associated with platform preferencing and biased intermediation. Hagiu and Jullien (2011) consider intermediaries' motives for diverting consumer search, including increasing revenue per customer and lowering sellers' prices. Bergemann and Bonatti (2011); Levin and Milgrom (2010) study a platform's trade-off between match quality and supply side market power, while (Hagiu and Jullien, 2014) considers the role of across-platform competition in driving these trade-offs

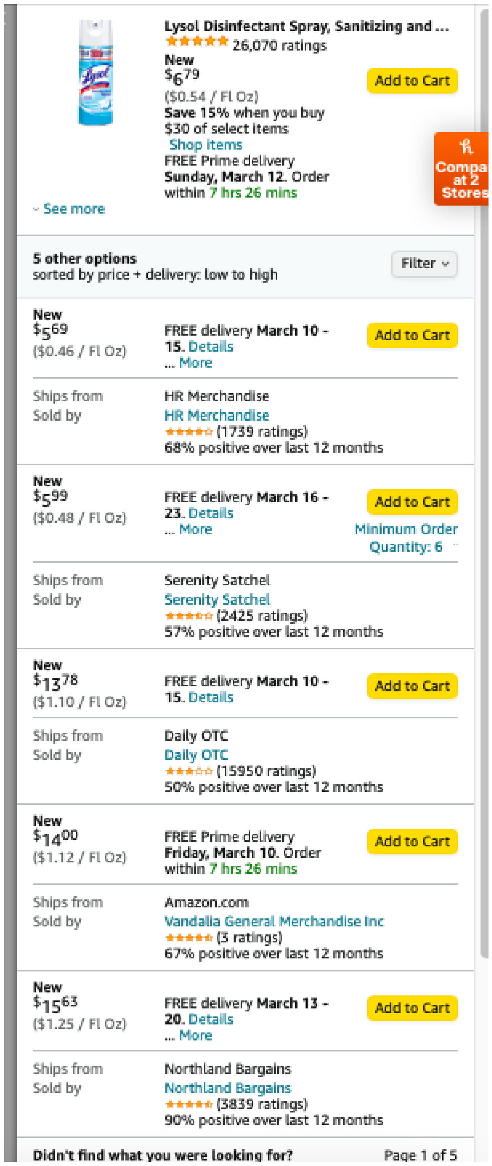

Figure 1.2: *Amazon product page with remaining offers not assigned the Buy Box.*

in platforms' decisions to divert consumer traffic via advertisements. Similarly, Gur et al. (2023) and Ke et al. (2022) use an information design perspective to study a platform's optimal preferencing rule, given similar trade-offs when sellers can learn the value of being preferenced. Additionally, Armstrong and Zhou (2011) compares the effects of paid advertisements and commission payments on equilibrium prices and finds that advertising payments decrease prices under high search costs.

More recent work has studied the harms of self-preferencing through under-investment in product discovery (Hagiu et al., 2022), and the effects of platform preferencing rules on supply-side entry decisions (Zou and Zhou, 2023), while Anderson and Bedre-Defolie (2021) endogenize entry decisions and optimal commission rates for hybrid and non-hybrid platforms, concluding that banning hybrid mode improves consumer welfare. More closely related to our work, de Cornière and Taylor (2019) finds that biased preferencing rules have varying effects depending on the extent to which seller and consumer utility co-vary in equilibrium. This paper instead studies optimal preferencing design by modeling the effects of a platform's preferencing rule on supply-side prices, and therefore consumer and producer welfare.

Our model maps closely to related market design settings, specifically within internet advertising auction design. Within the vast literature on advertising auctions, most closely related to our setting, Edelman et al. (2007) and Varian (2007) expand on the seminal auction literature (Vickrey, 1961; Clarke, 1971; Groves, 1973) by providing an equilibrium selection criterion that yields ex-post Nash equilibria whose positions and payments are equivalent to the Vickrey-Clarke-Groves (VCG) auction mechanism, while Varian and Harris (2014) extend the model to a setting with incomplete information. Athey and Ellison (2011) further considers the welfare implications of position auctions with a micro-founded model of consumer search. In a special case, we relate prices endogenous to the preferencing rule to equilibrium outcomes in Edelman et al. (2007) in section 1.2.1.

More broadly, several papers have studied auction environments in which bidders make joint price and quantity or quality decisions (Asker and Cantillon, 2008; Che, 1993;

de Cornière and Taylor, 2019; Nishimura, 2015). In contrast to other scoring auction studies, our model allows us to characterize prices endogenous to any scoring rule and general form of seller types. In the second part of the theoretical model, we use this notion to characterize the consumer-optimal preferencing rule, rather than the rule that maximizes revenue for the auctioneer, which under certain conditions, allows us to characterize a minimum level of welfare that the designer can ensure even with no information over sellers' types.

More recent empirical work has been dedicated to studying and detecting the extent to which platforms preference their own products beyond that which is consumer optimal, citing the need to empirically address questions of increasing regulatory concern (Aguiar et al., 2021; Bourreau and Gaudin, 2022; Reimers and Waldfogel, 2023). Jürgensmeier and Skiera (2023) build on this literature by proposing a framework to measure the extent to which a platform preferences their own products relative to third-party sellers, but they do not define nor measure the consumer-optimal rule given consumer preferences. Our theoretical model instead extends the framework described in Reimers and Waldfogel (2023) to a setting that endogenizes prices given supply-side competition and we perform this empirical test on the Amazon platform.

In other empirical work that has focused specifically on Amazon, Lam (2021) estimates a structural search model of consumer demand that embeds internet advertisements and finds that the effects of search behavior offset the welfare benefits from improved consumer-product matches, while Gutiérrez (2021) examines the welfare consequences of various regulatory interventions to suggest that policies that ban self-preferencing may be offset by other platform practices.

While Farronato et al. (2023) give descriptive evidence for Amazon's self-preferencing practices in keyword search, the authors do not conduct the empirical tests in the style of Reimers and Waldfogel (2023) to test for self-preferencing, nor do they embed an equilibrium model of prices. Additionally, Yu (2024) studies keyword search rankings and runs counterfactual exercises that eliminate sponsored advertising search positions using a structural model of consumer demand and platform equilibrium fees.

Previous empirical studies of preferencing on the Amazon platform have been limited by access to sufficient information about consumer demand to determine the extent to which Amazon engages in self-preferencing. Our richer data relative to previous papers permits us to estimate both Amazon's Buy Box assignment rule and consumer demand with more precision, and therefore enables us to conduct self-preferencing tests empirically.

1.2 Theory

We consider the platform's problem of assigning sellers $j \in J = \{1, 2, \ldots, N\}$ to a *ranking*, given by a permutation $\mathbf{r}$ of J.

The platform's *preferencing rule* is given by a mapping

$$g \colon J \mapsto \Delta(S_J)$$

where $\Delta(S_J)$ is the set of all probability distributions over S_J, the symmetric group of J, or the set of all permutations of J. For every permutation $\mathbf{r} \in S_J$, we will denote j's ranking by r_j, which is given by the jth element in $\mathbf{r}$.

Sellers $j \in J$ have vectors of fixed characteristics $X_j \in \mathbb{R}^d$, and choose prices $p_j \in P$ where $P \subset \mathbb{R}$. The tuple $(X_1, X_2, \ldots, X_N)$ and the vector $(p_1, p_2, \ldots, p_N)$ will be denoted by $\mathbf{X}, \mathbf{p}$, respectively.

We are concerned with ranking rules g over J that correspond with a *score* function,

$$f \colon \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}_{\geqslant 0}\, (p_j, X_j) \mapsto f(p_j, X_j),$$

which is monotonically non-increasing in price, $f(p_j, X_j) \geqslant f(p_j', X_j)$ for all $p_j \leqslant p_j'$, and for which $f_j(p_j, X_j) \geqslant 0$ for all $p_j \in P$. We will, in a slight abuse of notation, refer to $f(p_j, X_j)$ as $f_j(p_j)$, given that characteristics X_j are fixed.

We will restrict ourselves to preferencing rules, where for any ranking r given positive probability under the assignment rule,

$$r_j \leqslant r_k \iff f_j(p_j) \geqslant f_k(p_k)$$

and ties are broken uniformly. In other words, g assigns rankings in $\{1, 2, \ldots, N\}$, where lower ranks are awarded to higher scoring sellers, given the scoring rule f, and where any ties are broken with equal probability.

Sellers have payoff functions

$$\pi_j \colon \mathbb{R}^N \times \mathbb{R}^{d \times N} \times S_J \to \mathbb{R}, (\mathbf{p}, \mathbf{X}, \mathbf{r}) \mapsto \pi_j(\mathbf{p}, \mathbf{X}, \mathbf{r}),$$

for all $j \in J$.

Each player has a reservation utility, $\overline{\pi}_j \geqslant 0$, which is the utility that she obtains from obtaining the last ranking, $r_j = N$. We assume that $\pi_j(\mathbf{p}, \mathbf{X}, \mathbf{r}) > \pi_j(\mathbf{p}, \mathbf{X}, \mathbf{r}')$ for any $\mathbf{p}$, for any rankings $\mathbf{r}, \mathbf{r}'$ in which $r'_j > r_j$. Additionally, we assume that $\pi_j(\mathbf{p}, \mathbf{X}, \mathbf{r}) > \pi_j((p'_j, \mathbf{p}_{-j}), \mathbf{X}, \mathbf{r})$ for all $p_j > p'_j$ for any rankings $\mathbf{r}$ in which $r_j < N$, and $\pi_j(\mathbf{p}, \mathbf{X}, \mathbf{r}) \geqslant \pi_j((p_j, \mathbf{p}'_{-j}), \mathbf{X}, \mathbf{r})$ for all $\mathbf{p}'_{-j} \geqslant \mathbf{p}_{-j}$. That is, obtaining a lower rank strictly improves one's own payoff, and for any rankings in which $r_j < N$, payoffs are strictly increasing in own prices and non-decreasing in opponent prices.[2]

1.2.1 Nash Equilibria and Locally Envy-Free Profiles

Sellers simultaneously choose prices $\mathbf{p}$, given the scoring rule f, sellers' characteristics $\mathbf{X}$, and payoff functions π, all of which are all common knowledge.

A strategy profile $\mathbf{p} \in \mathbb{R}^N$ is a *Nash equilibrium* if for all agents $j \in J$ and any price p'_j,

$$\mathbb{E}[\pi_j(\mathbf{p}, \mathbf{X}, \mathbf{r})] \geqslant \mathbb{E}[\pi_j((p'_j, \mathbf{p}_{-j}), \mathbf{X}, \mathbf{r})].$$

where the expectation is with respect to the distribution over all rankings $\mathbf{r}$ under the preferencing rule.

In practice, sellers on platforms, including Amazon, may update their prices very frequently. Studying behavior in a setting where price updates are *dynamic*, and where sellers are not myopic, is the focus of chapter 2. In this chapter, we begin our study of

[2]The assumption that payoffs are strictly increasing in own-prices rules out settings in which some players' reservation prices lie above others' profit maximizing prices.

the effects of preferencing on prices and welfare in a setting in which sellers may update their prices dynamically, albeit according to myopic payoff functions. These agents may be representative of sellers who do not invest in sophisticated strategies nor pricing technologies that optimize over more complicated, non-myopic payoff functions. Therefore, outcomes in these environments may serve as a useful benchmark for considering behavior under various preferencing rules.

In particular, we are interested in the set of prices and corresponding rankings of the simultaneous one-shot game that should coincide with the potential *rest points* in pricing behavior. Given a scoring rule f, these rest points must satisfy a property known as *local envy-freeness* following Edelman et al. (2007).

Definition 1. *A strategy profile $\mathbf{p} \in \mathbb{R}^N$ of the simultaneous move preferencing game is* locally envy-free *if no player can improve their payoff by exchanging scores and ranks with the player ranked immediately above or below them.*

Formally, in a locally envy-free profile, the vector of prices $\mathbf{p}$ and (realized) ranks $\mathbf{r}$ satisfy:

$$\pi_j(\mathbf{p}, \mathbf{X}, \mathbf{r}) \geqslant \pi_j(\mathbf{p}', \mathbf{X}, \mathbf{r}'),$$

for all sellers j with $r_j < N$, where under $\mathbf{p}', \mathbf{r}'$, j obtains $r'_j = r_j + 1$ at p'_j such that $f_j(p'_j) = f_k(p_k)$, where k is the seller ranked immediately above j ($r_k = r_j + 1$). Seller k obtains $r'_k = r_j$ at price p'_k such that $f_k(p'_k) = f_j(p_j)$, and all other sellers $i \neq j, k$ prices and ranks remain unchanged, $p_i = p'_i, r_i = r'_i$.

For all sellers j for which $r_j > 1$, where under $\mathbf{p}', \mathbf{r}'$ k obtains $r'_j = r_j - 1$ at p'_j such that $f_j(p'_j) = f_k(p_k)$, where k is the seller ranked immediately below j ($r_k = r_j - 1$). Seller k obtains $r'_k = r_j$ at price p'_k such that $f_k(p'_k) = f_j(p_j)$, and all other sellers $i \neq j, k$ prices and ranks remain unchanged, $p_i = p'_i, r_i = r'_i$.

In Edelman et al. (2007), locally envy-free strategy profiles are a *refinement* of the set of Nash equilibria. However, in the preferencing game described above, since sellers' prices depend directly on their scores, a pure strategy Nash equilibrium may not be guaranteed to exist (Edelman and Ostrovsky, 2007). Nonetheless, strategy profiles that satisfy the locally envy-free property will serve as useful predictive measures of resting behavior in the settings that we are interested in.

In particular, studying this myopic setting allows us to map the preferencing problem closely to settings that have been studied in detail in the auction theory literature. Locally envy-free prices and ranks in this setting depend on the relativity of the payoff curves given $\mathbf{X}, \mathbf{r}$, and there may be many locally envy-free profiles. In order to characterize endogenous prices and ranks under stronger assumptions, we define an additional parameter.

We will define for each seller $j \in J$, for any preferencing rule f, for any given feasible such $p_j \in P$, the partial price profile $\underline{\mathbf{p}}_{-j}(p_j)$, such that for all $p_{-j} \in \underline{\mathbf{p}}_{-j}(p_j)$

$$p_{-j} = \inf\{p_{j'} | f_j(p_j) > f_{-j}(p'_{-j})\}$$

$\underline{\mathbf{p}}_{-j}(p_j)$ is the minimum opponent price vector that guarantees that j obtains a strictly higher score than all opponents, when she prices at p_j.

Independent Profits and the Generalized Second Price Auction

We consider the following independent profits case mostly for illustrative purposes, to demonstrate the relationship between the outcomes of the myopic preferencing game and the Generalized Second Price auction.

Consider a special case where $\pi_j((p_j, \mathbf{p}_{-j}), \mathbf{X}, \mathbf{r})$ does not depend on $\mathbf{p}_{-j}, \mathbf{r}_{-j}$. Additionally, for each seller $j \in J$ define

$$p_j^* := \inf\{p_j | \pi_j((p_j, \underline{\mathbf{p}}_{-j}(p_j)), \mathbf{X}, \mathbf{r}) \geqslant \overline{\pi}_j\}$$

which does not depend on $\mathbf{r}$. In other words, the minimum price that each player is willing to obtain is the same for any rank $1 \leqslant r_j < N$. Given the independence assumption with respect to $\mathbf{p}_{-j}, \mathbf{r}_{-j}$, p_j^* also does not depend on f.

We assume that $\overline{\pi}_j$ is sufficiently low such that p_j^* is guaranteed to exist in P for every seller j and $p_j^* < \sup P$. Additionally, we assume that the corresponding score, $f_j(p_j^*) \geqslant 0$ for all scoring rules f.

Under these assumptions, we consider a setting in which sellers' realized payoffs may be written as

$$\pi_j(\mathbf{p}, \mathbf{X}, \mathbf{r}) = \alpha_{r_j}\beta_{X_j}(f_j(p_j^*) - f_j(p_j))$$

where α_r does not depend on j, p_j, X_j, where $\alpha_r > \alpha_{r+1}$ $\forall r < N$ and, for simplicity, $\alpha_N = 0$. That is, the relative benefit to payoffs, or the increase in purchase probability of obtaining any rank $r < N$ does not depend on the price that any player obtains her rank at. $\beta_{X_j} \geqslant 0$, seller j's increase in purchase probability relative to her opponents, does not depend on p_j, r_j.

$f_j(p_j^*)$, or seller $j's$ maximum score that she is willing to obtain at any ranking $r < N$, corresponds with her value per purchase. $f_j(p_j^*) - f_j(p_j)$ is player j's *markup* in her score relative to $f_j(p_j^*)$. Note that when $f_j(p_j)$ is linear in characteristics and prices, and may be written as $f_j(p_j) = X_j'\hat{\beta} - p_j$ for some $\hat{\beta}$, then j's markup in her score corresponds exactly with the markup in her price, $p_j - p_j^*$.

To study locally envy-free profiles in preferencing games where payoffs may be written as above, we consider the related auction setting in which players submit bids b_j given by $f_j(p_j)$, and the resulting ranks are awarded in order of the highest scoring bids. Players' payoffs are given by

$$\pi_j(p_j, \mathbf{X}, r_j) = \alpha_{r_j}\beta_{X_j}(f_j(p_j^*) - t^{r_j})$$

where t^{r_j} is a transfer paid by the player awarded rank r_j.

When transfers are given by $t^{r_j} = b^{(r_j+1)}$, where $b^{(r_j+1)}$ denotes the bid of the player awarded rank $r_j + 1$, this setting corresponds exactly with the Generalized Second Price Auction (GSP) as described in Edelman et al. (2007), where $s_j = f_j(p_j^*)$, and K, the number of possible ranking positions, is equal to N. It is with this correspondence that we may characterize locally envy-free profiles in the preferencing setting.

Consider the following equilibrium strategy profile B^* in the GSP. For bidders j ordered in decreasing value of their per-purchase values, $f_j(p_j^*)$, equilibrium bids are given by

$$b_j^* = \begin{cases} \frac{t^{\mathrm{VCG}(j-1)}}{\alpha_{j-1}} & j > 1 \\ f_j(p_j^*) & j = 1 \end{cases}$$

where $t^{VCG(j)} = (\alpha_j - \alpha_{j+1}) f_{j+1}(p_{j+1}^*) + t^{\mathrm{VCG}(j+1)}$, and $t^{\mathrm{VCG}(N)} = 0$.

Corresponding equilibrum payoffs are given by

$$\pi_j(p_j, \mathbf{X}, r_j) = \alpha_j \beta_{X_j} \left(f_j(p_j^*) - \frac{t^{\mathrm{VCG}(j)}}{\alpha_j} \right)$$

By Edelman et al. (2007), Theorem 1, the equilibrium transfers under strategy profile B^*, constitute the locally envy-free equilibrium of the GSP in which $\sum_{j \in J} t^{r_j}$ are as low as possible. It is easy to see that the strategy profile B^* constitutes a locally envy-free profile in the corresponding preferencing game, where the resulting scores $f_j(p_j)$ are given by the equilibrium transfers $t^{r_j} = \frac{t^{VCG(j)}}{\alpha_j}$ under B^*.

Any locally envy-free profile of the preferencing game must satisfy

$$\alpha_{r_j} \beta_{X_j} (f_j(p_j^*) - f_j(p_j)) \geqslant \alpha_{r_k} \beta_{X_j} (f_j(p_j^*) - f_k(p_k))$$

for all pairs j, k such that $r_j > 1$, $r_k = r_j - 1$, and $r_j < N, r_k = r_j + 1$. These are precisely the conditions that characterize the locally envy-free equilibrium of the GSP in which transfers correspond with $f_j(p_j)$. Given our assumption that all admissible $f_j(p_j) \geqslant 0$, we obtain the following result.

Proposition 1. *The locally envy-free profile of the preferencing game in which score are given by the equilibrium transfers under profile B^*, is the locally envy-free profile in which the sum of the scores is as low as possible.*

Proof. First we show that in any locally envy-free profile of the preferencing game, players must be ranked in monotonically decreasing order of their per-purchase values, $f_j(p_j^*)$. Suppose not. Suppose that there is a locally envy-free profile in which two sellers, i, j obtain ranks $r_i > r_j, r_i = r_j + 1$ and $f_i(p_i^*) > f_j(p_j^*)$. For the profile to satisfy the locally envy-free property, it must hold that

$$f_j(p_j) \geqslant \frac{(\alpha_{r_j} - \alpha_{r_i})}{\alpha_{r_j}} f_i(p_i^*) + \frac{\alpha_{r_i}}{\alpha_{r_j}} f_i(p_i)$$

else seller i would envy j's assignment, and

$$f_j(p_j) \leqslant \frac{(\alpha_{r_j} - \alpha_{r_i})}{\alpha_{r_j}} f_j(p_j^*) + \frac{\alpha_{r_i}}{\alpha_{r_j}} f_i(p_i)$$

else seller j would envy i's assignment. But the above conditions hold if and only if $f_i(p_i^*) \leqslant f_j(p_j^*)$.

The remainder of the proof follows directly from the induction argument outlined in Edelman et al. (2007), Theorem 1, in which transfers p are replaced with scores $f_j(p_j)$.

□

As in GSP, scoring $f_j(p_j^*)$ is not a weakly dominant strategy in the preferencing game. When a player j has a higher $f_j(p_j^*)$ relative to an opponent i, since profits are strictly increasing in prices, she can obtain a strictly higher payoff by setting a score below $f_j(p_j^*)$ so that her opponent is indifferent between obtaining her current rank and increasing her own score, which will always be bounded above by $f_i(p_i^*)$.

Additionally, we note that the prescribed strategy profile is not a Nash equilibrium in the preferencing game, since a seller who obtains rank j can lower her resulting score to just above her next-best opponent, $f_{j+1}(p_{j+1}) + \epsilon$, increase her profits and maintain her rank. Such a deviation, however, would violate the locally envy-free conditions.

Since studying outcomes in full generality without the independence assumption are less tractable, we focus the remainder of our paper on a setting which is more amenable to studying both endogenous prices and consumer welfare.

Single-Winner Invariant Case

We consider a second special case in which only the *first* ranking non-trivially determines payoff functions and payoffs only depend on one's own ranking[3]. More precisely, $\pi_j(\mathbf{p}, \mathbf{X}, \mathbf{r}) = \pi_j(\mathbf{p}, \mathbf{X}, \mathbf{r}')$ for all sellers j, for any rankings $\mathbf{r}, \mathbf{r}'$ in which $r_j = r'_j$. We allow for winning payoffs to be interdependent, but the relative benefits to $r_j = 1$ are high enough to ignore opponent rankings. In this setting, we assume that $\max_{p_j} \pi_j(\mathbf{p}, \mathbf{X}, r_j)$ for any ranking $r_j > 1$ does not depend on $\mathbf{p}_{-j}$, and

$$\max_{p_j} \pi_j(\mathbf{p}, \mathbf{X}, r_j) \geqslant 0$$

for $r_j > 1$. Players ranked at $r_j > 1$ receive non-negative payoffs, but their maximum payoffs at these rankings do not depend on opponent prices nor rankings. $\max_{p_j} \pi_j(\mathbf{p}, \mathbf{X}, r_j)$, $r_j > 1$ is therefore treated as $\overline{\pi}_j$.

In this setting, for simplicity, $\pi_j^W(\mathbf{p})$ will denote $\pi_j(\mathbf{p}, \mathbf{X}, r_j)$ where $r_j = 1$, and $\pi_j^L(\mathbf{p})$ will denote $\pi_j(\mathbf{p}, \mathbf{X}, r_j)$ where $r_j > 1$, for the corresponding ranking $\mathbf{r}$ given f, which we will refer to as the *winner* and *loser* curves, respectively.

We will define

$$p_j^{f*} := \inf\{p_j | \pi_j^W(p_j, \underline{\mathbf{p}}_{-j}(p_j)) \geqslant \max_{p_j} \pi_j^L(p_j)\}.$$

p_j^{f*} is the minimum price that seller j is willing to obtain the first rank at, by beating her opponents by the smallest margin possible, given the scoring rule f. Note that by our assumptions on π_j, $\pi_j^W(\mathbf{p}) > \pi_j^L(\mathbf{p})$ for all $\mathbf{p}$ such that $\pi_j^W(\mathbf{p}) > \max_{p_j} \pi_j^L(p_j)$. Again we assume that $\max_{p_j} \pi_j^L(p_j)$ is sufficiently low such that p_j^{f*} is guaranteed to exist in P and $p_j^{f*} < \sup P$. We will denote the corresponding score $f_j(p_j^{f*})$ by f_j^*.

In this setting, we are interested in the locally envy-free profiles in which players maximize their profits conditional on their rankings.

Proposition 2. *In the locally envy-free profile with the highest total profit among all locally envy-free profiles, seller $i = \arg\max_j f_j^*$ wins. Her score satisfies $f_i(p_i) = \arg\max_{j \neq i} f_j^*$. All other sellers*

[3]Many platforms, including Amazon and Walmart, choose a ranking rule in which a single "featured offer" is awarded preferenced status.

$j \neq i$ score at prices that satisfy $\arg\max_{p_j} \pi_j^L(p_j)$.

Proof. First, we show that in any such locally envy-free profile, a seller with the highest f_i^* wins. Suppose there is $(p_i, \mathbf{p}_{-i})$ such that there is $j \neq i$ who obtains $\pi_j^W(p_j, \mathbf{p}_{-j})$ while i obtains $p_i = \arg\max_p \pi_i^L(p)$. Note that by our assumptions, any player who obtains the loser's curve and maximizes her profit conditional on losing, must obtain the loser's curve at $\max_{p_i} \pi_i^L(p_i)$. $f_j(p_j) \geqslant f_i^*$ since if not, i would prefer to obtain the winner's curve. But then $f_j(p_j) \geqslant f_i^* > f_j^*$, which contradicts the assumption that f_j^* is the maximum score that j is willing to obtain the winner's curve at.

Next we show that i must play p_i s.t $f_i(p_i) = \max_{j \neq i} f_j^*$. Suppose not. Suppose she plays p_i' s.t. $f_i(p_i') > \max_{j \neq i} f_j^*$. $\pi_j^W(p_j, \mathbf{p}_{-j}) < \max_{p_j} \pi_j^L(p_j)$ for any p_j s.t. $f_j(p_j) > f_i(p_i) = \max_{j \neq i} f_j^*$. Since $p_i > p_i'$ and payoffs on the winner's curve are monotonic, i can strictly improve her payoff by playing p_i, and still obtain the winner's curve.

Finally, we show that i cannot play any p_i s.t. $f_i(p_i) < \max_{j \neq i} f_j^*$. If so, then $\pi_j^W(p_j', \mathbf{p}_{-j}) > \arg\max_{p_j} \pi_j^L(p_j)$ for p_j' such that $f_j^* > f_j(p_j') > f_i(p_i)$ and hence $(p_i, \mathbf{p}_{-i})$ cannot be a locally envy-free profile. $\square$

Clearly, when seller $i = \arg\max_j f_j^*$ is unique, the proposed profile is unique, otherwise there is a set of possible winners. Note that the winner's score coincides with the winner's equilibrium transfer in a corresponding second price auction.

In this single-winner setting, the game is reminiscent of static Bertrand competition, where the extent that a player must undercut an opponent's price to obtain the winner's curve depends on the scoring rule f. In the extreme case, where $\pi_j^L(\mathbf{p}) = 0$ for all $j \in J$ and $\mathbf{p}$ and winner's curve payoffs are given by

$$\pi_j^W(\mathbf{p}) = D_j^W(p_j)(p_j - c_j)$$

for a demand function D_j^W such that π_j^W is non-decreasing in p_j, f_j^* is obtained at $p_j = c_j$ for all j.

When the scoring rule f is linear in characteristics, $f_j(p_j) = X_j'\hat{\beta} - p_j$, the winner i prices at

$$p_i = c_j + X_i'\hat{\beta} - X_j'\hat{\beta}$$

for $j = \arg\max_{j \neq i} f_j^*$. Player i prices at the minimum price that awards her the winner's curve, conditional on c_j. Therefore her markup is proportional to her cost advantage plus the extent to which her characteristics are favorable relative to her next best opponent under the scoring rule.

1.2.2 Defining Consumer Optimality

Following Reimers and Waldfogel (2023), we consider a setting in which consumers i choose among offers j and an outside option that yields zero utility, according to

$$u_{ij} = X_j'\beta - p_j + \lambda_r \mathbb{1}(r_j = r) + \epsilon_{ij}$$

where ϵ_{ij} is distributed Type 1 extreme value and X_j denotes a vector of offer j's characteristics, which are observable to both consumers and to the platform. λ_r is the effect of rank r on sales, where we assume that γ_r is strictly decreasing in r and $\lambda_r \geqslant 0$ for all $r \leqslant N$. λ_r may be interpreted as the additional convenience benefit, or reduction in search costs that a consumer enjoys from purchasing an offer ranked at r.

Offer j's corresponding market share is given by

$$s_j = \frac{\exp\Big(X_j'\beta - p_j + \lambda_r \mathbb{1}(r_j = r)\Big)}{1 + \sum_{k \in J} \exp\big(X_k'\beta - p_k + \lambda_r \mathbb{1}(r_k = r)\big)}$$

Given the above,

$$\mathbb{E}[\text{CS}] = \frac{M}{\alpha} \log\left(1 + \sum_{j \in J} \exp\Big(X_j'\beta - p_j\Big) \exp\big(\lambda_r \mathbb{1}(r_j = r)\big)\right)$$

for market size M, and α, the marginal utility of income. As in Reimers and Waldfogel (2023), since $\exp\big(\lambda_r \mathbb{1}(r_j = r)\big)$ is strictly decreasing in r, expected consumer surplus, given any *fixed* vector of characteristics and prices is maximized by ranking offers in order of expected rank-independent mean utility, $X_j'\beta - p_j$, where higher mean utility offers are awarded lower (better) ranks.

To study the assignment rule that maximizes expected consumer surplus when prices are *endogenous* given the scoring rule f, and in particular, coincide with the profit-maximizing locally envy-free profile of the perfect information simultaneous move game, we appeal to the setting in which $\lambda_r = 0$ for all $r > 1$, and where we assume that λ_1 is sufficiently high such that

$$\mathbb{E}[\text{CS}] \approx \frac{M}{\alpha} \log\left(1 + \exp\left(X_j'\beta - p_j\right) \exp\left(\lambda_1 \mathbb{1}(r_j = 1)\right)\right).^4$$

In this setting, we can reduce the platform's objective to maximizing

$$X_j'\beta - p_j$$

for the first-ranked offer only.

For the remainder, we will consider the single-winner invariant case studied in section 1.2.1. Here we will assume that γ_1 being sufficiently high implies that p_j^{f*} does not depend on f. In other words, the winner's curve is invariant enough to opponent prices such that differences in $\underline{\mathbf{p}}_{-j}(p_j)$ across scoring rules f do not meaningfully shift p_j^*.

For a brief discussion on consumer welfare under the special independent profits case studied in section 1.2.1, see A.2.1.

1.2.3 The Consumer-Optimal Static Scoring Rule

Proposition 3. *When prices correspond with the profit-maximizing locally envy-free profile, the designer maximizes expected consumer-welfare by choosing a scoring rule f that minimizes*

$$\min_f (p_i - p_i^*)$$

where $i = \arg\max_k f_k^* = \arg\max_k X_k'\beta - p_k^*$ *and* p_i *satisfies* $f_i(p_i) = \arg\max_{j \neq i} f_j^*$.

The optimal rule ensures the lowest possible equilibrium price markup subject to the

[4] This is analogous to a setting in which consumers only choose between the first-ranked offer and an outside option, or where the convenience benefits of choosing the first-ranked offer outweigh any disutility from increased prices or idiosyncratic preferences for offers ranked above the first.

constraint that the seller awarded featured status coincides with that which obtains the maximum possible expected consumer-welfare when she prices at her reservation price.

Proof. For any f, equilibrium welfare is proportional to

$$X_i'\beta - p_i$$

for the offer i for whom $f_i^* \geqslant f_j^*$ holds for all $j \neq i$ and p_i satisfies, $f_i(p_i) = f_j^*$ for offer $j = \arg\max_{k\neq i} f_k^*$.

When restricted to the class of *linear* scoring rules for which there exists $\hat{\beta}$ such that $f(p_j) = X_j'\hat{\beta} - p_j$, equilibrium welfare can be written as

$$X_i'\beta - (X_i'\hat{\beta} - X_j'\hat{\beta} + p_j^*)$$

and therefore under $\hat{\beta} = \beta$, is given by

$$X_j'\beta - p_j^*$$

for $j = \arg\max_{k\neq i} X_k'\beta - p_k^*$.

For any f for the which the featured offer $k \neq \arg\max_k X_k'\beta - p_k^*$, welfare is proportional to

$$X_k'\beta - p_k$$

where since $p_k \geqslant p_k^*$, $X_k'\beta - p_k \leqslant X_j'\beta - p_j^*$ for the offer $j = \arg\max_{k\neq i} X_k'\beta - p_k^*$. Therefore, the designer can always ensure weakly higher expected consumer welfare by choosing $f(p_j) = X_j'\beta - p_j$ relative to *any* other scoring rule that results in an assignment of featured status to any other seller other than that chosen under $f(p_j) = X_j'\beta - p_j$.

Therefore, the optimal f ensures that the featured offer coincides with that selected under $f(p_j) = X_j'\beta - p_j$, and hence equilibrium welfare is maximized under the f that minimizes

$$p_i - p_i^*,\ p_i : f_i(p_i) = \arg\max_{j\neq i} f_j(p_j^*)$$

while ensuring that $i = \arg\max_k X_k'\beta - p_k^*$ remains the featured offer.

□

Note that the optimal f depends on both vectors $\mathbf{X}$ and $\mathbf{p}^*$ in any given market, and hence requires the platform to perfectly observe all payoff-relevant information. This corresponds closely to a setting in which the designer can extract full or near-full surplus in any market, but in which she must be able to set a discriminatory scoring rule *across* marketplaces. In the extreme setting, when the designer has *no* information on the distribution F of $\mathbf{p}^*$, or even $\mathbf{X}$, she can always ensure a minimum level of welfare equal to the consumer welfare of the second-highest preferred offer when she prices at her reservation price, by choosing a linear scoring rule, $f_j(p_j) = X_j'\hat{\beta} - p_j$ and setting $\hat{\beta} = \beta$. Choosing any $\hat{\beta} \neq \beta$ risks awarding the first rank to a sub-optimal offer, or enabling increased markups.

If the designer has information over the distribution F of $\mathbf{p}^*$, where F may be conditional on $\mathbf{X}$, then, she can maximize *ex-ante* expected consumer welfare by solving for

$$\arg\min_{f} p_i - \mathbb{E}[p_i^*|\mathbf{X}]$$

such that

$$f_i(\mathbb{E}[p_i^*|\mathbf{X}]) = \arg\max_{j \neq i} f_j(\mathbb{E}[p_j^*|\mathbf{X}])$$

and

$$\arg\max_{k} f_k(\mathbb{E}[p_k^*|\mathbf{X}]) = \arg\max_{k} X_k'\beta - \mathbb{E}[p_k^*|\mathbf{X}].$$

The optimal f depends on the observed values of $\mathbf{X}$, even for any fixed distribution F. Note however, under the assumption that sellers observe the realized vector $\mathbf{p}^*$, the above is not guaranteed to coincide with the scoring rule that maximizes expected consumer welfare given the realized vector $\mathbf{p}^*$. The cost to consumer welfare from any other preferencing rule is directly proportional to increased markups due to either over-rewarding the optimal offer's characteristics relative to others' or preferencing a sub-optimal offer to an extent that allows him to price above his reservation price. When the platform is the most preferred offer, any self-preferencing comes at a cost to consumers through increased prices.

Equilibrium welfare under the class of linear scoring rules in which $\hat{\beta} = \beta$ coincides with welfare under the corresponding designer's problem, where the designer is restricted to assignment rules that satisfy ex-post incentive compatibility and where seller's types are given by the welfare that they generate by pricing at their reservation price. The payment rule in this mechanism corresponds with equilibrium transfers in a second price auction where players have incomplete information over opponent types, see A.2.2.

Considering welfare relative to the linear scoring rule in which $\hat{\beta} = \beta$ serves as a useful benchmark to study the extent to which platform's observed scoring rules deviate from a consumer-optimal rule, given that under our assumptions, implementing this rule requires no information on the sellers. Moreover, setting a discriminatory optimal rule *across* marketplaces may be too costly for even a platform that has near perfect information on all sellers.

Examining welfare relative to this rule will be the basis of our empirical test on the Amazon platform, which we describe in detail in the sections that follow.

1.3 Data

Our data were provided by a high-volume third-party seller who sold on Amazon from 2016-2022, after which their business dissolved. This seller sold over 1000 different products on Amazon, mostly in the consumer cyclical and staple categories including Beauty, Health and Household, Arts and Crafts, Office Products, Grocery, Pet, and Baby product categories. Their business operated only in the United States market. The data that we acquired include all events from Amazon's API which sends timestamped updates to sellers given any event change on a product page on which our seller sold. These events include a price, shipping time, fulfillment method, or rating or review update from any active seller on the product page, or a change in Buy Box ownership. Our data also includes all order information for our seller, product page views on a daily basis, cost data for our seller, and some keyword search information for certain product pages.

Additional historical information on product page rankings under keyword searches were obtained from a third-party company, Helium 10. These data are recorded for the product pages that Helium 10's customers track information for over time, where rankings are obtained by recording various keyword search result pages.

Due to the scale of our data, we select a random sample of a maximum of 20 000 product-page days in each product category, from December 2020 to December 2022 during which to study the Buy Box assignment rule[5]. Tables 1.1 and 1.2 present the summary statistics over these random samples within each product category. As reported in previous studies, the Buy Box is not the lowest price approximately 50 % of the time in many product categories and Amazon is awarded the Buy Box when it is an available offer, more than 50 % of the time in all product categories. Most notably, the Buy Box winner is "objectively worse" a substantial fraction of the time in all product categories. We define an offer as objectively worse if there is another offer on the page that is not identical in terms of observables and which satisfies all of the following; a weakly lower total price, weakly lower shipping time, weakly higher FBA status, weakly higher Prime status, weakly higher feedback or weakly higher reviews. We discuss this phenomenon, and further criterion that we use to select product page-days for estimation in the sections that follow.

Additionally, while the Buy Box is reported to have a substantial, positive effect on sales, many sales in our data occur off of the Buy Box. Approximately 37 % of our seller's sales are attributed to our seller's non-Buy Box winning offers, when they have another offer available. We discuss how we select and use variation in our seller's data to identify parameters of interest in demand estimation in the sections that follow.

Table 1.1: *Buy Box winner descriptive statistics by product category (part 1).*

	Beauty	Health & Beauty	Office Products
Average distinct Buy Box winners per day	1.43	1.42	1.43
Average distinct offers per choice set	3.92	3.83	3.53
Fraction of choice sets including Amazon	0.09	0.08	0.12
Fraction of Amazon Buy Box wins	0.51	0.64	0.65
Fraction Buy Box offer is lowest price	0.56	0.54	0.54
Fraction objectively worse Buy Box winners	0.17	0.16	0.15

[5] We use data starting in 2020 as many variables necessary for demand estimation are only available from 2020 onwards.

Table 1.2: *Buy Box winner descriptive statistics by product category (part 2).*

	Arts & Crafts	Baby Products	Grocery	Pet Products
Average distinct Buy Box winners per day	1.60	1.50	1.55	1.43
Average distinct offers per choice set	4.97	3.69	4.10	4.09
Fraction of choice sets including Amazon	0.15	0.10	0.02	0.12
Fraction of Amazon Buy Box wins	0.85	0.43	0.63	0.39
Fraction Buy Box offer is lowest price	0.62	0.55	0.63	0.62
Fraction objectively worse Buy Box winners	0.24	0.14	0.18	0.20

1.4 Model

1.4.1 Product Page Preferencing: The Buy Box Assignment Rule

Industry knowledge and previous work suggest that Amazon's scoring rule can be modeled as a multinomial logit, specific to a product category, where categories may include Beauty and Personal Care, Health and Household, Pet Supplies, Office Supplies, Books and Music (Raval, 2022).

For each active offer j on page m in product category c, where we suppress m, c from notation, Amazon assigns a score to offer j according to

$$f_j = X_j'\psi + \gamma \log p_j + \epsilon_j$$

for all eligible offers j, where ϵ_j is distributed Type 1 extreme value and represents characteristics that are observed by Amazon but not by the econometrician. For example, these might include more detailed information on feedback variables, such as how quickly the seller responds to return requests. p_j denotes offer j's total price which is equal to the offer's unit price plus shipping price. In most of our data, shipping prices are zero, and therefore we focus our estimation on total prices. X_j is a vector of observable characteristics which includes; shipping time in days, as well as dummy variables for whether the offer is sold by a non-Amazon third-party but Fulfilled-by-Amazon (FBA), or Fulfilled-by-Merchant (FBM). Since the correlation between Prime offers and non-Amazon FBA offers is close to one in our data, we do not include an indicator for whether the offer is Prime eligible in our regression. Therefore the non-Amazon FBA indicator measures the average difference in log odds between an FBA and Prime offer being assigned the Buy Box relative to an Amazon offer, conditional on observables, while the FBM indicator measures the average difference in log odds of an FBM offer being assigned the Buy Box relative to an Amazon offer conditional on observables. Additionally, since Prime status is highly correlated with shipping time, differences in Prime status among FBM offers are mostly accounted for by shipping time. Given that lifetime feedback and rating data for Amazon were not reported after 2018, we impute a lifetime feedback count of 5 million for Amazon and a perfect rating of 100[6]. Following Raval (2022), in order to measure the log odds of a non-Amazon offer being awarded the Buy Box relative to an Amazon offer, we include variables defined as the difference in log feedback between an offer's seller and a log feedback of 5 million, and the difference in rating between an offer's seller and 100. If higher feedback and ratings imply a higher score, the coefficients on these variables should be positive. We exclude the condition of the product (whether it is new, used-good, etc.) as all offers on the product pages on which our seller sold are in new condition.

Amazon assigns featured status to an offer, given the scoring rule, according to

$$\mathrm{BB}_j = 1 \iff f_j(p_j, X_j) \geqslant f_k(p_k, X_k), \text{ for all } k \in m$$

where BB_j denotes the outcome indicator variable for whether offer j was assigned the Buy Box. Only one offer can be assigned the Buy Box at a time, and we assume that ties are broken at random with equal probability. If all scores are sufficiently low, the scoring rule may not award the Buy Box to any offer j. Offers may be ineligible for the Buy Box if some of their metrics do not meet threshold criterion, in particular, if its price exceeds a *competitive price threshold*, known as CPT[7]. Amazon might also consider an offer ineligible if it does not have sufficiently high ratings or reviews.

[6] Amazon's average rating in 2018 was 91 %. The maximum lifetime feedback of any seller in our data is 4.3 million.

[7] These price thresholds may correspond with the current lowest price of the item sold on other platforms, for example, on Walmart.

1.4.2 Consumer Demand

To model consumer demand, markets are given by product pages on which an *identical* product is sold by various sellers, whose offers may differ on characteristics such as shipping time or product condition. On product page m we denote offer j, on day d, and time t. We assume that consumer preferences for products $m \in c$, where c is a product category, such as Health and Beauty, or Office supplies, are governed by the same parameters. However, the only products in an individual's choice set lie on a given product page m.

A consumer i arrives on product page m exogenously and chooses between all active offers j on page m at time t on day d, and a fixed outside option that yields zero utility, according to the following utility function:

$$u_{imjdt} = X'_{jmdt}\beta - \alpha \log p_{jdmt} + \phi_{md} + \xi_{jmd} + \epsilon_{imdt}$$

where p_{jmdt} is the total price of offer j, equal to the sum of the unit and shipping price, X_{jmdt} are observable product characteristics, ϕ_{md} is a product page-day fixed effect, ξ_{jmd} is unobservable offer product quality, and ϵ_{imdt} is distributed Type 1 extreme value. As in our Buy Box scoring rule model, the vector X_{jmdt} includes; shipping time, the difference in log lifetime feedback relative to Amazon, the difference in ratings relative to Amazon, as well as indicator variables for whether the offer is a third-party offer that is FBA or FBM. The vector X_{jmdt} *additionally* includes whether or not offer j on page m owns the Buy Box on day d at time t. The coefficient on the Buy Box indicator measures the difference in log-odds of a consumer choosing the Buy Box offer relative to a non-Buy Box offer conditional on observables. As in our theoretical model, we view the presence of the Buy Box not as limiting the choice sets of consumers, but rather as providing some convenience benefit that may be interpreted as a reduction in search costs, wherein that benefit does not depend on the offer assigned the Buy Box nor it's characteristics, including price. Since the correlation between FBA and Prime in our data is approximately one, we omit the Prime indicator from the model, and we exclude the condition of the product (whether it is new, used-good, etc.) as all offers on the product pages on which our seller sold are in new condition.

Consumers choose the offer among all active offers that maximizes their utility and therefore offer j's share on page m, on day d at time t among all active offers $k \in m$ is given by,

$$s_{jmdt} = \frac{\exp\left(X'_{jmdt}\beta - \alpha \log p_{jmdt} + \phi_{md} + \xi_{jmd}\right)}{1 + \sum_{k \in m} \exp\left(X'_{kmdt}\beta - \alpha \log p_{kmdt} + \phi_{md} + \xi_{kmd}\right)}$$

We note that we allow unobservable offer quality ξ_{jmd} to vary on the daily level d but not across time intervals t within any day d for reasons that we discuss more in detail in the estimation section.

Additionally, since we only observe product page views on the daily level, and therefore can only estimate the total shares for offers other than our seller's on a given day at the aggregate daily level, we define *weights* w_{dt}, that represent the fraction of total consumers who visit the page on day d in time period t, such that $s_{jmd} = \sum_{t \in d} w_{dt} s_{jmdt}$ where $\sum_{t \in d} w_{dt} = 1$, and $s_{jmdt} = \frac{q_{jmdt}}{w_{dt} M_{md}}$, where q_{jmdt} denotes offer j's total quantity of sales on day d at time t, and M_{md} denotes the total number of page visitors to page m, or the market size, on day d.

Weights, w_{dt} follow the parameterization

$$w_{dt}(\gamma) = \frac{\exp\left(\tilde{t}'\gamma\right)}{\sum_{\tilde{\tau} \in d} \exp(\tilde{\tau}'\gamma)}$$

$\tilde{\tau} \in \{0, 1, \ldots, 6, 7\}$ for each possible three-hour interval in d, where $\tilde{t}'$ is a second-order polynomial in $\tilde{t}$.

1.4.3 Keyword Search

In the consumer demand model, we assume that consumers arrive exogenously on product pages. In reality, consumers arrive on product pages via keyword searches, where for any given keyword search, Amazon deploys an algorithm that determines how product pages that appear in the keyword search results are ranked among one another, where pages with lower ranks appear earlier.

Any consumer-optimal preferencing rule therefore, must jointly determine both a keyword search ranking rule as well as a product-page preferencing rule. Determining such an optimal rule requires knowledge of consumer search patterns, namely, how consumers behave when they leave keyword search pages or product pages. Since our data to do not allow us to observe these search patterns, we will consider consumer-optimal keyword search ranking rules that are agnostic about consumers' outside options, which will nevertheless be a useful first step in studying optimal preferencing rules and the costs of sub-optimal preferencing rules in these environments.

Although keyword search rankings depend on prices on the product page, we will maintain our previous assumption that sellers consider the number of product page visitors to be exogenous and therefore equilibrium prices coincide with those studied above. This assumption is not unrealistic, as the platform's keyword search ranking algorithm may be highly complicated and difficult for firms to additionally optimize prices over. Instead, firms attempt to increase the number of consumers who visit their product pages while they are featured by purchasing advertising, including sponsored keyword search rankings. As previously noted, industry knowledge suggests that when an individual seller on a given product page purchases advertising, that seller is automatically awarded the Buy Box when a consumer visits the produce page after clicking on the sellers' sponsored keyword search result. Therefore, any welfare distortions that arise from advertising are due not only to the potentially sub-optimal assignment of the Buy Box, but also to the distortions that arise from sub-optimal keyword search ranks which influence the probabilities of consumers visiting various product pages.

Consider a setting in which consumers i draw keywords k exogenously and choose among product pages, denoted by m, that are ranked under keyword search k and a fixed outside option yielding zero utility, according to

$$v_{imk} = X_m'\zeta - \eta \log p_m + \kappa r_{mk} + \xi_{mk} + \epsilon_{imk}$$

where ϵ_{imk} is distribution Type 1 extreme value. X_m is a vector of observables for the featured offer *only* on page m at the time of search including; offer unit price, shipping time, whether the offer is Prime eligible, and rating metrics. We note that these observable characteristics are a subset of the product's observables once a consumer has arrived on the product page m. r_{mk} denotes page m's ranking under keyword search page k, and ξ_{mk} denotes the perceived quality of product page m under keyword search k, which we allow to vary for a given product page under the choice sets governed by different keyword searches, and which may differ from quality terms in the consumer demand model, conditional on landing on product page m.

Page m's corresponding click through rate, or the fraction of consumers who click on page m under keyword search k, is given by

$$\text{CTR}_{mk} = \frac{\exp(X_m'\zeta - \eta logp_m + \kappa r_{mk} + \xi_{mk})}{1 + \sum_{m'k} \exp\left(X_{m'}'\zeta - \eta logp_{m'} + \kappa r_{m'k} + \xi_{m'k}\right)}$$

for all product pages m' ranked under keyword search k.

Consumer Welfare

A true measure of consumer welfare would require understanding how consumers behave upon choosing the outside option once they've landed on a product page and on a search results page. Under the assumption that consumers derive no inherent utility from searching or from clicking on a page, but only from final purchase decisions, and given the agnostic assumption about consumers' outside option, then the expected consumer welfare conditional on arriving on a product page m coincides with coincides with that from the demand model

$$\mathbb{E}[\text{CS}_m] \propto \log\left(1 + \sum_{j \in m} \exp\left(X_j'\beta - \alpha \log p_j + \phi_m + \xi_{jm}\right)\right)$$

Given our assumption about the outside option, choice probabilities conditional on landing on a page do not depend on *how* the consumer arrived on the page, and $\mathbb{E}[\text{CS}_m]$ given keyword search k, is the same for all keyword searches k in which m appears. Indeed, our data suggest that this assumption is reasonable. The mean standard deviation in observed conversion probabilities conditional on landing on a given page from a given keyword search k is .008, for pages on which we observe conversion probabilities from multiple keyword searches.

This implies that a platform that optimizes for expected welfare that coincides with $\mathbb{E}[\text{CS}_m]$ must choose a ranking rule for any keyword search k that maximizes

$$\sum_{m \in k} \text{CTR}_{mk}\, \mathbb{E}[\text{CS}_m]$$

Clearly, any ranking rule that optimizes the above objective function must imply the highest CTR_{mk} on the highest $\mathbb{E}[\text{CS}_m]$ pages. If rank-independent expected consumer utility on page m under keyword search k is consistent with conversions on product pages, in other words, if

$$X_m' \zeta - \eta \log(p_m) + \xi_{mk}$$

is always increasing in $\mathbb{E}[\text{CS}_m]$ where $\mathbb{E}[\text{CS}_m]$ is given under the consumer-optimal Buy Box assignment rule, then for any $\kappa < 0$, clearly the platform cannot do any better than by employing the consumer-optimal Buy Box assignment rule and ranking product pages in increasing order of $\mathbb{E}[\text{CS}_m]$.

On the other hand, if rank-independent expected consumer utility on page m under keyword search k is not consistent with conversions, by employing the consumer-optimal Buy Box rule, the platform still ensures the highest possible consumer utility on any given page m. Conditional on the resulting characteristics and prices of the featured offer, the platform can always rank products in such a way so as to ensure the highest resulting CTR on the expected consumer surplus pages. This is preferable to any other preferencing rule, assuming that any losses in total conversions that arise from the above are strictly less than any that result from having a sub-optimal preferenced offer on a product page.

Following Reimers and Waldfogel (2023), one could test for the presence of self-preferencing in keyword search rankings by measuring the coefficients on FBA and FBM indicator variables in a regression of expected consumer surplus on the product-page level on observed ranks. Unfortunately, our data are too sparse for this regression, and contain very few product pages on which Amazon competes. However, in the absence of this test, one can provide evidence for suboptimal preferencing without estimating an explicit ranking algorithm, by examining the extent to which observed click-through-rates that are implied by ranks are systematically higher for higher expected consumer surplus pages.

Our data suggest that the current keyword ranking rule does not optimize for end-conversions, nor expected consumer surplus. For all product-page pairs ranked for a given keyword, a higher end-conversion page is not awarded a higher organic rank 40.25 % of the time, and a higher end-conversion page does not have a higher click-through-rate 32.25 % of the time.

Table 1.3 summarizes additional regressions of expected consumer surplus on click-through-rates and average organic rank, on the few product page-keyword pairs on which we are able to estimate expected consumer surplus.

Table 1.3: *Expected consumer surplus and keyword rankings estimation results.*

	CTR	Average Organic Rank
Expected Consumer Surplus	-0.03^{***}	0.07
	(0.01)	(0.07)
R-squared	0.23	0.03
Observations	37	37

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Standard errors in parentheses

Expected consumer surplus does not appear to be positively correlated with implied click-through-rates, nor negatively correlated with average organic ranks, as would be expected if higher welfare pages were awarded lower organic ranks.

1.5 Estimation

1.5.1 Product Page Preferencing

Contrary to other authors who estimate the scoring rule using data on all non-used offers available on a given product page (Raval, 2022; Musolff, 2022), we eliminate offers that are likely to be ineligible. We believe that Amazon first selects offers that are eligible for the Buy Box, and then chooses the highest scoring offer according to the rule prescribed above among these eligible offers. Therefore, including ineligible offers in the estimation may bias estimates and potentially over-estimate the presence of self-preferencing. Since we do not observe Amazon's eligibility criterion, we eliminate offers whose prices are greater than the competitive price threshold, as well as those whose lifetime feedback or ratings are below the 25th percentile of all lifetime feedback or ratings. Additionally, we eliminate offers that were never assigned the Buy Box during the month that they were active, and do not estimate the scoring rule on any set of offers in which no offer was assigned the Buy Box. For this reason, we do not include a constant in our multinomial logit specification.

Industry information also suggests that in recent years, Amazon has assigned the Buy Box to offers that have paid for advertising, when a consumer arrives on the product page after clicking that offer's advertisement. Indeed, given advertising data from our seller, on the majority of the product-page days on which our seller was assigned the Buy Box and was an objectively worse offer than a competing offer, these product pages were visited directly from sponsored ads that were purchased by our seller.

Since we do not have advertising data on all sellers, we cannot eliminate all instances of Buy Box assignments that are due to advertising. To mitigate the effects of this on our estimates, we do not estimate a scoring rule on any product page during which an objectively worse offer was assigned the Buy Box. Therefore, the assignment rule that we estimate is the average rule over all assignments due to Amazon's true scoring rule and any assignments due to advertising or ineligibility that we were not able to eliminate.

In order to ensure that the underlying sample identifies the parameters of interest, we sub-select a sample of product-page days on which both an Amazon and a non-Amazon FBA offer were awarded the Buy Box, both an Amazon and an FBM offer were awarded the Buy Box, both a non-Amazon FBA and an FBM offer were awarded the Buy Box, and multiple FBM offers with various shipping times were awarded the Buy Box, since most, but not all of the variation in shipping time is across FBM offers.

We estimate the Buy Box assignment rule via Maximum Likelihood Estimation on the product category level, for the Beauty and Personal Care and Health and Household product categories only, as the resulting sub-samples in these categories are large enough to identify our parameters of interest.

1.5.2 Consumer Demand

A challenge with estimating demand in this setting is that we only observe sales for some offers j on any given page m, but not all. Nevertheless, we impute total product page sales on the daily level using sales ranks, a measure of aggregate sales on the product-page level, in a procedure described in A.1. Then, given that we observe total page visitors M_{md} for a given page m on day d, s_{0md}, the corresponding outsize share, and residual shares for the remaining offers are determined.

While we do not exactly observe the weights $w_{dt}(\gamma)$ that represent the fraction of total consumers who visit the page on day d in time period t, industry knowledge suggests that w_{dt} follows a bell curve over a daily period in which pages receive the majority of page views during the late morning into the afternoon, falling substantially during the late evening hours. Industry knowledge also suggests that the fraction of consumers who shop during certain time intervals of the day is very consistent over time and across product pages, and therefore it is safe to assume that $w_{mdt} = w_{dt}$. Proprietary information that we acquired from our data provider as well as sales data over various product pages during days on which prices had very low variation, informed our estimation of γ from the data. The distribution implied by these resulting weights over each three-hour daily interval is shown in Figure 1.3 where we note that the additional visitors between midnight and 3am relative to 3 to 6am are most likely due to shoppers on the west coast.

Given $w_{dt}(\gamma)$, we estimate parameters $\theta = (\beta, \alpha)$ using the following procedure following Berry et al. (1995).

For any guess of $\hat{\theta}$:

1. Given observed shares denoted by s^o, solve for $\xi_{jmd} + \phi_{md}, \xi_{fmd} + \phi_{md}$ such that

$$|s^o - s(\hat{\theta}, \gamma, \xi, \phi)| \leqslant \text{tol}$$

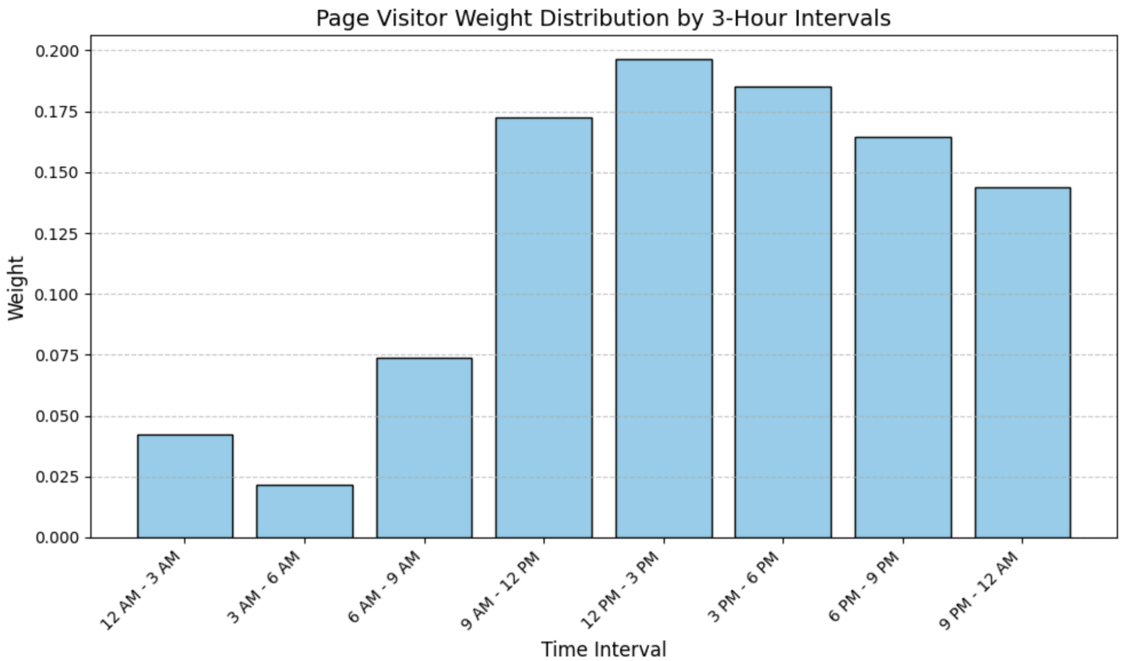

Figure 1.3: *Page visitor daily weight distribution implied by γ, Eastern Time.*

using the BLP contraction following Conlon and Gortmaker (2020)[8] where s is given by

$$s_{jmd} = \sum_t w_{dt}(\gamma) s_{jmdt}(\hat{\theta}, \xi, \phi)$$

for offers j for which s_{jmdt} is observed and

$$s_{fmd} = \sum_{k \in f} \sum_t w_{dt}(\gamma) s_{kmdt}(\hat{\theta}, \xi, \phi)$$

for all offers k in the subset for which s_{kmdt} is not observed, denoted by f.

2. Form sample moments, which are demeaned on the t-level in order to absorb product page-day level fixed effects, ϕ_{md}

$$g(\hat{\theta}) = \left[\frac{1}{N} \sum_m \sum_{j \in m} \sum_d \sum_t (X_{jmdt} - \overline{X}_{mdt})(\xi_{jmd}(\hat{\theta}, \gamma) - \overline{\xi}_{md}(\hat{\theta}, \gamma)) \right]$$

[8] For a proof of the BLP contraction in this setting, refer to Elliott et al. (2023), appendix B.1.2, where here the $\tilde{u}_{jmdt}$ are given by $X^d_{jmdt} = X_{jmdt} - \overline{X}_{jmd}$ and the distribution of the random coefficients $u_{jmdt} = X_{jmdt}$ is given by the distribution implied by the weights w_{dt}.

and the GMM objective function

$$g(\hat{\theta})Wg(\hat{\theta})'$$

Update $\hat{\theta}$ to solve for the value of θ that minimizes the GMM objective function.

Importantly, we assume a common value ξ_{fmd} for all $k \in f$, offers for which we do not directly observe daily sales. Given that each seller sells the same underlying product, additional observables control for most differences in product quality across sellers on the page, including those arising from varying fulfillment methods or seller reviews. Therefore, measures in unobservable offer quality should be very similar across offers on a given product page. While we nonetheless allow quality terms to vary between the seller for whom we observe sales and those for whom we do not, we are not substantially harming our estimation by assuming a common ξ_{fmd} across sellers $k \in f$. In this setting, ϕ_{md} represents the difference in perceived product quality between the product on page m and the outside option, which may indeed vary across days d, and thus is important for estimating the fraction of consumers who choose the outside option in a counterfactual setting.

We limit demand estimation to product pages m on days d in which there is substantial price variation arising from competition for the Buy Box.[9] Indeed, as previous literature documents, when sellers compete for the Buy Box using pricing algorithms, pricing patterns closely resemble Edgeworth cycles, and result in high-frequency variation in not only which seller is assigned the Buy Box, but also at which prices (Musolff, 2022) (see also chapter 2). Given that the pricing algorithms that determine these prices are set well in advance of individual daily-level demand shocks, pricing patterns on these product pages should be uncorrelated with the perceived unobserved product quality on a daily level. This level of variation allows us to identify not only price elasticity, but also the coefficient on the Buy Box indicator variable without additional instruments for prices. The variation in t-level characteristics, including prices, relative to daily average characteristics, are therefore

[9]We define substantial variation in which the standard deviation of Buy Box prices is at least two times the average Buy Box price, each seller that wins the Buy Box has at least four price changes on the day, and there are at least two different Buy Box winners.

uncorrelated with individual offer-level quality terms. Demeaning the moments on the time-level rather than the daily level, critically preserves this variation in prices across time.

Finally, while total daily page visitors M_{md} may depend on Buy Box prices, prices on product pages are likely to be unresponsive to these demand shifts. This is due to the fact that these prices arise from competing algorithms that are set in advance of daily shifts, and sellers do not optimize prices to obtain higher page visitors, given that visitors arise from complicated keyword search ranking algorithms. Moreover, the pricing patterns that we observe on these pages appear to be relatively invariant to time shifts throughout the day, and consumers likely do not adjust shopping patterns due to daily-level fluctuations in prices on a given page. It is therefore safe for us to restrict ourselves to a specification in which w_{dt} does not depend on prices.

We carefully select a sub-sample of product page-days during 2020-2022, that allows us to identify the demand parameters of interest. Fortunately, on a subset of product page-days, our seller has multiple active offers at given a time, including at least one FBA offer and at least one FBM offer. Among our sample of product page-days for which we observe page visitors, we sub-select pages on which; Amazon and a non-Amazon FBA offer from our seller are awarded the Buy Box for at least four hours during the day, both Amazon and an FBM offer from our seller are awarded the Buy Box for at least four hours during the day, our seller has at least one FBA offer and at least one FBM offer, our seller has at least two offers off of the Buy Box for at least three hours during the day, our seller has multiple FBM offers, one of which is awarded the Buy Box for at least 10 hours during the day and one of which is not awarded the Buy Box during at least 7 hours during the day, and lastly, days on which Amazon is not awarded the Buy Box for at least four hours of the day and one of our seller's offers is awarded the Buy Box for at least 7 hours during the day.

The variation in our data allows us to identify the FBA-non-Amazon and FBM coefficients relative to Amazon, although there are very few days in the data where an FBM offer competes for the Buy Box with an Amazon offer. Additionally, this allows us to identify preferences for shipping time, since FBM offers from our seller differ only on this dimension, apart from Prime. Given that nearly all FBA and Amazon offers are Prime, and since all Prime offers tend to have a shipping time of two days or less, there is not sufficient variation to distinguish differences in shipping time from Prime status. Finally, having multiple FBM offers both on and off of the Buy Box allows us to identify the Buy Box coefficient conditional on shipping time. Due to limited samples in other product categories that satisfy these criterion, we only estimate demand in the Health and Household product category.

1.5.3 Product Page Self-Preferencing Test

In the theoretical model, we motivated measuring the extent to which the platform's preferencing rule differs from the linear rule whose coefficients on observables corresponds with the coefficients of consumer's rank-independent demand parameters. We can directly assess this difference by examining the coefficients from the above estimation procedures. Additionally, by implementing a test similar to that of Reimers and Waldfogel (2023) conditional on observables test, we can measure the explanatory power of preferencing indicators on Buy Box assignment, conditional on rank-independent purchase probabilities, which account for consumer preferences over unobservable product quality.[10] That is, we can measure the extent to which self-preferencing explains Buy Box assignments relative to the rule that would maximize consumer purchase probabilities at any instant of time. By using purchase probabilities, we can separately estimate the explanatory power of preferencing indicators from consumer's own preferences for Amazon, FBA, and FBM offers.

Recall that on any product page m expected consumer surplus is proportional to

$$\log\left(\sum_{j\in m}\exp\left(X_{jm}'\beta-\alpha\log p_{jm}+\phi_m+\xi_{jm}\right)+1\right)$$

given offers $j\in m$, which for any fixed vector of prices and characteristics, is maximized when featured status is awarded to the offer with highest preferenced-independent mean-utility

[10]The distinction of the test proposed here and Reimers and Waldfogel (2023)'s conditional-on-observables (COO) test is that our test estimates assignments conditional on mean consumer utility, including unobserved product quality. This mitigates some of the challenges of using observable characteristics alone to characterize mean utility, as explained in Reimers and Waldfogel (2023), p.12.

offer, under our assumption that the log-odds of choosing the featured offer conditional on observables is the same for any possible featured offer j.

We test for self-preferencing by examining the extent to which preferenced-independent mean consumer utilities predict an offer's featured status. In particular, we estimate a model in which

$$V_j = s_j(X, p) + \delta \mathbb{1}_{\text{FBA} \setminus \text{AMZ}} + \rho \mathbb{1}_{\text{FBM}} + \epsilon_j$$

for all active offers j on a product page m, where ϵ_j is distributed Type 1 extreme value, and corresponding Buy Box assignments are given by $\text{BB}_j = 1 \iff V_j \geqslant V_k$ for all $k \in m$.

$$s_j(X, p) = \frac{\exp\left(X'_{jm}\beta - \alpha \log p_{jm} + \phi_m + \xi_{jm}\right)}{\sum_{k \in m} \exp\left(X'_{km}\beta - \alpha \log p_{km} + \phi_m + \xi_{km}\right)}$$

is offer j's choice probability among all active offer's on the same page, using the consumer demand parameters we estimated above, *excluding* the coefficient and indicator for whether the offer was awarded the Buy Box or not. These shares are strictly increasing in rank-independent mean utility.

As in our original estimation of the scoring rule, we include indicator variables, for whether the offer is FBA and not sold by Amazon, or FBM, to determine the extent of these indicators' explanatory power on Buy Box assignment, conditional on rank-independent choice probabilities. If Amazon awards the Buy Box according to the rule that maximizes consumer surplus, Buy Box assignments should be independent of these indicator variables conditional on unbiased choice probabilities. If the coefficients δ, ρ are both negative, this suggests that Amazon engages in self-preferencing when assigning the Buy Box.

We remind the reader that given that we cannot fully control for assignments that arise from advertising, if these assignments occur for any non-Amazon offers within our data, than these coefficients will be downward biased relative to the true Buy Box assignment rule.

1.5.4 Keyword Search

The true weekly click-through-rate, CTR_{mkw}, satisfies

$$\text{CTR}_{mkw} = \sum_{t \in w} w_t \frac{\exp(X'_{mt}\zeta - \eta \log p_{mt} + \kappa r_{mkt} + \xi_{mk})}{1 + \sum_{m' \in k} \exp\left(X'_{m't}\zeta - \eta \log p_{m't} + \kappa r_{m'kt} + \xi_{m'kt}\right)}$$

given weights w_t, for all periods t in week w during which choice sets and observable characteristics may differ, for all pages m' ranked under k.

We observe click-through-rates per keyword search at the weekly level, for three separate weeks in 2021 and 2022, for a subset of product pages m that are ranked under keyword search k.[11] Therefore we are unable to estimate the parameters that govern click-through-rates using the entire choice set of pages m ranked under search k at all periods t in a given week w during which choice sets vary. However, we can approximate

$$\text{CTR}_{mkw} \approx \frac{\exp\left(\overline{X}'_{mw}\zeta - \eta \overline{\log p}_{mw} + \kappa \overline{r}_{mkw} + \overline{\xi}_{mwk}\right)}{1 + \sum_{m'k} \exp\left(\overline{X}'_{mw'}\zeta - \eta \overline{\log p}_{mw'} + \kappa \overline{r}_{m'kw} + \overline{\xi}_{m'kw}\right)}$$

where all variables are given by their average values during week w. Given that for any week in our data, we observe CTR_{0k}, or the fraction of consumers who searched for keyword k but did not click on any page m ranked under search k, we can employ the Berry inversion (Berry, 1994) to estimate

$$\log\left(\frac{\text{CTR}_{mkw}}{\text{CTR}_{0kw}}\right) = \overline{X}'_{mw}\zeta - \eta \overline{\log p}_{mw} + \kappa \overline{r}_{mkw} + \overline{\xi}_{mkw}$$

using linear IV, where the subscript w denotes a given week, and all variables in the regression are the average values of variables on the weekly level. Although this estimation based on averages is less than ideal, we nonetheless find it useful to estimate the extent to which observables determine click-through-rates.

The major challenge with using these estimates in a truly structural model is the endogeneity of both prices *and* ranks. Even on pages in which sellers use pricing algorithms to set prices, average prices at the weekly level will likely be correlated with relative quality.

[11] We observe 140 keywords, 80 product pages, and 284 product page, keyword, week tuples.

To mitigate this issue, we instrument for average prices with FBA fees. FBA fees vary across weeks within our sample on the same product pages, and across product pages within the same week as FBA fees are determined by weight. Therefore FBA fees should be uncorrelated with relative product quality to the extent that consumers do not perceive products with varying weights as being more or less desirable. Our data suggests that FBA fees are correlated with prices, and prices respond to increases in FBA fees across product pages, even for those product pages on which no offers are FBA.

Given that ranks are determined by platforms according to previous click-through-rates, rankings are inherently correlated with perceived product quality, assuming that perceived product quality on the weekly level is somewhat persistent over time. Unfortunately, our panel is not sufficiently long to employ the approach used by Yu (2024) and others to mitigate the endogeneity of ranks. Absent an alternative instrument for ranks, our estimates should be interpreted with caution. Nonetheless, observing the correlation between ranks and CTR is useful in understanding consumer behavior given $\mathbb{E}[\mathrm{CS}_m]$.

1.6 Results

1.6.1 Product Page Preferencing

Results for the Buy Box assignment rule are given in Table 1.4 where standard errors are computed from the Fischer information matrix.

In the Health and Household category, the Buy Box assignment rule grants a 4 % price premium to an Amazon offer relative to an FBA offer, an 8.7 % price premium to an Amazon offer relative to an FBM offer, and a 22 % price premium to an Amazon offer with zero shipping time relative to an FBM offer with two-day shipping time. The negative coefficient on rating difference may be due to assignments arising from advertising or other ineligible offers that we were not able to eliminate from the data. We remind the reader that due to the potential presence of these assignments in our data, the implied price premia for Amazon offers relative to non-Amazon offers will be less than those under the true preferencing rule.

Table 1.4: *Buy Box assignment rule estimation results.*

	Beauty and Personal Care	Health and Household
Log total price	-19.44^{***}	-17.21^{***}
	(0.25)	(0.22)
Shipping time	-0.88^{***}	-0.99^{***}
	(0.01)	(0.01)
FBA Non-Amazon	-0.17^{***}	-0.72^{***}
	(0.05)	(0.04)
FBM	-1.19^{***}	-1.44^{***}
	(0.06)	(0.05)
Log Feedback Difference	0.06^{***}	0.09^{***}
	(0.01)	(0.01)
Rating Difference	0.00	-0.02^{***}
	(0.00)	(0.00)
Out-of-Sample Accuracy	0.74	0.75
Observations	51,828	124,387

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Standard errors in parentheses.

The estimated rule performs predictably fairly well, with an out-of-sample accuracy of 75 percent.

In the Beauty and Personal Care category, the corresponding price premia are 0.08 % for an Amazon offer relative to an FBA offer, 6 % for an Amazon offer relative to an FBM offer, and 16 % for an Amazon offer with zero shipping time relative to an FBM offer with two-day shipping time.

1.6.2 Consumer Demand

Table 1.5: *Consumer demand estimation results.*

	Health and Household
Log total price	-32.42^{***}
	(0.04)
Shipping time	-1.42^{***}
	(0.01)
FBA Non-Amazon	-0.83^{***}
	(0.03)
FBM	0.84^{***}
	(0.03)
Rating difference	-0.18^{***}
	(0.00)
Log Feedback difference	-0.14^{***}
	(0.00)
Owns Buy Box	2.38^{***}
	(0.01)
Markets	388
Observations	430,349

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Standard errors in parentheses

Results for demand estimates are reported in Table 1.5. On average, consumers are willing to tolerate a 2.5 % price premium for an Amazon offer relative to a non-Amazon FBA offer, and 6.43 % price premium for an Amazon offer with less than 24 hour shipping time, relative to an FBM offer with two day shipping time. This is substantially lower than the corresponding 22 % price premium implied by the Buy Box assignment rule. It is important to note that most, but not all, Amazon and FBA offers in our sample have shipping time values in days, of zero, and therefore the shipping time coefficient mostly accounts for preferences in shipping time between FBM offers, who exhibit much greater variation in shipping time. Most of the FBM offers have positive shipping time values, which is why the FBM coefficient alone is positive. The coefficient on the Buy Box indicator implies a 7 % price premium for a Buy Box offer, relative to a non-Buy Box offer, all else equal. Consumers, on average, appear to be far less sensitive to differences in rating and review metrics among offers, conditional on fulfillment method, shipping time, and whether the offer owns the Buy Box or not.

1.6.3 Product Page Self-Preferencing Test

We conduct the product page self-preferencing test using all data from the Health and Household category for which we are able to estimate offer-level unobservable quality terms. Table 1.6 summarizes these estimates. Indeed, the indicator variable coefficients imply that Amazon engages in self-preferencing, and conditional on preferenced-independent mean utility, Amazon penalizes FBA offers and FBM offers relative to their own offers when assigning the Buy Box. We note that this does not necessarily imply that Amazon's assignment rule is suboptimal in terms of overall consumer welfare. As discussed earlier, on many product pages on the platform, sellers engage in dynamic price competition for which the model studied in the theoretical section above may not be appropriate. A more detailed study of the preferencing rule and its effects on welfare in these settings is relegated to chapter 2. What these results do suggest, is that Amazon's assignment rule differs, on average, from the rule that would maximize purchase probabilities, conditional on landing on a product page, at any moment in time.

Table 1.6: *Product page self-preferencing estimation results.*

	Health and Household
Choice Probabilities	3.93***
	(0.04)
FBA Non-Amazon	−0.78***
	(0.04)
FBM	−2.00***
	(0.04)
Out-of-Sample Accuracy	0.81
Observations	53,223

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Standard errors in parentheses

Table 1.7: *Regression of FBA fees on logarithmic total price.*

	Log Total Price
FBA fee	0.82***
	(0.02)
R-squared	0.90
F-statistic	2006.70
Observations	211

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Standard errors in parentheses

Table 1.8: *Click-through rate estimation results.*

	First Stage: Log-relative CTR	2SLS: Log-relative CTR
Cons	−5.50	−7.00
	(4.84)	(7.82)
Log Total Price	−0.10	−3.17***
	(0.17)	(1.02)
Shipping Time	−0.37	0.18
	(0.20)	(0.36)
Prime	−0.70	−0.45
	(0.32)	(0.52)
Log Feedback	−0.06	−0.13
	(0.09)	(0.14)
Rating	0.05	0.16*
	(0.05)	(0.09)
Average Organic Rank	−0.06***	−0.04***
	(0.00)	(0.01)
R-squared	0.52	−0.25
F-statistic	228.44	97.65
Observations	211	211

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Standard errors in parentheses

1.6.4 Click-Through Rates

Results for the two-stage least squares (2SLS) regressions instrumenting FBA fees for prices are reported in Table 1.8 and the first stage regression of FBA fees on prices is reported in Table 1.7. The resulting coefficients imply that rank-independent CTR do not appear to agree with expected consumer-surplus conditional on landing on product pages. Ranks do significantly influence product page visits and overall purchase probabilities, although these coefficients should be interpreted with caution.

Given that the observed ranking rule does not appear to maximize for either end-conversions or expected consumer surplus, this suggests that regardless of whether self-preferencing is prevalent in search result rankings, there is a welfare improvement from a ranking rule that optimizes for consumer surplus. Any additional losses in welfare arising from sponsored advertisements could be significant in that the average improvement in rankings for a sponsored ranking relative to an organic ranking, in our data, is 18 positions. Overall, a more detailed study of search rankings is necessary to fully determine the welfare effects of various sponsored ranking policies and organic rankings that optimize for welfare.

1.7 Discussion

Our empirical results suggest that Amazon engages in self-preferencing, awarding themselves a score that is greater than that implied by consumer preferences alone, relative to non-Amazon offers. Penalties are particularly severe for FBM offers, relative to the score premiums implied by consumer preferences. As discussed earlier, these results do not necessarily imply that Amazon's scoring rule harms consumers overall, since on many, but not all, product pages, prices do not correspond with the resting behavior studied in our model. On product pages where sellers do not adjust their prices dynamically, however, the results suggest that Amazon's self-preferencing does, in fact, harm consumers. Studying the welfare effects of preferencing rules in environments where prices update dynamically is the focus of chapter 2, in which we assess the counterfactual welfare implications of allowing for dynamically-updated prices versus ensuring the static outcome prices studied in this setting under the consumer-optimal Buy Box rule. This empirical study is nonetheless useful as a first-pass test of self-preferencing to measure the extent to which a platform's preferencing rule may deviate from the consumer-optimal rule under static competition, and may be applicable to a variety of other pricing platforms.

While data limitations prevented us from being able to conduct similar empirical assessments with keyword search rankings, in our sample, it is evident that Amazon may not choose rankings to maximize for either conversions or expected consumer surplus, the welfare consequences of which could be substantial.

Chapter 2

Platform Preferencing and Price Competition II

2.1 Introduction

Chapter 1 studied the optimal preferencing rule, within the class of linear scoring rules, under a myopic environment in which prices are determined in a locally envy-free profile. These prices correspond with *rest* points in pricing behavior, when firms may update their prices dynamically, albeit according to myopic payoff functions. This static setting may appropriately model environments in which sellers do not wish to invest in technologies that continually update prices based on opponents' behavior. However, as noted in previous work, many prices on product pages on Amazon do not correspond with such resting behavior, but rather, they are continually updated using strategies that result in highly distinctive price cycling patterns (Musolff, 2022).

The optimal preferencing rule, and the effects of preferencing rules on consumer and producer welfare in such a setting, are far from obvious. In order to the study the welfare trade-offs associated with various preferencing rules, in particular, the welfare effects of policy proposals to eliminate platform self-preferencing, this paper first studies how a platform's preferencing rule determines the dynamic pricing strategies that sellers employ.

Given that there are a limited class of strategies typically employed by sellers on Amazon and other platforms, the model restricts to studying the decisions over the parameters that govern these strategies, how they depend on seller primitives, and critically, how sellers fare relative to their opponents under the platform's preferencing rule. This paper is not only the first detailed empirical study of price cycling patterns on Amazon, but also the first, to my knowledge, that studies how preferencing rules govern sellers' decisions under cycling behavior.

2.2 Related Literature

Very few theoretical papers have studied preferencing rules in non-myopic environments. More recently, in the closely related auction literature, Banchio et al. (2024) studies a theoretical auction model where score quality measures are updated dynamically to analyze the resulting trade-offs between market thickness and information. For a broader review of the related dynamic auction literature see Bergemann and Said (2010), where previous models mostly focus on agents' evolving private information. This paper instead, studies preferencing environments in which firms optimize over non-myopic payoff functions, given restrictions on their strategy spaces.

This paper is closely related to the study of Edgeworth cycles, in particular, the work of Maskin and Tirole (1988), which studies Edgeworth cycles as Markov Perfect Equilibria (MPE) in an alternating-move Bertrand competition game with short-term commitment. Later computational work expanded on this seminal paper and analyzed the structure of Edgeworth cycles with more than two players (Noel, 2008), and empirically studied Edgeworth cycles in retail gasoline markets (Noel, 2007, 2012; Lewis and Noel, 2011). These authors, however, did not attempt to estimate the parameters that govern these cycles. Additionally, Eckert (2003) studies how firm's relative sizes influence the presence of cycling versus non-cycling equilibria, while Leisten (2024) models algorithmic pricing with managerial override, which empirically resemble gasoline markets. Notably, some of the assumptions of these models do not correspond with the environment on the Amazon platform. For example, the frequencies of price updates and cycle re-sets observed on Amazon are too high to justify the frequency of fluctuations in seller's costs in Noel (2008) and manual pricing re-sets in Leisten (2024).

Relatedly, Zhang and Feng (2005) describes evidence for the presence of pricing cycles in advertising auctions, environments that are closely related to the preferencing setting, as explored in Chapter 1 Other recent empirical work on platforms and dynamic pricing has shown that asymmetry in seller's reaction times drive differences in profits (Brown and MacKay, 2021). In their setting, seller's reaction times are taken to be exogenous, which in practice, are due to many factors including technological investment and the number of product pages on which a seller sells.

In the most closely related work to this paper, Musolff (2022) uses an event study design to provide evidence that automated pricing strategies lead to supracompetitive prices, and proposes a model of algorithm choice in which Edgeworth-like cycles may exist in equilibrium, although his model substantially deviates from Maskin and Tirole (1988). Importantly, Musolff (2022) assumes that sellers compete with homogeneous products, and does not embed the preferencing mechanism in his model. This paper critically endogenizes how pricing strategies and resulting welfare are determined by the preferencing rule, and importantly, seller's *relative* preferencing levels given by this rule. This allows me to estimate the primitives that govern sellers' payoffs and therefore estimate welfare under various counterfactual preferencing mechanisms.

Finally, this paper builds on the tacit algorithmic collusion literature. This more recent, burgeoning theoretical literature has shown that Reinforcement Learning algorithms can learn to price supracompetitively without intentionally colluding (Calvano et al., 2020, 2021; Asker et al., 2021; Banchio and Skrzypacz, 2022). Notably, (Klein, 2021) considers the outcomes of Q-learning algorithms in Maskin and Tirole (1988)'s model of alternating moves and finds that such models yield supracompetitive prices after training. Other theoretical work considers the platform design problem of assigning a seller preferenced status, and concludes that dynamic pricing rules are able to mitigate supracompetitive prices when sellers learn to set prices via learning algorithms relative to preferencing rules that condition only on current prices (Johnson et al., 2020). While seller's algorithms on platforms such as Amazon appear to be much simpler than the Reinforcement Learning strategies studied in these papers, by estimating seller's costs, I am able to estimate the extent to which prices and markups under dynamic pricing strategies differ from the static setting in Chapter 1, and estimate the relative costs of allowing for dynamic price updates.

2.3 Data

For a detailed description of the data used in this study, please refer to chapter 1.

2.3.1 Descriptive Statistics and Observed Pricing Patterns

On many platforms including Amazon, sellers have access to technologies that enable rapid responses to changes in the pricing environment, which result in pricing patterns that often do not converge to the resting behavior studied in chapter 1. This motivates my examination of price variation, particularly variation that arises among Buy Box prices, in order to study how preferencing rules govern pricing behavior in these settings, and their effects on welfare.

First, I define a product page-day as having *significant variation* in Buy Box prices if there are at least two distinct sellers who are awarded the Buy Box on that day, where for each such seller, the standard deviation in their Buy Box-winning prices is at least .2 times their average Buy Box winning-price, and there are at least six changes in Buy Box ownership. In our data, approximately 53 % of product pages in the product categories Beauty, Health and Household, Arts and Crafts, Office Products, Grocery, Pet, and Baby products have days during 2018-2022 in which the variation in Buy Box prices satisfies this criterion.

In general, for any given product page-day, data and industry knowledge suggest that sellers can be partitioned into those who *actively* compete for the Buy Box and those who do not. I define an offer as actively competing for the Buy Box on a given product page-day if they are awarded the Buy Box at least 90 percent of the time when they update their price.

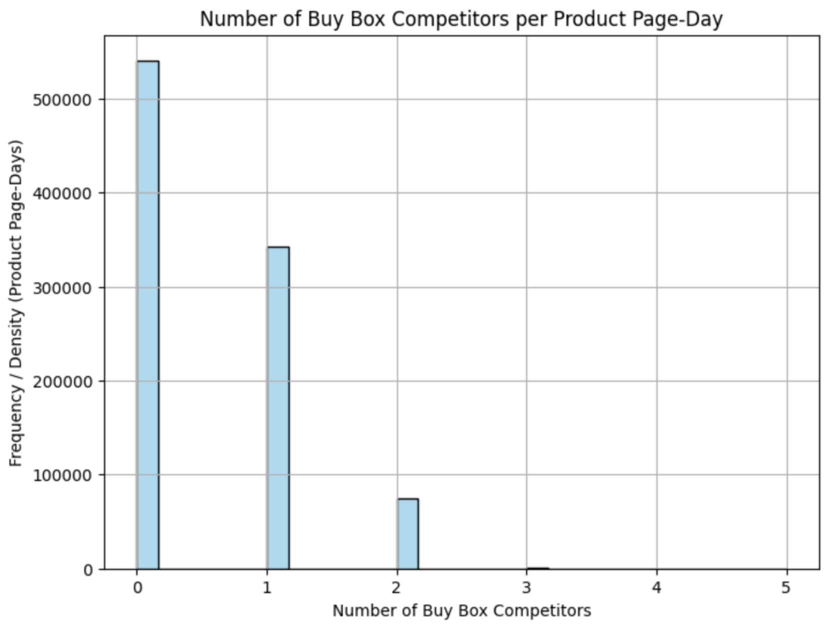

Figure 2.1: *The distribution of the number of Buy Box competitors on the product page-day level.*

The distribution of the number of sellers who actively compete for the Buy Box and those who do not are given in Figures 2.1 and 2.2. The average number of offers that actively compete for the Buy Box across all product-page days with price variation is .51, while the average total number of offers is 9.2, and the average number of offers without any price updates on a given day is 5.5. In only .1 percent of product page-days is the number of offers who actively compete for the Buy Box greater than 2. Indeed, there are many offers on product pages that do not actively compete for the Buy Box, and these offers have fewer price updates, on average 3.5 updates per day. The average number of price updates for active Buy Box competitors is given in Table 2.1. In general, the average number of price updates per seller, on the product-page day level, is higher for sellers in the set of Buy Box competitors when the size of this set is 2, versus when there is a single seller actively competing for the Buy Box, in particular above the median number of price updates.

A natural question is, where does Buy Box price variation come from in cases where the number of offers who actively compete for the Buy Box is less than two? This variation appears to arise from several different factors. First, the price variation criterion requires

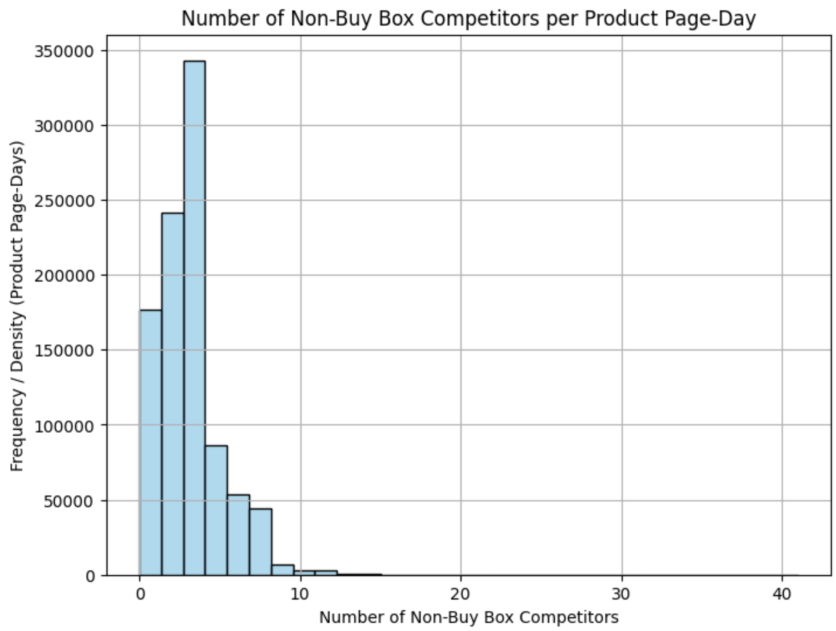

Figure 2.2: *The distribution of the number of Non-Buy Box competitors on the product page-day level.*

Table 2.1: *Average Price Updates per Seller Among Buy Box Competitors.*

Number of Buy Box Competitors	1	2
25th percentile	5	5
50th percentile	7	8
75th percentile	16	50.5

only that there are at least six changes in Buy Box ownership on the date of question. On many of these selected dates, the number of updates to Buy Box prices is close to this lower bound and hence many sellers simply do not make the 90 percent accuracy criterion. In some cases, frequent entry or exit of offers within the day also induces change in Buy Box ownership and prices, and eliminates sellers from the active Buy Box competition criterion. In other cases, particularly in earlier years, some sellers simply do not appear to have invested in pricing technologies that sufficiently undercut their opponents' price in order to obtain the Buy Box with 90 percent accuracy.

Since much of the Buy Box price variation appears to arise from Buy Box competition, and there are extremely few instances where the set of Buy Box competitors is greater than two, I will focus the study on pricing behavior on product page - days with two Buy Box competitors. Although there is nothing apparent that should preclude a seller from actively competing for the Buy Box with two other Buy Box-competitors, most sellers may simply find it unprofitable.

I then further examine pricing behavior to detect the presence of cycling as noted in previous work, within the product page-days that satisfy the criterion for significant Buy Box price variation (Musolff, 2022). I define a *complete price cycle* on a given day, as a series of weakly decreasing Buy Box prices consisting of at least four changes in Buy Box ownership, having at least four price changes between the maximum and minimum prices, and where there is a subsequent price increase immediately following the series of weakly decreasing prices, such that the difference between the maximum price and the price was subsequently raised to following the cycle is less than forty percent of the range of prices.

Table 2.2 summarizes the findings with respect to these pricing cycles among product-page days with significant price variation. In general, cycling behavior is very prevalent. On average, there is at least one complete pricing cycle per day on pages with price variation and many product page-days in which there are several complete cycles on a given day. Indeed, pricing cycles in our data exhibit a substantial range in prices of on average $4.52. Additionally, while the theoretical model in Musolff (2022) suggests that lower bounds in pricing cycles should be well above costs in order to sustain average payoffs equal to monopoly profits, this does not appear to be the case. Our seller's lower bounds on Buy Box prices are on average only \$.32 above cost, while in many cases, markups relative to costs on lower bounds are zero. If all other sellers' costs were identical to ours, their lower bounds, on average, would be \$0.32 below their costs, suggesting that most of our sellers' competitors have weakly higher costs and may set lower bounds in pricing cycles close to their own costs.

Musolff (2022) also suggests that sellers systematically raise prices to re-start pricing cycles overnight and re-commence undercutting one another in the morning hours. While there are pages in our sample in which raises take place after midnight, this is not systematically true. The distribution of average price updates per seller for both Buy Box competitors and non-Buy Box competitors, is shown in Figures 2.3 and 2.4. Although sellers who actively compete for the Buy Box on average, have slightly fewer price updates during the early morning hours, the number of updates is not significantly different across time. The time period during which pricing cycles re-start depends on a range of factors, including the frequency of sellers' price updates when Buy Box ownership changes, and the maxima and minima that govern the bounds on the pricing cycle.

Table 2.2: *Price Cycle Summary Statistics For Product Page-Days with Significant Buy Box Price Variation.*

Average Number of Complete Price Cycles Per Day	1.20
Average Number of Complete Price Cycles Per Day on Days with $\geqslant 1$ Cycle	3.00
Average Standard Deviation of Buy Box Prices	1.42
Average Markup Relative to Cost (All Buy Box Competitors)*	8.70%
Average Markup Relative to Cost (Our Seller)	8.80%
Average Markup Relative to Cost (All Non-Buy Box Competitors)*	14.80%
Average Minimum Buy Box Price Less Cost (Our Seller)	0.32
Average Maximum Buy Box Price Less Cost (Our Seller)	3.06
Average Range of Buy Box Prices	4.52
Median Number of Price Updates/Day (Buy Box Competitors)	40.50
Median Number of Price Updates/Day (Non-Buy Box Competitors)	3.00

* Note: If all sellers' costs were identical to our seller's, accounting for necessary differences in fulfillment and referral fees.

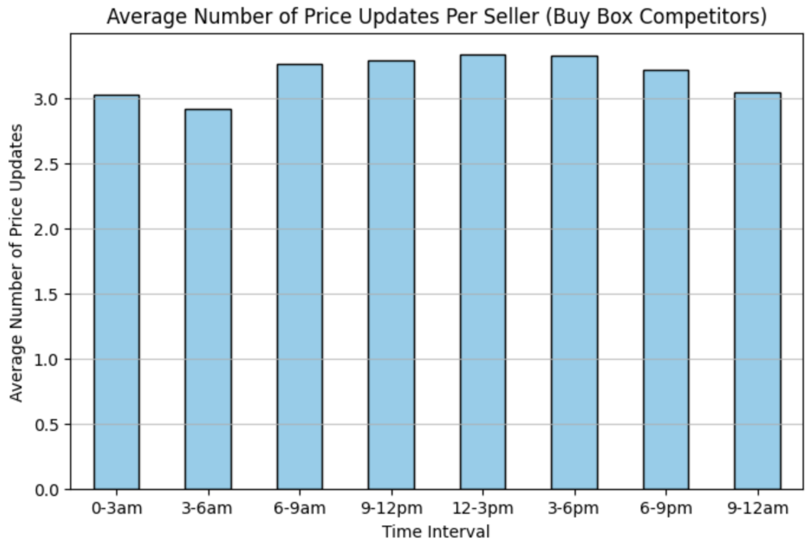

Figure 2.3: *Average Price Updates per 3-Hour Interval, Buy Box competitors.*

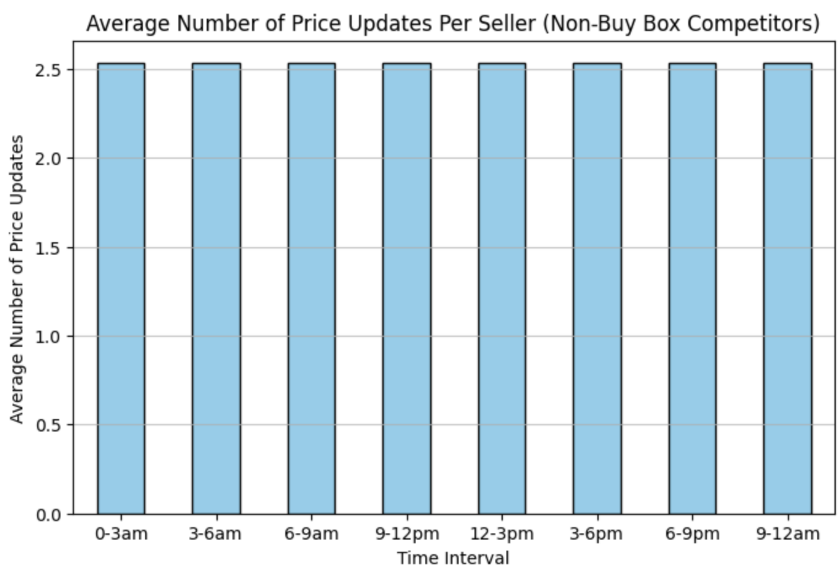

Figure 2.4: *Average Price Updates per 3-Hour Interval, Non-Buy Box competitors.*

2.3.2 Pricing Algorithms and Pricing Cycles

As discussed in previous work by Musolff (2022), pricing cycles on platforms such as Amazon arise from pre-set algorithms that determine dynamic pricing strategies. These algorithms require sellers to program maximum and minimum bounds on prices, and an amount by which they will undercut their fellow Buy Box competitor, which could be a positive or negative amount. This maximum price is also typically set when sellers face no Buy Box competition. Amazon, in fact, offers free automated pricing tools to their sellers that enable such pricing strategies, as shown in Figures 2.8 and 2.9.

Other third-party pricing companies offer pricing tools that are quite similar to Amazon's automated and custom-automated rules, set to undercut opponents by a fixed amount, along with pre-specified maximum and minimum pricing boundaries, and strategies to re-set prices when opponents' prices fall outside of the pre-specified bounds. For example, see Figures 2.10 and 2.11 for screenshots of pricing tools from third-party companies Seller Engine and Aura.

Most sellers in our data set *both* maxima and minima, rather than one simply always committing to undercutting the other, as is evident when one seller raises to a maximum price significantly above their opponent's maximum price. While Amazon allows for sellers to deploy other customized pricing rules, the strategy defined in the following section is most representative of the pricing strategies employed by sellers who actively compete for the Buy Box.

These strategies elicit pricing patterns that are highly reminiscent of Edgeworth cycles (Maskin and Tirole, 1988) with the exception that, as noted in Musolff (2022), we do not observe sellers deploying mixed strategies. What critically differs in this setting, relative to the theoretical setting in Maskin and Tirole (1988), is the minimum amount by which an opponent must undercut their rival, in order to obtain the Buy Box, depends on the preferencing rule, which assigns the Buy Box based on observable characteristics as well as prices. The preferencing rule, therefore, can give rise to asymmetry in payoffs from pricing cycles, even for sellers who otherwise have the same underlying primitives.

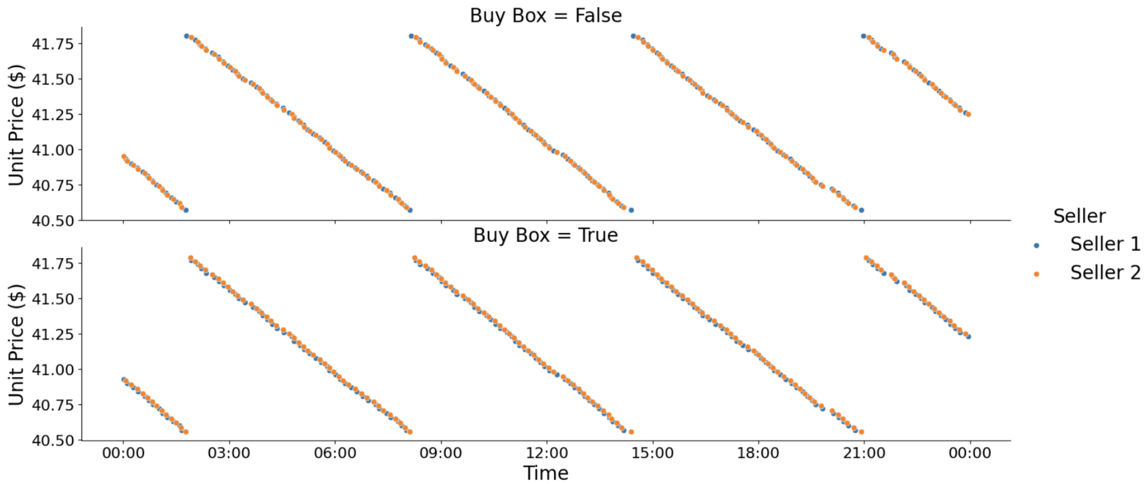

Figure 2.5: *Product page-day A, two non-Amazon sellers compete in pricing cycles.*

Examples of pricing cycles observed on Amazon for three different product page-days are shown in Figures 2.5, 2.6 , 2.7. In Figure 2.5, two non-Amazon sellers compete in prices, each undercutting the other to obtain the Buy Box, until the first seller raises to her maximum, to re-start the cycle. In Figure 2.6, Amazon competes with a non-Amazon seller, matching their price at every price update. Here we observe that Amazon re-starts the cycle first, and raises to a maximum price that significantly exceeds her opponent's maximum. Finally, in Figure 2.7, Amazon again competes with a non-Amazon offer and sets the amount by which to undercut their opponent at a significant, negative amount. Indeed Amazon appears to obtain the Buy Box at approximately $1 above their opponent's price. Sellers may undercut by amounts greater than the preferencing rule requires, but they cannot undercut below this amount.

As one might expect, the composition of Buy Box competitors appears to be related to those who score similarly in terms of the preferencing rule. As shown in Table 2.3, on the majority of product page-days with cycles, Amazon competes with an FBA offer, while most other pages consist of two FBM offers competing against one another. It is exceptionally rare for an FBA offer to compete with an FBM offer, and slightly less rare for an FBM offer to compete with an Amazon offer.

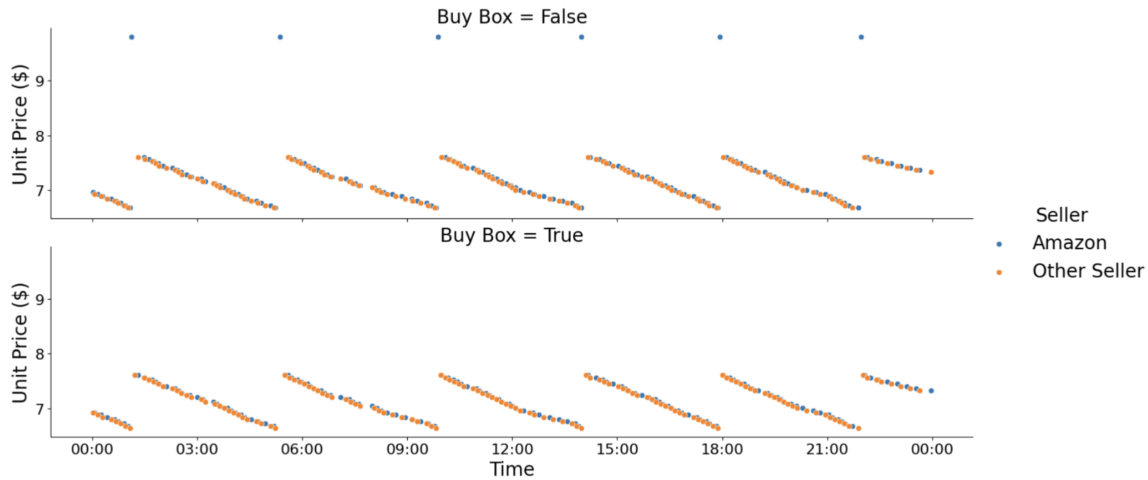

Figure 2.6: *Product page-day B, Amazon competes with a non-Amazon seller in pricing cycles.*

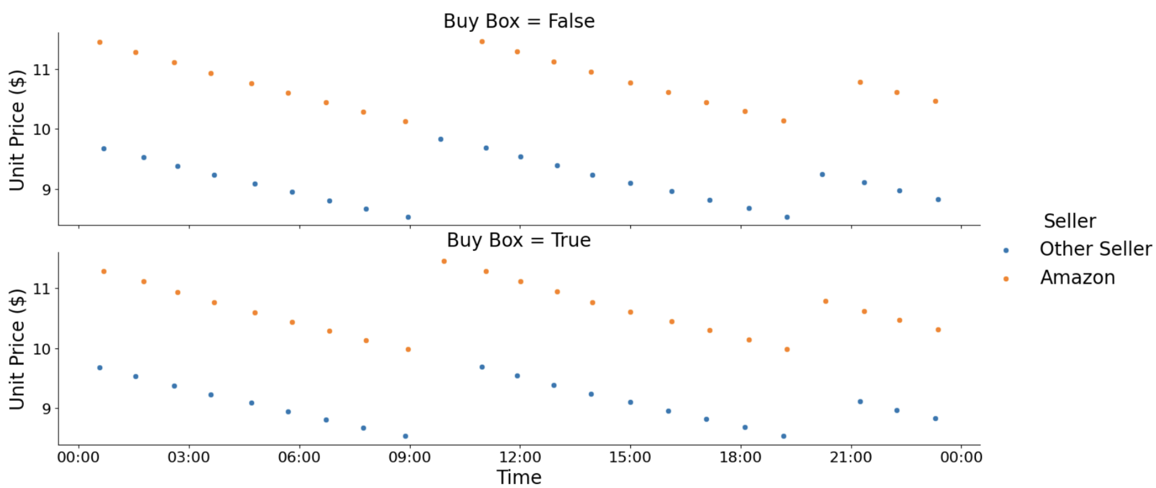

Figure 2.7: *Product page-day C, Amazon competes with a non-Amazon seller in pricing cycles.*

Table 2.3: *Composition of Buy Box Competitors Among Product-Page Days with Cycles.*

Competitors	Sample Composition
FBA-non-Amazon *vs.* FBA-non-Amazon	6%
FBM *vs.* FBM	36%
FBA-non-Amazon *vs.* FBM	1%
FBA-non-Amazon *vs.* Amazon	53%
FBM *vs.* Amazon	3%

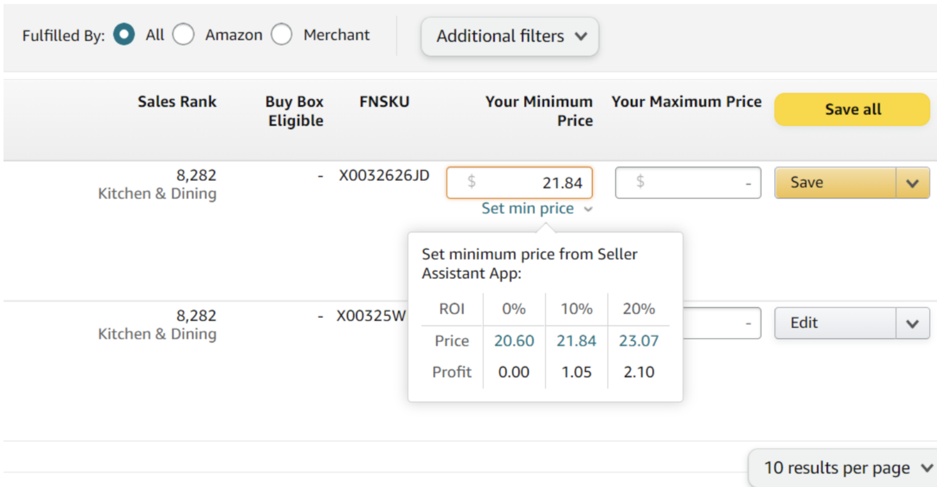

Figure 2.8: *Amazon seller's interface specifying product-page pricing strategies Seller Assistant App (2024).*

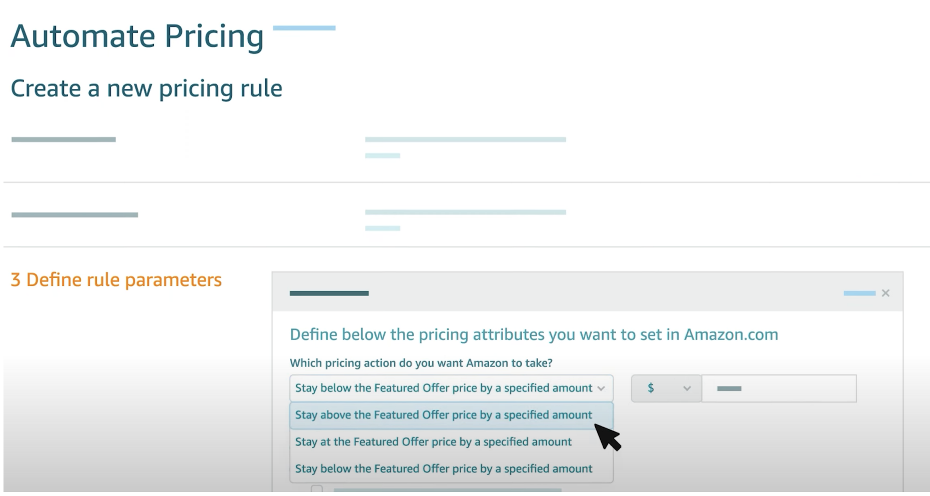

Figure 2.9: *Amazon Seller University's demonstration of programming product-page pricing strategies Amazon Seller University (2024).*

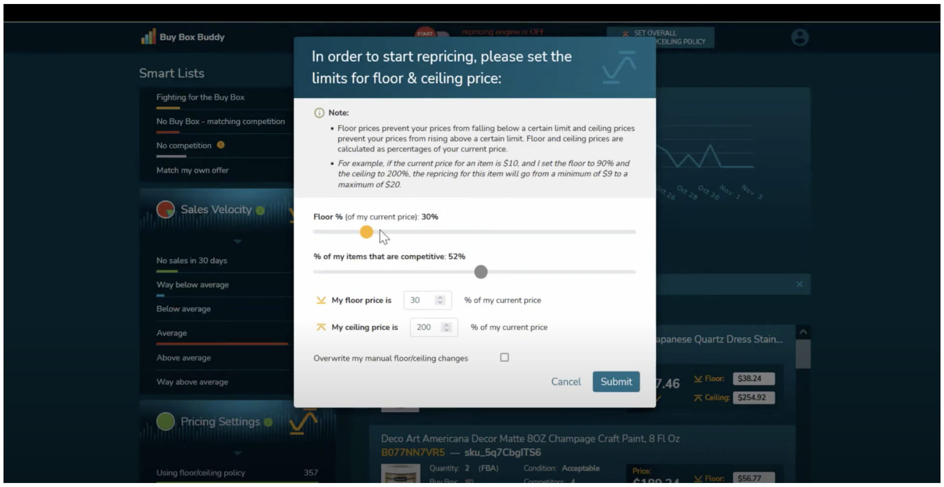

Figure 2.10: *Seller Engine's Buy Box Buddy SellerEngine (2024).*

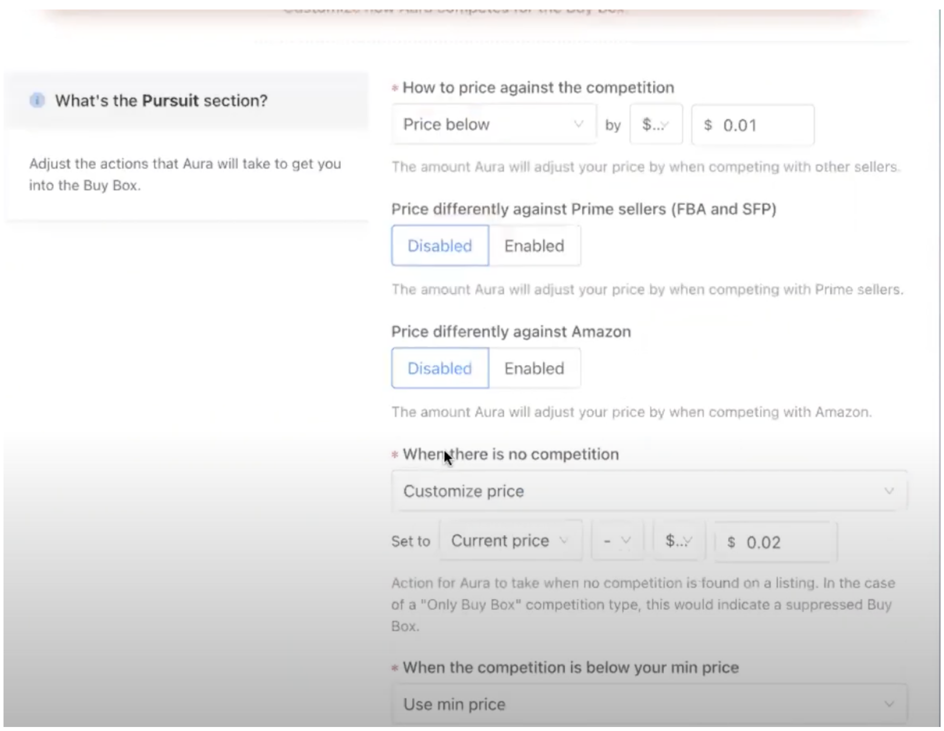

Figure 2.11: *Aura's Repricing interface Aura (2024).*

Table 2.4: *Entry and Exit Rates Among Non-Amazon Sellers After Amazon Entry.*

Exit Rate Non-Amazon Buy Box Competitors	87%
Exit Rate Non-Amazon Non-Buy Box Competitors	32%
Entry of Buy Box Competitors to Non-Buy Box Competitors	42%
Entry Buy Box Competitors	5%
Entry Non-Buy Box Competitors	33%

Indeed, there appears to be substantial exit on product pages among Buy Box competitors after Amazon enters a product page. Table 2.4 summarizes entry and exit patterns among product pages for the one month period after Amazon first begins selling on a given product page. Most non-Amazon Buy Box competitors stop competing on the product page altogether, while many of these sellers remain on the page as non-Buy Box competitors. There is substantial turnover among non-Buy Box competitors as well. Many of these sellers may simply find it unprofitable to compete against Amazon for the Buy Box given the preferencing rule and their available pricing strategies.

2.4 Model

Consider a setting in which a market is given by a product page m on a platform on a particular day d, where m, d will be suppressed from notation. In any market, there is a set of sellers $J = \{1, ..., N\}$, where all sellers $j \in J$ have vectors of fixed characteristics $X_j \in \mathbb{R}^d$. Sellers choose prices $p_j \in \mathbb{P}$ where $\mathbb{P} \subset \mathbb{R}$. We will denote the tuple $(X_1, X_2, \ldots, X_N)$ by $\mathbf{X}$ and the vector $(p_1, ..., p_N)$ by $\mathbf{p}$.

J is partitioned into the set of Buy Box-optimizers, $M \subseteq J$, and the set of non-Buy Box optimizers $J \backslash M$, which are fixed and observed by all sellers $j \in J$. We will restrict ourselves to a setting in which $|J| \geqslant 2$, and $|M| = 2$.

Following chapter 1, we will be concerned with preferencing rules that are given by linear scoring rules. The platform commits to a scoring rule, given by $f_j = X_j' \psi + \gamma p_j$, which determines which seller $j \in J$ is awarded the Buy Box, denoted by the vector of indicator variables $\mathbf{BB} = (\text{BB}_1, \text{BB}_2, \ldots, \text{BB}_N)$. $\text{BB}_j = 1$ if and only if $f_j \geqslant f_k$ for any $\forall k \in J \neq j$. Only one offer can be assigned the Buy Box at a time, and we assume that ties are broken at random with equal probability. Critically, the scoring rule determines the extent to which a seller must undercut an opponent's price in order to obtain the Buy Box.

For sellers $j \in M$, who compete for the Buy Box, $\eta_j(\psi)$ will denote the amount that seller j undercuts her opponent $-j \in M$, where this amount is fixed and is guaranteed to obtain the Buy Box with probability one, given the preferencing rule governed by ψ, γ. For notational simplicity, we suppress the dependence on the price coefficient γ and her fixed opponent $-j$. $\eta_j(\psi)$ is observed by all sellers.

These prices determine payoff functions, which at any fixed vector of prices $\mathbf{p}$ and corresponding Buy Box assignments $\mathbf{BB}$, will be given by

$$\pi_j(\mathbf{p}, \mathbf{X}, \mathbf{BB}, c_j) = s_j(\mathbf{p}, \mathbf{X}, \mathbf{BB}, \beta)(p_j - c_j)$$

where $s_j(\mathbf{p}, \mathbf{X}, \mathbf{BB}, \beta)$ are the corresponding logit shares governed by the consumer demand parameters β, as modeled in Chapter 1, and c_j is seller j's marginal cost.

For each seller $j \in M$, their pricing strategy will be given by a policy function, or *algorithm* whose state is given by her opponent, seller $-j \in M$'s price. Under parameters $\underline{p}_j, \overline{p}_j$, their strategy is given by

$$R_j(p_{-j}) = \begin{cases} p_{-j} - \eta_j(\psi) & p_{-j} \in (\underline{p}_j + \eta_j(\psi), \overline{p}_j + \eta_j(\psi)] \\ \overline{p}_j & p_{-j} \leqslant \underline{p}_j + \eta_j(\psi), p_{-j} > \overline{p}_j + \eta_j(\psi), p_{-j} = \varnothing \end{cases}$$

$\underline{p}_j, \overline{p}_j$ govern the boundaries within which, seller j undercuts her opponent's price by the fixed amount $\eta_j(\psi)$. For all opponent prices that exceed these bounds, seller j sets her price at $\overline{p}_j$. $\underline{p}_j$ is therefore the lowest price that j will obtain, and $\overline{p}_j$ is the maximum price that j will obtain under the above pricing strategy. $p_{-j} = \varnothing$ denotes the state in which her Buy Box competitor has exited the market and she faces no Buy Box competition. In this state, she prices at her maximum $\overline{p}_j$.

Each seller $j \in J \backslash M$, employs a pricing strategy that statically best responds to non-Buy

Box optimizing opponents, assuming that she herself does not obtain the Buy Box.

Formally,

$$p_j = argmax_p s_j(\mathbf{p}, \mathbf{X}, \mathbf{BB}, \beta)(p - c_j)$$

where s_j depends only on sellers in $J \backslash M$, and where $BB_j = 0$ for all such sellers. We assume that non-Buy Box competitors only respond to other non-Buy Box competitor prices, and that these prices indeed ensure that these sellers are never awarded the Buy Box when Buy Box competitors are in the market.

Each seller $j \in J$ updates their price according to the prescribed strategies after independently distributed exponential time with fixed rate λ_j, which is observed by all sellers. Sellers *exogenously* exit the market and re-enter with independently distributed exponential rates χ_j, κ_j.

Each seller $j \in M$ chooses $\underline{p}_j, \overline{p}_j$ *independently* to maximize her expected payoff, which is a weighted average of her expected payoff from a pricing cycle, assuming that her opponent $-j$ *always* undercuts her price by $\eta_{-j}(\psi)$, and her expected payoff from facing no Buy Box competition, accounting for price-update rates λ, entry and exit rates for non-Buy Box optimizers $\kappa_{J\backslash M}, \chi_{J\backslash M}$, non-Buy Box optimizer prices $\mathbf{p}_{J\backslash M}$, and her belief over the probability that her opponent $-j$ is not in the market, denoted by x_j[1]. These parameters solve

$$\max_{\overline{p}_j, \underline{p}_j} (1 - x_j) \left\{ \sum_{p_j, p_{-j}, \mathbf{p}_{J\backslash M}} \omega_{p_j, p_{-j}, \mathbf{p}_{J\backslash M}} (\lambda, \kappa_{J\backslash M}, \chi_{J\backslash M}, \overline{p}_j, \underline{p}_j, \eta(\psi)) \pi_j(p_j, p_{-j}, \mathbf{p}_{J\backslash M}, \mathbf{X}, \mathbf{BB}, \beta, c_j) \right\}$$
$$+ x_j \left\{ \sum_{\mathbf{p}_{J\backslash M}} \omega_{\mathbf{p}_{J\backslash M}} (\lambda_{J\backslash M}, \kappa_{J\backslash M}, \chi_{J\backslash M}) \pi_j(\overline{p}_j, \varnothing, \mathbf{p}_{J\backslash M}, \mathbf{X}, \mathbf{BB}, \beta, c_j) \right\} \quad (2.1)$$

where weights $\omega_{\mathbf{p}}$ denote the average daily weight over the relevant price vector $\mathbf{p}$, given the selected strategies, where these weights depend upon entry rates, exit rates, price update rates, and when opponent $-j$ is believed to be in the market, the selected bounds $\underline{p}_j, \overline{p}_j$ as well as $\eta_j(\psi), \eta_{-j}(\psi)$. Since seller j assumes that her opponent *always* undercuts her price by

[1]We allow Buy Box optimizers' beliefs over opponent entry and exit rates to differ from from the realized rates.

$\eta_{-j}(\psi)$, expected profits from pricing cycles are computed beginning in the state in which $p_j = \overline{p}_j$ and p_{-j} is the lowest price that $-j$ obtains in the cycle given $\underline{p}_j, \eta_j(\psi), \eta_{-j}(\psi)$.

The seller $j \in M$ with the weakly higher lower bound, $max_{j \in M} \underline{p}_j + \eta_j(\psi)$ will be referred to as the *leader* in the resulting pricing cycle, and will raise to their maximum first. The remaining seller $-j \in M$ will be referred to as the *follower*. The prices and corresponding Buy Box assignments will be those arising from $R_j(p_{-j}), R_{-j}(p_j)$ beginning in the state in which $p_j = \overline{p}_j$ and p_{-j} is $-j$'s lowest price sustained in a cycle governed by these strategies.

Seller j's resulting expected payoffs are given by the weighted average

$$\sum_{\mathbf{p}} \omega_{\mathbf{p}}(\lambda, \kappa, \chi, \underline{p}, \overline{p}, \eta(\psi)) \pi_j(\mathbf{p}, \mathbf{X}, \mathbf{BB}, \beta, c_j)$$

where weights ω_p denote the daily average over each *realized* price vector and Buy Box assignment given price-update, entry, and exit rates, and strategies resulting from $\overline{p}, \underline{p}, \eta(\psi)$.

In this model, we allow for the parameters that govern pricing cycles to be selected in a single-agent, non-equilibrium setting. In particular, we allow *both* players' strategies to entail raising to maxima conditional on opponent prices reaching a minimum threshold. Given that in observed pricing cycles, we observe the leader in the pricing cycle choosing a maximum price sufficiently far above the follower's subsequent price, it is evident that both players are choosing these strategies more often than would justify a model as in Musolff (2022), where the follower simply chooses to undercut their opponents' price in every state. Accounting for the possibility of one's Buy Box competitor exiting appears to be important in selecting these maxima.

Note that I assume that players use undiscounted average payoff functions to determine their maxima and minima prices, as the average price in a profit-maximizing cycle will depend on both a discount factor and a belief that one faces no Buy Box competition, and therefore both will not be separately identified from data alone. Since the frequency of price updates can be too reasonably high to justify a discount factor less than one, I instead choose a model in which payoffs depend on beliefs about opponent exit rates, and therefore the ability to price at monopoly prices. As noted previously, these exit rates are critical in determining chosen bounds in cycles, as without them, players should always choose cycles with average prices centered around monopoly prices, as is true in Musolff (2022)'s model by design. However, observed bounds in pricing cycles in our data appear too often centered below monopoly profits to be justified without considering additional exit rates. Since true exit rates will depend on selected prices, I allow players to hold beliefs over exit rates, that may or may not correspond with observed exit rates.

It is worth noting that strategies that coincide with Markov Perfect Equilibria (MPE) are not necessarily ruled out in this setting. When sellers have heterogeneous characteristics, and therefore different η_j, even for otherwise identical primitives, sellers payoff from equilibrium pricing cycles may differ. Indeed these equilibrium strategies may not involve mixed strategies at the lower bounds in cycles as in Maskin and Tirole (1988).

Nonetheless, MPE may not be appropriate in this setting for other reasons. In particular, for states in which a seller undercuts their opponent, they must have a *sufficiently low* discount factor less than one in order to prefer doing so to undercutting further and inducing their opponent to raise to their maximum price. Such a discount factor may not be reasonable given the frequency of price updates. Additionally, the class of MPE for a given set of seller primitives may in fact be very large. Without pinning down at least one of the seller's parameters $\overline{p}_j, \underline{p}_j$, it may be exceptionally difficult to solve for and select an MPE, even in a full information setting. It is unlikely that individual sellers choose parameters in this way given this challenge. In fact, industry knowledge suggests that individual sellers select these bounds based on their markups and their beliefs over opponent exit rates, rather than internalizing equilibrium conditions arising from opponents' upper and lower bound decisions.

2.4.1 Comparative Statics

Before outlining the estimation procedure for seller primitives, it is useful to consider the effects of varying parameters in the model on Buy Box optimizers' selected pricing bounds, as well as consumer and producer welfare. In general, the *center* of the cycle governs the average payoff for fixed η_j, η_{-j} values, which for any given c_j, will depend on x_j, the belief that j faces no Buy Box competition and thus allows her to price at her maximum. x_j, c_j govern the bounds $\underline{p}_j, \overline{p}_j$ that define this center. The *dispersion*, or width of the cycle, given by $\underline{p}_j, \overline{p}_j$, for any given x_j, c_j values depend on η_j, η_{-j} that maximize average payoffs given the optimal center.

The following figures illustrate the comparative statics of optimally chosen bounds $\underline{p}_j, \overline{p}_j$ in a simulation for a fixed seller j on a product page, with respect to varying parameters $\eta_j, \eta_{-j}, x_j, c_j, \lambda_j$.

As η_j alone increases, as shown in Figure 2.12, optimally chosen bounds tend to *expand*, up until a point at which obtaining the Buy Box at more than one price becomes suboptimal. Intuitively, for a fixed η_{-j}, an increased η_j implies that at fixed bounds, each step in a cycle is at a greater distance from profit-maximizing prices. Expanded bounds may imply fewer steps in a cycle and a lower range of Buy Box prices, but a higher range of overall prices, and decreased profits. In an extreme case, if η_j is sufficiently high, a seller may prefer setting $\underline{p}_j = \overline{p}_j$ and simply remain off of the Buy Box until their opponent exits.

Similarly, in Figure 2.13, as η_{-j} increases, for a fixed η_j, optimal bounds also *expand* and profits decrease, since a higher opponent η_{-j} requires lower on-Buy Box prices for the same selected bounds.

As previously noted, x_j governs the *center* of the cycle. For $x_j = 0$, where the seller does not account for opponent exit, the selected bounds are centered around monopoly profits, as expected. As x_j increases, in Figure 2.14, the selected maximum approaches the monopoly price, and the chosen bounds may contract around the optimal price. Higher x_j values tend to increase consumer welfare, although these beliefs may not coincide with true entry and exit rates.

Intuitively, increasing c_j monotonically increases profit-maximizing prices for any fixed parameters. As shown in Figure 2.15, since marginal cost also governs the minimum possible lower bound in a cycle, as well as the upper bound, increasing costs can contract the optimal chosen bounds, all else equal.

On the other hand, λ_j does not shift selected bounds in cycles, but merely changes the relative time spent in seller j's off-Buy Box prices. Higher λ_j, or a slower reaction time, decreases average profits as shown in Figure 2.16, while increasing λ_{-j} has a similar, opposite effect of increasing one's own average profits. These rates together will determine the *number* of cycles on a given day for fixed $\overline{p}, \underline{p}, \eta$.

To estimate the effects of varying *preferencing* rules on consumer and producer welfare it is important to consider how various rules shift the bounds and prices in a cycle. To do so, consider a setting in which η_j, η_{-j} depends on the preferencing rule governed by ψ in a particular manner.

For any seller j and opponent $-j$, the extent to which j must undercut p_{-j} in order to achieve $-j$'s score, under the scoring rule, is given by $\frac{X_j'\psi - X_{-j}'\psi}{\gamma}$, where since $\gamma < 0$, a higher relative $X_j'\psi$ implies a *lower* value. Indeed, this value may be negative. In particular, consider a setting in which η_j is defined by

$$\eta_j(\psi) = \frac{X_j'\psi - X_{-j}'\psi}{\gamma} + \epsilon_j$$

for some fixed $\epsilon_j \geqslant .01$, which guarantees that j undercuts her opponent by at least the minimum amount guaranteed to award her the highest score. η_{-j} is defined analogously, for some fixed $\epsilon_{-j} \geqslant .01$.

If a scoring rule increases $\eta_j(\psi)$, then $\eta_{-j}(\psi)$ proportionally decreases, for any fixed ϵ values. Figure 2.17 shows optimally chosen bounds for seller j for given minimum undercut values given by $\frac{X_j'\psi - X_i'\psi}{\gamma}$, and corresponding $\eta_j(\psi), \eta_{-j}(\psi)$ with fixed $\epsilon_j, \epsilon_{-j}$ values. Conditional on j remaining the *leader* in the pricing cycle, increasing her own minimum undercut value implies that her opponent prices *higher* for the same pricing bounds. This can increase j's profits and shift the center of her cycle closer to a profit-maximizing price that she would obtain if she were the only Buy Box competitor. While on average, this may come at a cost to consumer welfare, for minimum undercut values relatively close to one another, the effects are more ambiguous, since higher prices on path for one seller may indeed be compensated by *lower* prices on path for another given adjusted bounds. Whether or not these rules improve consumer welfare will depend on the extent to which changes in prices are compensated for by consumer's preferences for the various sellers, given their characteristics. However, if j's minimum undercut value becomes sufficiently high, her optimal prices may contract sufficiently to no longer improve profits, and may come at a disadvantage due to fewer on-Buy Box prices. Therefore, conditional on remaining the leader, paradoxically, it may benefit the seller to have a *lower* relative score insofar as the decrease in one's own score does not limit one from obtaining the Buy Box sufficiently often.

Importantly, since the leader in the resulting pricing cycle is determined by the seller with the maximum $\underline{p}_j + \eta_j$, a scoring rule that results in an increased lower bound for an opponent can change the leader and follower in a counterfactual setting. Whether or not the resulting follower's profits improve in the counterfactual, depends on whether the parameters that govern the opponent's cycle allow for prices that are sufficiently close to one's own profit-maximizing prices, which will depend on the sellers' relative primitives as well as the resulting η values.

There is no guarantee that there exists a *unique* preferencing rule governed by ψ that optimizes for either consumer or producer welfare, given the often off-setting effects of a rule that improves profits for one seller over another. Nonetheless, counterfactual simulations will be useful to examine the *relative* welfare costs among various rules, and the trade-offs of policy proposals.

2.4.2 A Note on Non-BuyBox Competition

Recall that non-Buy Box competitors, $j \in J \backslash M$, play static best responses, among only fellow non-Buy Box competitors when updating their prices, and assume that they will not obtain the Buy Box. Their prices, therefore, do not respond to Buy Box competitor behavior. This simplifies the analysis in that in a counterfactual setting, I assume that prices from off-Buy Box competitors remain fixed, and they do not obtain the Buy Box while Buy Box competitors are active. Since I am mainly interested in changes in welfare due to the varying pricing cycles from Buy Box competitors, this setting allows me to lessen the computational

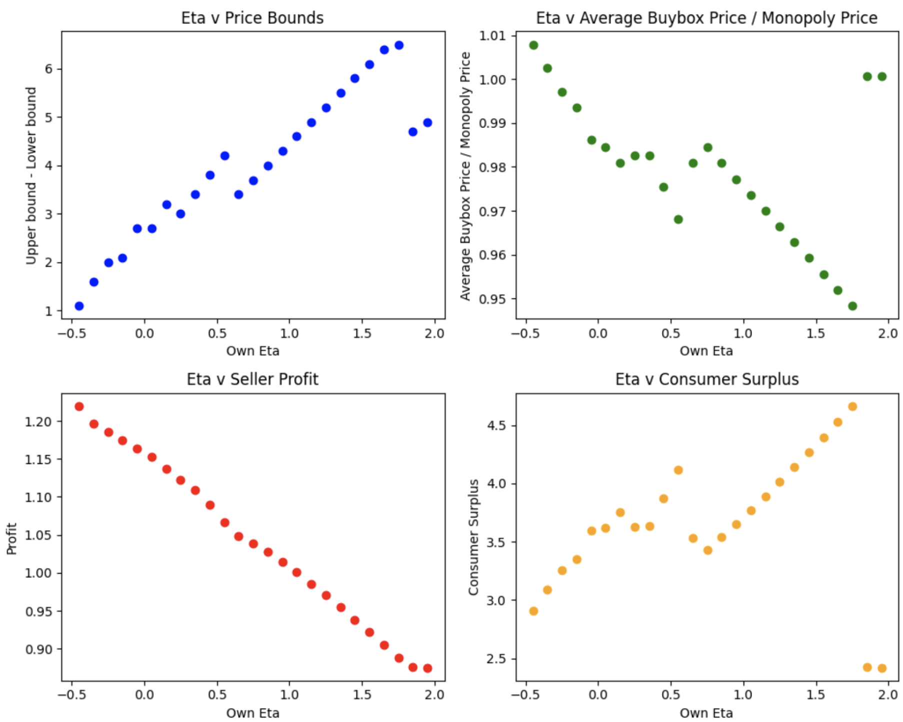

Figure 2.12: *Simulation of an Amazon seller's optimally chosen bounds and resulting welfare conditional on her remaining the leader, for varying η_j, given fixed values $\eta_{-j} = 0.5$, $c_j = 10$, $x_j = 0$, $\lambda_j = 5$, $\lambda_{-j} = 5$ for an FBA opponent.*

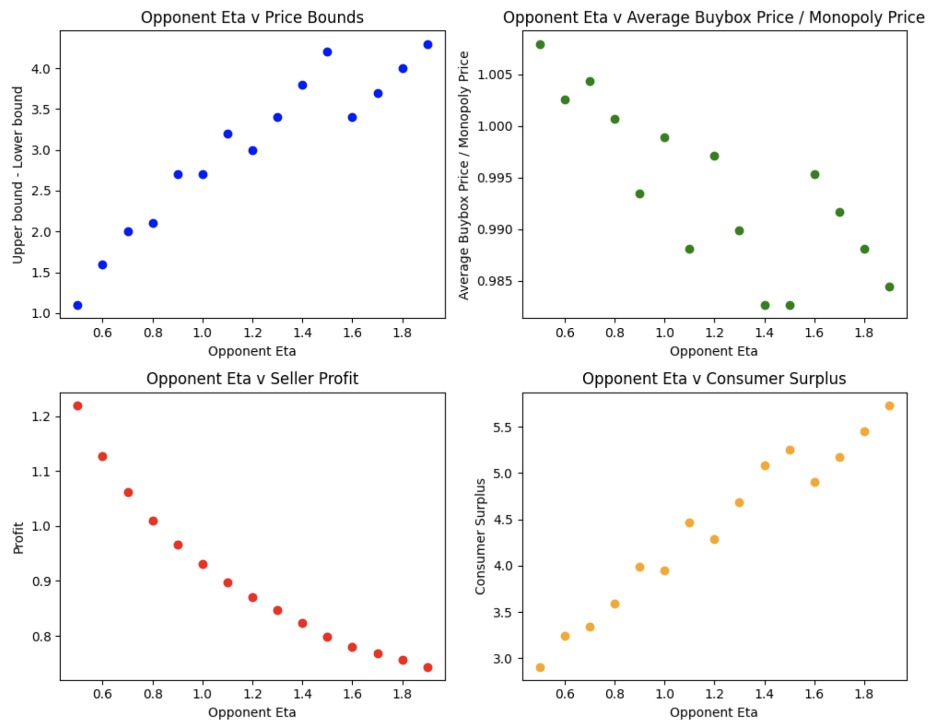

Figure 2.13: *Simulation of an Amazon seller's optimally chosen bounds and resulting welfare conditional on her remaining the leader, for varying η_{-j}, given fixed values $\eta_j = -0.45$, $c_j = 10$, $x_j = 0$, $\lambda_j = 5$, $\lambda_{-j} = 5$ for an FBA opponent.*

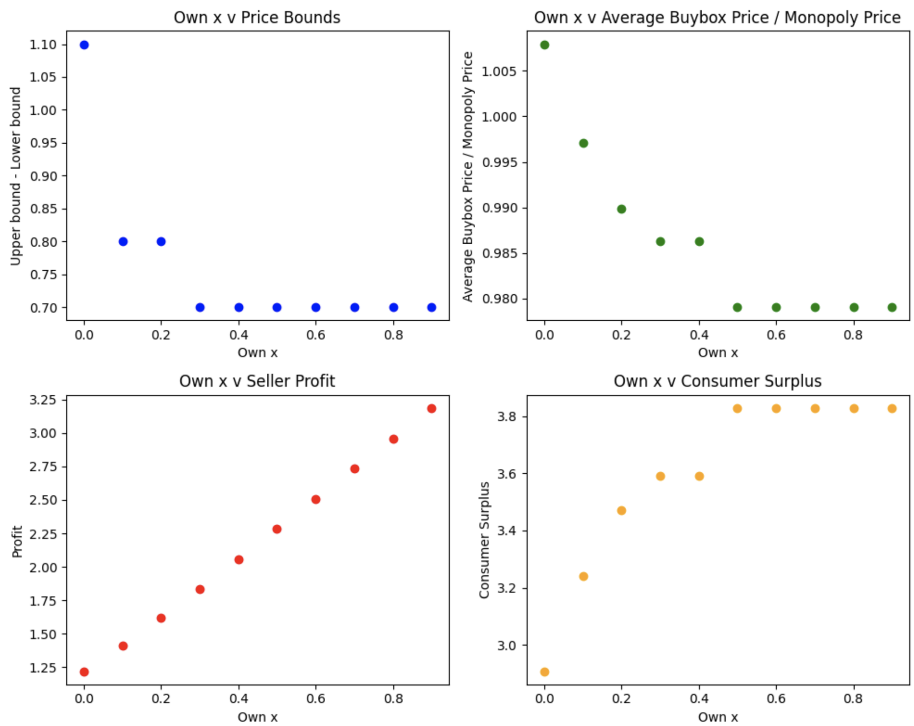

Figure 2.14: *Simulation of an Amazon seller's optimally chosen bounds and resulting welfare conditional on her remaining the leader, for varying x_j, given fixed values $\eta_j = -0.45$, $\eta_{-j} = 0.5$, $c_j = 10$, $\lambda_j = 5$, $\lambda_{-j} = 5$ for an FBA opponent.*

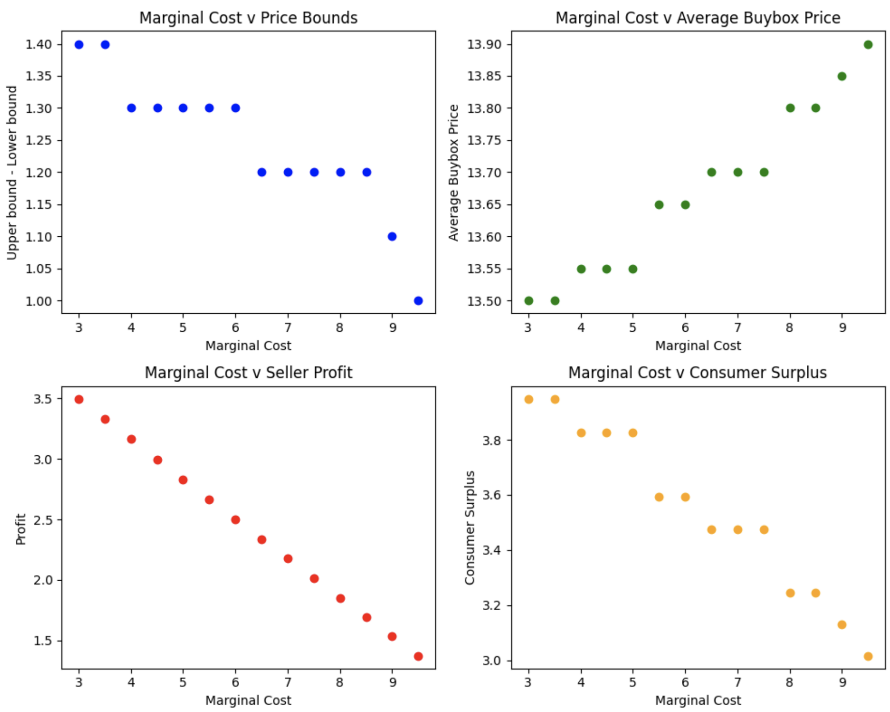

Figure 2.15: *Simulation of an Amazon seller's optimally chosen bounds and resulting welfare conditional on her remaining the leader, for varying c_j, given fixed values $\eta_j = -0.45$, $\eta_{-j} = 0.5$, $x_j = 0$, $\lambda_j = 5$, $\lambda_{-j} = 5$ for an FBA opponent.*

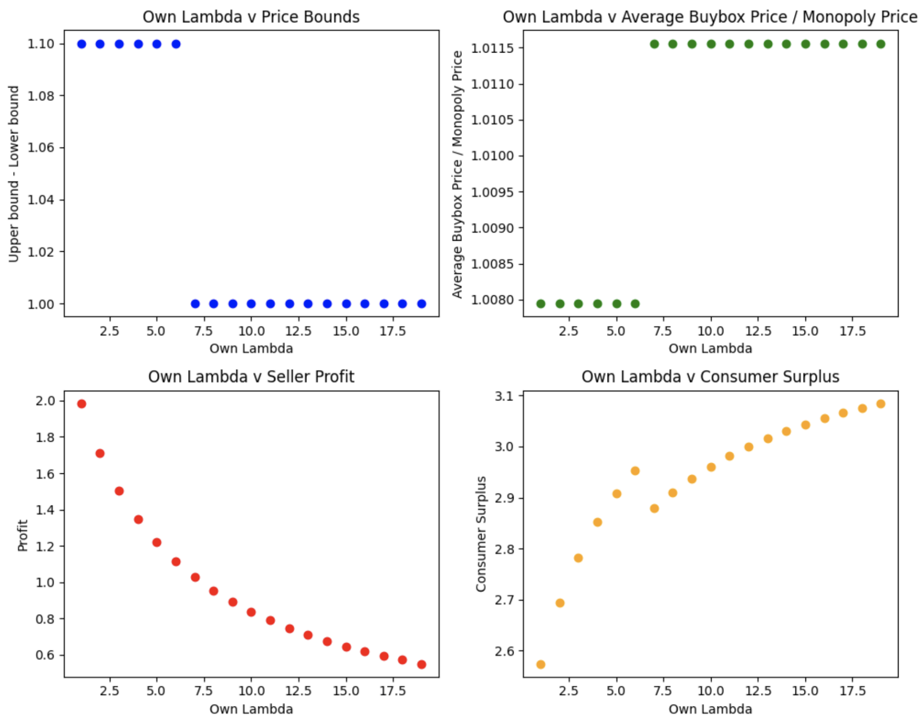

Figure 2.16: *Simulation of an Amazon seller's optimally chosen bounds and resulting welfare conditional on her remaining the leader, for varying λ_j, given fixed values $\eta_j = -0.45$, $\eta_{-j} = 0.5$, $c_j = 10$, $x_j = 0$, $\lambda_{-j} = 5$ for an FBA opponent.*

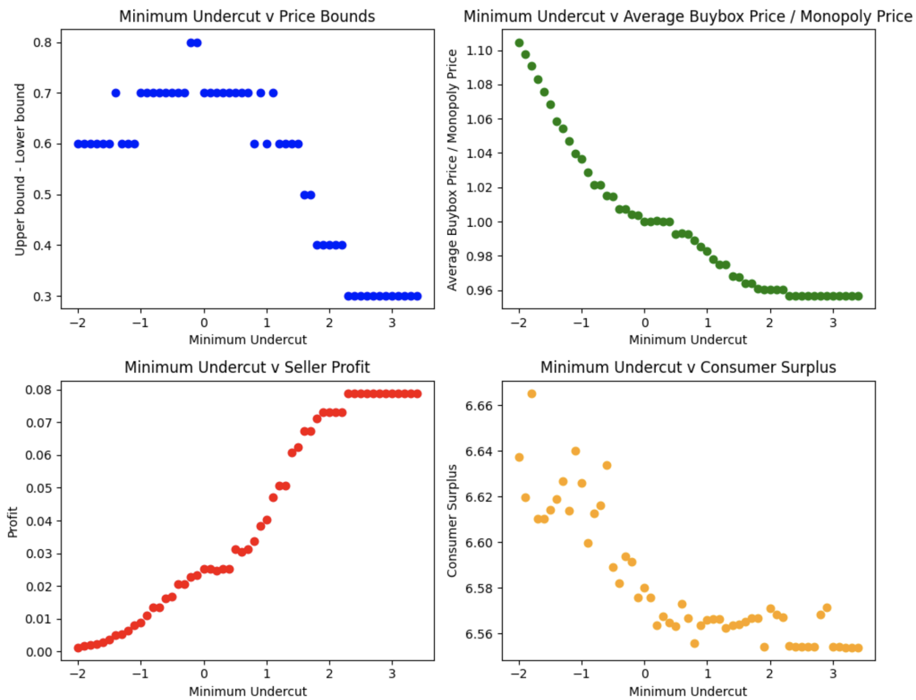

Figure 2.17: *Simulation of an Amazon seller's optimally chosen bounds and resulting welfare conditional on her remaining the leader, for varying minimum undercut values, given fixed values $x_j = 0$, $c_j = 10$, $\lambda_j = 5$, $\lambda_{-j} = 5$, for an FBA opponent.*

burden from simulating equilibrium prices from off-Buy Box competitors that would ensure that they fail to obtain the Buy Box at every estimated preferencing rule.

Given the consumer preference for the Buy Box as estimated in chapter 1, it is unlikely that changes in off-Buy Box competitor prices would significantly shift profit-maximizing prices for Buy Box competitors, while fluctuations in prices off-Buy Box that statically profit maximize assuming one does not obtain the Buy Box should also be mild given changes in on-Buy Box prices. This setting nonetheless allows me to estimate off-Buy Box competitor primitives and account for their changes in welfare under various Buy Box pricing cycles.

2.5 Estimation

In order to estimate seller primitives, for any given product page-day, I first identify the set $M \subseteq J$, and $J \backslash M$. If a seller obtains the Buy Box when she updates her price at least 90 percent of the time, she is in the set M, else she is in $J \backslash M$. We limit our estimation to product-page days on which $|M| = 2$, and therefore in which I observe price cycling behavior.

2.5.1 Calibrating Seller Primitives: Buy Box Competitors

Among the set M, I identify the *leader* in the observed data as the seller who raises to their maximum first. For this seller, we observe $\underline{p}_j, \overline{p}_j$ as well as η_j, η_{-j}. The price update rates λ, entry and exit rates κ, χ are estimated using Maximum Likelihood Estimation.

$\overline{p}_j, \underline{p}_j$ will be jointly determined by both c_j, x_j. For any fixed c_j, x_j determines the profit-maximizing center of the cycle, relative to monopoly prices, which governs $\underline{p}_j, \overline{p}_j$, while for any fixed x_j, c_j determines the profit-maximizing $\underline{p}_j, \overline{p}_j$. Given the observed data, $\underline{p}_j, \overline{p}_j, \eta_j, \eta_{-j}, \lambda, \kappa_{J\backslash M}, \chi_{J\backslash M}, \mathbf{p}_{J\backslash M}, \mathbf{X}, \mathbf{BB}$, I am interested in the parameters $\theta_j = (c_j, x_j)$ for which the corresponding $\underline{p}_j, \overline{p}_j$ maximize expected profits.

I solve for a θ_j that minimizes

$$min_{\theta_j=(c_j,x_j)} \sqrt{(\overline{p}_j - \overline{p}_j(\theta_j))^2 + (\underline{p}_j - \underline{p}_j(\theta_j))^2}$$

where $(\overline{p}_j(\theta_j), \underline{p}_j(\theta_j))$ solves

$$max_{\overline{p}_j,\underline{p}_j}(1-x_j)\{\sum_{p_j,p_{-j},\mathbf{p}_{J\backslash M}} \omega_{p_j,p_{-j},\mathbf{p}_{J\backslash M}}(\lambda,\kappa_{J\backslash M},\chi_{J\backslash M},\overline{p}_j,\underline{p}_j,\eta)\pi_j(p_j,p_{-j},\mathbf{p}_{J\backslash M}\mathbf{X},\mathbf{BB},\beta,c_j)\}$$
$$+ x_j\{\sum_{\mathbf{p}_{J\backslash M}} \omega_{\mathbf{p}_{J\backslash M}}(\lambda_{J\backslash M},\kappa_{J\backslash M},\chi_{J\backslash M})\pi_j(\overline{p}_j,\emptyset,\mathbf{p}_{J\backslash M},\mathbf{X},\mathbf{BB},\beta,c_j)\}$$

where expected payoffs are as in equation 2.1. While a unique pair (c_j, x_j) that minimizes the objective function may not be guaranteed, due to the sensitivity of both $\underline{p}_j, \overline{p}_j$ with respect to both parameters, any differences among identified pairs should be slight.

For corresponding *followers* in a pricing cycle, I observe only a bound on their chosen $\underline{p}_j$, and possibly observe their chosen $\overline{p}_j$, or a lower bound on $\overline{p}_j$[2]. Denote the observed bounds by $\underline{p}_j^b, \overline{p}_j^b$.

I calibrate the set $\{\theta_j : \overline{p}_j^b \leqslant \overline{p}_j(\theta_j), \underline{p}_j^b \geqslant \underline{p}_j(\theta_j)\}$, and additionally calibrate $min_{\theta_j=(c_j,x_j)}\sqrt{(\overline{p}_j - \overline{p}_j(\theta_j))^2}$ when applicable, where $(\overline{p}_j(\theta_j), \underline{p}_j(\theta_j))$ are given by the maximization problem above.

2.5.2 Calibrating Seller Primitives: Non-Buy Box Competitors

For Non-Buy Box optimizers $j \in J\backslash M$, I estimate $\theta_j = c_j$ given observed price updates p_j, which solve

$$argmax_{p_j}\frac{exp(X_j'\beta - \alpha p_j + \phi + \xi_j)}{\sum_{k\in J\backslash M} exp(X_k'\beta - \alpha p_k + \phi + \xi_k)}(p_j - c_j)$$

given off-Buy Box competitor characteristics $\mathbf{X}$ and prices $\mathbf{p}$ at the time of the observed price update, and estimated demand parameters β, α, ϕ, ξ from chapter 1. If a seller in $J\backslash M$ does not have a price update on the day of interest, I use their most recent prior price update to estimate their marginal costs.

[2] I observe $\overline{p}_j$ exactly when the leader raises to a maximum price above $\overline{p}_j + \eta_j$.

2.5.3 Counterfactual Estimation Procedure

The main interest in counterfactual estimation is to determine the trade-offs among consumer and producer welfare from various preferencing rules. Although welfare-optimal rules may not be unique for any given welfare measure, I begin by estimating rules that coincide with various policy proposals, and additionally, compare outcome measures to those that optimize over consumer and producer welfare. This allows me to estimate the distribution of the *relative costs* of one preferencing rule versus another, in terms of various welfare measures.

I restrict to a setting in which for any seller $j \in M$, given the preferencing rule governed by ψ and opponent $-j \in M$,

$$\eta_j(\psi) = \frac{X_j'\psi - X_{-j}'\psi}{\gamma} + \epsilon_j$$

In the data, given observed η_j and estimated $\hat{\psi}, \hat{\gamma}$ from the Buy Box assignment rule in staticpaper, I observe that $\epsilon_j \sim N(.032, 1.12)$, where these errors could account for error in the estimated Buy Box assignment rule. Evidently, on average, sellers undercut their opponents by only a few cents more than would award them an equal score. Before estimating a counterfactual scoring rule, I assign each seller a random draw from the observed distribution of ϵ errors truncated to be greater than or equal to .01, which ensures that under any assignment rule, η_j is at least as large as that which would ensure Buy Box assignment, and remains fixed under all counterfactual settings. Therefore variation in η_j arises from variation in ψ alone.

Recall that the observable characteristics in the scoring rule include shipping time, indicator variables for FBA Non-Amazon, FBM, and log feedback and rating differences, in addition to total price[3]. Since I am mostly concerned with assignment rules that weigh Amazon's characteristics differently relative to non-Amazon sellers, I will limit counterfactual ψ to those that vary the coefficients on FBA Non-Amazon, FBM, and shipping time,

[3]The coefficients from the corresponding Buy Box assignment rule with total price levels, as estimated in staticpaper as opposed to log total prices is reported in Table **??**.

and keep the remaining coefficients fixed.[4].

For any given ψ, I estimate welfare via the following procedure:

1. For every product page-day in our sample, for any seller $j, \in M$, compute
$$\eta_j(\psi) = \frac{X_j'\psi - X_{-j}'\psi}{\gamma} + \epsilon_j$$
given assigned ϵ_j, opponent $-j$.

2. For each seller $j \in M$ solve for $\underline{p}_j(\theta_j), \overline{p}_j(\theta_j)$ that maximize expected profits in equation 2.1, given calibrated θ_j, and $\eta_j(\psi), \eta_{-j}(\psi)$.

3. Identify the leader in the resulting cycle as the seller with the maximum resulting $\underline{p}_j(\theta_j) + \eta_j(\psi)$. Expected profits for each seller $j \in J$ are given by the weighted daily average.
$$\sum_p \omega_p(\lambda, \kappa, \chi, \underline{p}, \overline{p}, \eta)\pi_j(\mathbf{p}, \mathbf{X}, \mathbf{BB}, \beta, c_j)$$
where weights ω_p denote the daily average over each price vector and resulting Buy Box assignments given price-update, entry, and exit rates, and strategies resulting from $\overline{p}, \underline{p}, \eta$, taking non-Buy Box competitor prices as in the data, as they do not respond to variation in preferencing rules ψ.

4. Expected consumer welfare for a representative consumer is given by the weighted daily average
$$\sum_p \omega_p(\lambda, \kappa, \chi, \underline{p}, \overline{p}, \eta) \log\left(\sum_{j \in J} exp(X_j'\beta - \alpha p_j + \phi + \xi_j) + 1\right)$$
given the demand parameters estimated in Chapter 1, where the observable characteristics include Buy Box assignment.

[4]Recall that most variation among FBM offers is due to differences in shipping time, and any additional variation for Prime products is mostly captured by variation in shipping time.

For each counterfactual ψ of interest, for sellers j with a *set* of calibrated θ_j, I conduct the estimation procedure for both the θ_j in which x_j is the lowest, and corresponding c_j is the lowest, and under the highest x_j and corresponding highest c_j in the set.

I estimate counterfactual welfare under the current rule, the static-optimal rule that coincides with consumer demand parameters as estimated in chapter 1, and under a rule in which there is no additional preferencing for Amazon offers relative to non-Amazon offers. Under this rule, coefficients on FBA Non-Amazon, FBM are suppressed to zero, and shipping time coincides with the current rule. Additionally, in order to investigate rules that improve upon consumer or producer welfare relative to the current rule, I estimate ψ that optimize over average consumer, producer, and non-Amazon producer surplus, via a Bayesian optimizer, where the resulting measures weight each producer and product page-day equally.

Any resulting relative welfare measures critically depend on the composition of the sample. For computational feasibility, I take a representative sub-sample of approximately 210 product page-days out of over 6700 product page-days between 2018 - 2022, on which I observe price cycling, and fix this sample in all counterfactuals. As expected, the composition of Buy Box competitors is very similar to the overall sample shown in Table 2.3. Amazon competes on 52 percent of these product page-days, and on only two of them, Amazon competes with an FBM seller. On all remaining product-page days, Amazon competes with FBA sellers. The remaining sample is composed of approximately a total of 12 percent of product page-days on which two FBA sellers compete, and 30 percent of product page-days on which two FBM sellers compete. Nowhere in the sample does an FBA seller compete with an FBM seller, as is exceptionally rare in price cycling data overall. Therefore, any differences in welfare will largely arise from differences in scoring rules between Amazon and FBA sellers, and shipping time variation between FBM sellers. There is little to no variation in shipping time in our sample among FBA sellers or between Amazon and FBA sellers.

For any given counterfactual rule ψ, I estimate a 95 % Bootstrapped confidence interval of the average difference between given welfare measures under the current rule less the counterfactual rule of interest.

2.6 Results

The coefficients that govern the various counterfactual preferencing rules ψ are given in Table 2.5. For each counterfactual, 95 % Bootstrapped confidence intervals for differences in expected consumer surplus, expected producer surplus, Amazon producer surplus, non-Amazon producer surplus, as well as mean differences in cycle lower bounds, upper bounds, and Buy Box prices are reported in the relevant tables. Measures are also reported for the sub-samples in which Amazon competes with a non-Amazon offer (Amazon sample), two FBA offers compete (FBA sample), and two FBM offers compete (FBM sample).

I focus my discussion of overall results on the first set of counterfactuals in which all calculations are made fixing θ_j at the lowest x_j value and lowest corresponding c_j value for sellers j whose calibrated primitives consist of a set. The results for the remaining counterfactuals where these θ_j are fixed at the highest x_j and highest corresponding c_j value, are also reported in Tables 2.12, 2.14, 2.15, 2.16, 2.17, 2.18. The general results are all very similar. The magnitudes of the counterfactuals mainly differ due to the fact that under the second set of counterfactuals, when x_j approaches one, j's profit-maximizing upper bound is close to monopoly prices, and hence her overall optimal bounds may be less sensitive to changes in η than they are under the first set of counterfactuals.

2.6.1 Policy Proposal Counterfactuals

I first estimate the distribution of mean welfare differences between the current rule and a proposed policy rule in which all Amazon and non-Amazon offers are scored identically, conditional on shipping time and rating and review metrics, referred to as the "No-Preferencing" Rule in Table 2.6. Under this preferencing rule, both FBA and FBM sellers have *higher* η_j values relative to the current rule, while in the counterfactual, Amazon most often maintains a lower η_j value relative to their opponent, due to its higher ratings and reviews. Perhaps counterintuitively, in this setting on average, Amazon's producer surplus improves at a cost to non-Amazon producers. Most of this overall relative cost to non-Amazon sellers is due to the fact under the current rule, Amazon is often the follower in the pricing cycle and remains the follower under the counterfactual rule, since the change in η values is often not significant enough to induce them to become the leader. Amazon's higher η_j values, relative to the current rule, imply that their prices are *lower*, for the same leader bounds, which implies *lower* profits for non-Amazon producers, similar to the simulated setting described in the comparative statics section. These lower average Buy Box prices benefit consumer surplus. When Amazon is the leader in the current setting and remains the leader, non-Amazon sellers may improve on profits for the same reason, but not in all cases.

In the smaller sub-sample where Amazon is the follower in the current setting and becomes the leader in the counterfactual setting, the changes in welfare are somewhat ambiguous. In some cases, Amazon may improve from becoming the leader, while in other cases they do not. Amazon becomes the follower in only a very small sub-sample of cases, and their welfare improves at the cost of non-Amazon producers, who are induced to become the leader with lower Buy Box prices than under the current rule.

There are no changes to welfare in the non-Amazon sample when ψ alone is varied, due to the fact that there are no differences in relative scores between competitors in these samples, since FBA and FBM sellers only compete against sellers with the same fulfillment method.

Results for the "Static Optimal" rule, or a preferencing rule that corresponds with the consumer demand parameters estimated in chapter 1 are reported in Table 2.7. For this rule, since demand parameters are estimated using log total prices, I estimate the relevant η values using the average log total price of the opponent. As reported in chapter 1, this amounts to a 1.5 % decrease in η_j for an FBA offer relative to an Amazon offer, an 11.3 % decrease in η_j for FBM offers relative to Amazon offers, and a 1.9 % decrease in η_j per additional day of shipping time. In this setting, non-Amazon producers, in particular, FBM sellers, are scored *less* harshly relative to Amazon than under the current rule, but differences in resulting η are smaller relative to the "No-Preferencing" rule. Overall, the effects on consumer surplus are somewhat ambiguous, while producer surplus improvements are mostly awarded to Amazon at the cost to non-Amazon producers. Similar to the "No-Preferencing" counterfactual environment, most of the Amazon sub-sample consists of Amazon remaining the follower in both the current and counterfactual setting. The decrease in Amazon's relative η_j implies *lower* Amazon prices, which imply lower profits for non-Amazon producers, given their optimal bounds do not significantly shift. In some cases, when Amazon is induced to become the leader in the counterfactual, some non-Amazon producers improve their profits, at the cost to consumers.

Since η were simulated somewhat differently in this setting, FBA sellers' relative η values shift slightly, but overall counterfactual changes are mild. In the FBM sub-sample, higher shipping time values are scored *less* severely than under the status quo. In the majority of this sub-sample, the leader in the counterfactual remains fixed, and any relative improvements in producer surplus from one seller are largely off-set by decreases to the other. When the leader changes, the counterfactual that weighs shipping time less comes at a cost to consumer surplus, potentially for enabling a higher shipping time seller to obtain higher prices.

Given that the "No-Preferencing" rule largely comes at a cost to non-Amazon sellers by inducing *lower* Amazon prices when non-Amazon sellers lead, and reducing shipping time penalties for FBM sellers may come at a cost to consumer surplus, the question arises: what are the preferencing rules that improve upon consumer, producer, and in particular, non-Amazon producer surplus? I select three such preferencing rules resulting from a Bayesian optimizer that optimizes over each respective expected welfare function under the full sample.

2.6.2 Additional Preferencing Rules

Under the "CS Optimal" rule as shown in Table 2.8, improvements in consumer surplus mainly come from the sub-sample in which Amazon becomes the leader in the counterfactual.

Under the counterfactual preferencing rule, FBA sellers get a significantly *higher* score relative to Amazon sellers, all else equal. When Amazon is the leader under both settings, their significantly higher η_j value relative to the current rule implies lower own prices relative to the status quo, while non-Amazon sellers obtain higher Buy Box prices for the same Amazon bounds. While overall Buy Box prices may increase, in cases where the preferencing rule implies lower prices for Amazon, consumer welfare improves. In the much smaller sub-sample where Amazon does not become the leader in the counterfactual, the cost to non-Amazon sellers is more significant. While Amazon prices lower for any same leader bounds, overall Buy Box prices increase due to more significant changes in bounds.

In the non-Amazon FBM sample, higher shipping times are punished more severely than under the current rule. When the leader remains the same in the counterfactual, if followers have lower shipping times, they obtain higher Buy Box prices in the counterfactual setting, which improves one seller's surplus at the expense of the other. When the leader changes in the counterfactual, consumer welfare also decreases due to increased Buy Box prices. Intuitively, a rule that weighs shipping time at an absolute value closer to zero, while keeping the leader in tact, should reduce Buy Box prices and imply compensatory average producer surplus differences.

The "PS Optimal" rule shown in Table 2.9 is one such rule, where the shipping time coefficient is positive, and FBA and FBM coefficients are close to zero, as in the "No Preferencing" rule. When Amazon is the follower in both settings, this rule decreases average Buy Box prices overall, which provides some improvements to consumer surplus and Amazon surplus at the expense of non-Amazon surplus. When Amazon leads in both settings, non-Amazon welfare can improve from having a lower η_j. Improvements to consumer surplus appear to largely come from decreases in average Buy Box prices. In the non-Amazon FBM sample, the costs to consumer surplus relative to other preferencing rules that punish shipping time more harshly are less pronounced. While this ψ may induce lower prices for higher shipping times in some cases, there are off-setting effects from the relative decreases in prices from another.

Finally, I investigate a rule that improves upon non-Amazon surplus, the "non-Amazon PS" rule shown in Table 2.10. Under this rule, Amazon has an even higher scoring advantage, and an even *lower* η_j value relative to the current rule. When Amazon is the follower in a cycle in both settings, this implies higher prices for Amazon sellers, which improves non-Amazon surplus at a mild cost to consumer surplus. When Amazon is the leader under both the current rule and the counterfactual rule, while non-Amazon sellers price *lower* for the same Amazon bounds, which in some cases improves non-Amazon surplus, and in others, not. Shipping time for FBM sellers is again weighed similar to the "CS Optimal" rule. Some of these producers therefore benefit from inducing higher opponent prices when opponents have higher shipping times, but these effects come at a cost to consumer welfare via higher Buy Box prices.

2.6.3 Dynamic Pricing Strategies

Overall, in a setting with price cycles, the welfare results of different preferencing rules may vary greatly from the environment studied in chapter 1. In general, in counterfactuals where non-Amazon sellers remain the *leader* in the pricing cycle, preferencing rules that imply *higher* relative scores for Amazon may increase Amazon's prices and improve non-Amazon producer surplus at the expense of consumer welfare. By contrast, when Amazon remains the leader, higher relative scores for non-Amazon sellers imply *lower* Amazon prices, improving consumer welfare, and in some cases, improving non-Amazon surplus by allowing for higher non-Amazon prices. The trade-offs between various rules depend on the extent to which the change in the preferencing rule is significant enough to shift each player's optimally-chosen boundaries and induce a change in leadership. Since differences in sellers on product pages where two non-Amazon sellers compete are mostly due to shipping time, scoring rules that punish shipping time more harshly are not necessarily better for consumers, as this may imply higher prices when leaders remain fixed. Consumer welfare in these settings may improve if lower score differences in shipping time imply lower prices.

Therefore, when the platform allows for pricing strategies that induce price cycles, both consumer welfare and non-Amazon welfare can improve when Amazon leads in the cycle, and commits to a *lower* score premium relative to non-Amazon sellers. This improves non-Amazon welfare when the lower score premium decreases Amazon prices, and allows non-Amazon sellers to sustain prices that are high enough to improve profits, while the off-setting effects of these price differences may improve consumer surplus on net. Such a preferencing rule may not coincide with that derived in chapter 1 and may require an even larger relative score premium to non-Amazon sellers. In other circumstances, non-Amazon sellers benefit when Amazon's score premium implies *higher* Amazon prices for any fixed bounds.

Of course, the optimal coefficients that would govern such rules depend critically on the composition of sellers. While the current rules do not appear to induce non-FBA sellers to compete with Amazon, and non-FBA sellers to compete with FBA sellers, an alternative rule that ensures higher profits for the lesser-preferred seller can improve this seller's profits while also improving consumer surplus. The long run effects of inducing such competition may benefit the platform overall.

Lastly, in order to estimate the magnitudes of the relative costs of dynamic pricing strategies, I simulate locally envy-free prices and Buy Box assignments using the consumer-optimal rule in the static setting from chapter 1 in Table 2.11. The estimated costs to consumer welfare relative to this setting are far larger than those of any other counterfactual preferencing rule under price cycling. These improvements are substantial in both Amazon and non-Amazon sub-samples, and reduce Buy Box prices by several dollars on average. Those who obtain the Buy Box obtain positive profits, and under this rule obtain an average markup of $2.83.

Given the substantial improvements to consumer welfare from shutting down the price-updating mechanism, a potential improvement relative to the "No Preferencing" rule or any other of the counterfactuals studied above, would be to only allow sellers to update prices once the Buy Box owner has exited, presumably when they have run out of stock.

If Amazon exits often enough, this could induce positive non-Amazon producer surplus while ensuring substantially lower Buy Box prices. Given that many sellers currently fail to actively compete with Amazon for the Buy Box, the rule that induces static prices and assignments whenever the set of sellers changes could potentially induce more competition, although careful analysis of this setting should be relegated to further study.

Table 2.5: *Counterfactual Preferencing Rule Coefficients.*

	Current Rule	No-Preferencing Rule	CS rule	PS rule	Non-Amazon Rule
Shipping Time	−0.96	−0.96	−4.98	0.13	−5.00
FBA Non-Amazon	−0.68	0.00	1.96	0.00	−1.21
FBM	−1.04	0.00	−4.85	0.05	−3.15

Table 2.6: *Counterfactual Estimation, Current Rule Less No-Preferencing Rule 95 % Bootstrap CI (Min x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader Lower Bound	Leader Upper Bound	Average Buy Box Price	Fraction of Sample
Full	[-0.38, -0.06]	[-0.17, 0.01]	[-0.48, -0.12]	[0.04, 0.27]	[-0.23, -0.02]	[-0.10, 0.04]	[0.09, 0.25]	1.00
Amazon	[-0.74, -0.15]	[-0.33, 0.00]	[-0.51, -0.12]	[0.02, 0.28]	[-0.42, -0.03]	[-0.19, 0.07]	[0.19, 0.48]	0.52
Non-Amazon	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.44
Amazon Leader Pre-Post	[-1.73, -0.01]	[-0.56, 0.37]	[-1.17, 0.11]	[-0.51, 1.22]	[0.15, 1.57]	[0.07, 0.68]	[-0.19, 0.69]	0.02
Amazon Leader Pre- Follower Post	[-3.22, -2.47]	[0.08, 0.68]	[-0.04, -0.02]	[0.10, 0.73]	[3.99, 4.01]	[0.39, 0.41]	[0.41, 0.70]	0.01
Amazon Follower Pre-Post	[-0.84, -0.21]	[-0.39, -0.01]	[-0.59, -0.15]	[0.07, 0.31]	[-0.22, -0.02]	[-0.14, -0.01]	[0.23, 0.56]	0.42
Amazon Follower Pre- Leader Post	[-0.61, 1.27]	[-0.38, 0.29]	[-0.35, 0.70]	[-0.47, 0.09]	[-2.50, -1.77]	[-1.35, 0.70]	[-0.39, 0.38]	0.07
FBA	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Same	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Different	–	–	–	–	–	–	–	0.00
FBM	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.32
FBM Leader Same	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.32
FBM Leader Different	–	–	–	–	–	–	–	0.00

Table 2.7: *Counterfactual Estimation, Current Rule Less Static-Optimal Rule 95 % Bootstrap CI (Min x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader $\underline{p}$	Leader $\overline{p}$	Mean BB Price	Fraction of Sample
Full	[-0.11, 0.27]	[-0.27, -0.08]	[-0.69, -0.27]	[0.14, 0.41]	[-0.43, -0.19]	[-0.22, -0.04]	[0.03, 0.20]	1.00
Amazon	[-0.41, 0.16]	[-0.38, -0.02]	[-0.71, -0.30]	[0.12, 0.45]	[-0.54, -0.15]	[-0.20, 0.05]	[0.14, 0.42]	0.52
Non-Amazon	[0.15, 0.57]	[-0.23, -0.06]	[0.00, 0.00]	–	[-0.46, -0.12]	[-0.34, -0.07]	[-0.18, 0.01]	0.44
Amazon Leader Pre-Post	[-1.65, -0.33]	[-0.18, 0.30]	[-1.58, -0.14]	[0.31, 1.44]	[0.51, 2.11]	[-0.11, 0.50]	[0.09, 0.69]	0.04
Amazon Leader Pre-Follower Post	–	–	–	–	–	–	–	0.00
Amazon Follower Pre-Post	[-0.77, -0.16]	[-0.43, 0.03]	[-0.85, -0.22]	[0.16, 0.53]	[-0.28, -0.04]	[-0.29, -0.05]	[0.19, 0.55]	0.35
Amazon Follower Pre- Leader Post	[0.56, 1.66]	[-0.57, -0.13]	[-0.49, -0.13]	[-0.29, 0.16]	[-1.80, -1.10]	[-0.23, 0.49]	[-0.23, 0.20]	0.12
FBA	[0.07, 0.43]	[-0.06, 0.05]	–	–	[-0.24, -0.06]	[-0.20, -0.05]	[-0.06, 0.01]	0.12
FBA Leader Same	[0.00, 0.13]	[-0.10, 0.08]	–	–	[0.00, 0.00]	[-0.00, 0.00]	[-0.11, -0.01]	0.08
FBA Leader Different	[0.15, 1.06]	[-0.03, 0.05]	–	–	[-0.61, -0.34]	[-0.51, -0.28]	[-0.03, 0.07]	0.04
FBM	[0.10, 0.69]	[-0.31, -0.09]	–	–	[-0.55, -0.14]	[-0.43, -0.05]	[-0.22, 0.04]	0.32
FBM Leader Same	[-0.05, 0.60]	[-0.35, -0.08]	–	–	[-0.31, 0.08]	[-0.28, 0.02]	[-0.26, 0.09]	0.22
FBM Leader Different	[0.07, 1.30]	[-0.38, -0.07]	–	–	[-1.30, -0.39]	[-0.95, -0.07]	[-0.31, 0.09]	0.10

Table 2.8: *Counterfactual Estimation, Current Rule Less CS Rule 95 % Bootstrap CI (Min x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader Lower Bound	Leader Upper Bound	Average Buy Box Price	Fraction of Sample
Full	[-0.22, 0.40]	[-0.01, 0.40]	[-0.99, -0.27]	[0.51, 1.26]	[-0.04, 0.37]	[-0.03, 0.36]	[-1.37, -0.69]	1.00
Amazon	[-1.16, -0.03]	[-0.09, 0.55]	[-1.06, -0.30]	[0.52, 1.32]	[-0.27, 0.36]	[-0.25, 0.31]	[-1.24, -0.27]	0.52
Non-Amazon	[0.61, 1.18]	[-0.12, 0.37]	[0.00, 0.00]	–	[-0.03, 0.57]	[0.09, 0.59]	[-1.78, -0.92]	0.44
Amazon Leader Pre-Post	[-4.27, 1.47]	[0.69, 3.81]	[-2.11, 0.63]	[0.32, 5.84]	[0.33, 3.78]	[-0.14, 1.28]	[-1.42, 1.39]	0.04
Amazon Leader Pre- Follower Post	[-2.30, -2.30]	[-0.60, -0.60]	[-7.31, -7.31]	[6.71, 6.71]	[0.42, 0.42]	[0.11, 0.11]	[0.29, 0.29]	0.00
Amazon Follower Pre-Post	[-0.95, 0.66]	[0.37, 2.89]	[-1.07, 1.71]	[0.26, 2.62]	[-2.53, -0.40]	[-1.76, -0.23]	[-8.09, -1.43]	0.06
Amazon Follower Pre- Leader Post	[-1.12, 0.10]	[-0.41, 0.17]	[-1.17, -0.37]	[0.25, 0.99]	[-0.31, 0.42]	[-0.20, 0.46]	[-0.52, -0.03]	0.41
FBA	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Same	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Different	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.00
FBM	[0.88, 1.64]	[-0.16, 0.50]	–	–	[-0.05, 0.77]	[0.13, 0.78]	[-2.51, -1.30]	0.32
FBM Leader Same	[0.79, 1.61]	[-0.25, 0.53]	–	–	[-0.63, -0.01]	[-0.03, 0.57]	[-2.62, -1.18]	0.27
FBM Leader Different	[0.70, 2.54]	[-0.39, 1.09]	–	–	[2.94, 4.44]	[0.36, 3.33]	[-2.50, -1.16]	0.05

Table 2.9: *Counterfactual Estimation, Current Rule Less PS Rule 95 % Bootstrap CI (Min x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader Lower Bound	Leader Upper Bound	Average Buy Box Price	Fraction of Sample
Full	[-0.39, 0.06]	[-0.24, -0.01]	[-0.47, -0.08]	[0.02, 0.26]	[-0.42, -0.11]	[-0.17, 0.04]	[0.02, 0.25]	1.00
Amazon	[-0.85, -0.26]	[-0.30, 0.02]	[-0.49, -0.08]	[0.02, 0.29]	[-0.43, -0.01]	[-0.16, 0.22]	[0.22, 0.51]	0.52
Non-Amazon	[-0.08, 0.50]	[-0.33, 0.04]	[0.00, 0.00]	–	[-0.58, -0.11]	[-0.31, -0.06]	[-0.24, 0.04]	0.44
Amazon Leader Pre-Post	[-1.76, 0.03]	[-0.58, 0.35]	[-1.21, 0.24]	[-0.55, 1.24]	[0.16, 1.70]	[0.07, 0.65]	[-0.25, 0.67]	0.03
Amazon Leader Pre- Follower Post	[-3.22, -2.48]	[0.08, 0.68]	[-0.04, -0.02]	[0.10, 0.72]	[3.40, 4.01]	[0.39, 0.41]	[0.42, 0.71]	0.01
Amazon Follower Pre-Post	[-0.78, -0.21]	[-0.31, 0.06]	[-0.54, -0.07]	[0.07, 0.32]	[-0.22, -0.02]	[-0.16, -0.01]	[0.24, 0.56]	0.40
Amazon Follower Pre- Leader Post	[-1.93, 0.84]	[-0.63, 0.13]	[-0.58, 0.51]	[-0.43, 0.07]	[-2.23, -1.19]	[-0.78, 1.61]	[-0.30, 0.78]	0.08
FBA	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Same	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Different	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.00
FBM	[-0.10, 0.72]	[-0.46, 0.04]	–	–	[-0.81, -0.17]	[-0.50, -0.10]	[-0.31, 0.05]	0.32
FBM Leader Same	[-0.27, 0.55]	[-0.53, 0.21]	–	–	[-0.34, 0.32]	[-0.14, 0.11]	[-0.33, 0.17]	0.20
FBM Leader Different	[-0.10, 1.43]	[-0.58, -0.03]	–	–	[-1.81, -0.77]	[-1.02, -0.32]	[-0.48, 0.05]	0.12

Table 2.10: *Counterfactual Estimation, Current Rule Less Non-Amazon PS Rule 95% Bootstrap CI (Min x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader Lower Bound	Leader Upper Bound	Average Buy Box Price	Fraction of Sample
Full	[0.08, 0.53]	[0.00, 0.31]	[0.11, 0.52]	[-0.22, -0.03]	[0.11, 0.49]	[0.06, 0.32]	[-1.14, -0.59]	1.00
Amazon	[-0.56, 0.15]	[-0.01, 0.36]	[0.09, 0.53]	[-0.25, -0.04]	[0.11, 0.55]	[-0.05, 0.16]	[-0.82, -0.10]	0.51
Non-Amazon	[0.61, 1.18]	[-0.14, 0.40]	[0.00, 0.00]	[0.00, 0.54]	[-0.08, 0.59]	[0.07, 0.59]	[-1.83, -0.86]	0.49
Amazon Leader Pre-Post	[-0.49, 1.49]	[0.05, 0.30]	[-0.30, 0.51]	[-0.22, 0.35]	[-2.23, 0.00]	[-0.74, 0.00]	[-0.78, 0.07]	0.01
Amazon Leader Pre- Follower Post	[-2.70, -0.26]	[-0.42, 0.30]	[-0.02, 0.08]	[-0.47, 0.30]	[3.66, 4.57]	[0.04, 1.42]	[-0.08, 0.53]	0.03
Amazon Follower Pre-Post	[-0.56, 0.22]	[-0.01, 0.43]	[0.13, 0.60]	[-0.26, -0.05]	[-0.01, 0.24]	[-0.06, 0.12]	[-0.95, -0.11]	0.47
Amazon Follower Pre- Leader Post	–	–	–	–	–	–	–	0.00
FBA	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Same	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Different	–	–	–	–	–	–	–	0.00
FBM	[0.86, 1.63]	[-0.17, 0.49]	–	–	[-0.12, 0.77]	[0.12, 0.82]	[-2.48, -1.34]	0.33
FBM Leader Same	[0.77, 1.61]	[-0.26, 0.44]	–	–	[-0.59, -0.02]	[-0.02, 0.53]	[-2.61, -1.20]	0.27
FBM Leader Different	[0.79, 2.49]	[-0.28, 1.04]	–	–	[2.95, 4.32]	[0.41, 3.18]	[-2.52, -1.32]	0.05

Table 2.11: *Counterfactual Estimation, Current Rule Less Static Equilibrium Bootstrap 95% CI (Min x).*

	Confidence Interval
Full Sample CS difference	[-15.95, -13.07]
Full Sample PS difference	[-4.92, -2.27]
Full Sample Amazon PS difference	[-2.07, -0.17]
Full Sample Non-Amazon PS difference	[-4.41, -1.65]
Full Sample Average Buybox Price difference	[4.43, 5.77]
Amazon Sample CS difference	[-16.70, -12.57]
Amazon Sample PS differences	[-5.33, -2.20]
Amazon Sample Amazon PS difference	[-2.20, -0.10]
Amazon Sample Non-Amazon PS difference	[-4.11, -1.41]
Amazon Sample Average Buybox Price difference	[3.95, 5.39]
Non-Amazon Sample CS difference	[-16.65, -12.57]
Non-Amazon Sample PS difference	[-6.02, -1.36]
Non-Amazon Sample Average Buybox Price difference	[4.63, 7.11]

Table 2.12: *Counterfactual Estimation, Current Rule Less No Preferencing Rule 95% Bootstrap CI (Max x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader Lower Bound	Leader Upper Bound	Average Buy Box Price	Fraction of Sample
Full	[-0.16, 0.31]	[-0.04, 0.08]	[-0.18, 0.07]	[0.00, 0.21]	[-0.25, -0.03]	[-0.06, 0.07]	[0.03, 0.17]	1.00
Amazon	[-0.27, 0.64]	[-0.07, 0.16]	[-0.19, 0.09]	[-0.01, 0.22]	[-0.49, -0.05]	[-0.10, 0.12]	[0.06, 0.32]	0.52
Non-Amazon	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.44
Amazon Leader Pre-Post	[-1.73, 0.95]	[0.09, 0.92]	[-0.66, 0.25]	[-0.08, 1.32]	[0.00, 0.79]	[0.00, 0.17]	[-0.30, 0.73]	0.04
Amazon Leader Pre- Follower Post	[-3.22, -1.05]	[0.09, 0.11]	[-0.54, -0.02]	[0.11, 0.64]	[3.97, 4.51]	[-0.02, 0.39]	[0.70, 0.75]	0.01
Amazon Follower Pre-Post	[-0.73, 0.36]	[-0.09, 0.17]	[-0.21, 0.13]	[-0.02, 0.20]	[-0.27, 0.05]	[-0.16, 0.00]	[0.14, 0.45]	0.39
Amazon Follower Pre- Leader Post	[1.47, 3.41]	[-0.32, 0.00]	[-0.23, 0.13]	[-0.28, 0.06]	[-2.21, -1.54]	[-0.32, 0.84]	[-0.49, -0.15]	0.08
FBA	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Same	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Different	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.00
FBM	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.32
FBM Leader Same	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.20
FBM Leader Different	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12

Table 2.13: *Counterfactual Estimation, Current Rule Less Static-Optimal Rule 95 % Bootstrap CI (Max x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader Lower Bound	Leader Upper Bound	Average Buy Box Price	Fraction of Sample
Full	[-0.08, 0.44]	[-0.19, 0.00]	[-0.31, -0.01]	[-0.04, 0.2]	[-0.51, -0.22]	[-0.14, 0.03]	[0.05, 0.28]	1.0
Amazon	[0.03, 0.90]	[-0.24, 0.05]	[-0.33, -0.03]	[-0.04, 0.21]	[-0.69, -0.25]	[-0.15, 0.12]	[0.08, 0.42]	0.519
Non-Amazon	[-0.37, 0.15]	[-0.20, 0.06]	[-0.0, 0.0]	–	[-0.44, 0.042]	[-0.20, 0.04]	[-0.07, 0.20]	0.438
Amazon Leader Pre-Post	[-0.43, 2.59]	[-0.75, 0.45]	[-0.76, -0.06]	[-0.51, 1.15]	[0.0, 1.30]	[-0.11, 0.15]	[-1.11, 0.28]	0.047
Amazon Leader Pre- Follower Post	[-1.1, -1.1]	[0.12, 0.12]	[-0.55, -0.55]	[0.67, 0.67]	[4.51, 4.51]	[-0.02, -0.02]	[0.69, 0.69]	0.0047
Amazon Follower Pre-Post	[-0.44, 0.68]	[-0.25, 0.18]	[-0.43, 0.07]	[0.01, 0.29]	[-0.40, -0.06]	[-0.24, -0.02]	[0.21, 0.67]	0.29
Amazon Follower Pre- Leader Post	[0.37, 1.88]	[-0.40, -0.06]	[-0.28, 0.02]	[-0.25, 0.08]	[-1.64, -1.14]	[-0.22, 0.48]	[-0.14, 0.27]	0.176
FBA	[-0.10, 0.17]	[-0.16, -0.00]	–	–	[-0.22, 0.0]	[-0.01, 0.0]	[-0.08, 0.07]	0.119
FBA Leader Same	[-0.12, 0.09]	[-0.15, 0.00]	–	–	[0.0, 0.0]	[-0.0, 0.0]	[-0.03, 0.08]	0.109
FBA Leader Different	[-0.03, 1.27]	[-0.52, 0.03]	–	–	[-1.77, -0.17]	[-0.09, -0.00]	[-0.71, 0.00]	0.009
FBM	[-0.53, 0.15]	[-0.26, 0.11]	–	–	[-0.50, 0.07]	[-0.28, 0.05]	[-0.09, 0.27]	0.319
FBM Leader Same	[-0.68, 0.16]	[-0.23, 0.31]	–	–	[-0.16, 0.45]	[-0.03, 0.10]	[-0.14, 0.33]	0.238
FBM Leader Different	[-0.49, 0.55]	[-0.61, -0.14]	–	–	[-1.93, -0.95]	[-1.06, 0.01]	[-0.13, 0.31]	0.08

Table 2.14: *Counterfactual Estimation, Current Rule Less Static-Optimal Rule 95% Bootstrap CI (Max x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader Lower Bound	Leader Upper Bound	Average Buy Box Price	Fraction of Sample
Full	[-0.08, 0.44]	[-0.19, 0.00]	[-0.31, -0.01]	[-0.04, 0.20]	[-0.51, -0.22]	[-0.14, 0.03]	[0.05, 0.28]	1.00
Amazon	[0.03, 0.90]	[-0.24, 0.05]	[-0.33, -0.03]	[-0.04, 0.21]	[-0.69, -0.25]	[-0.15, 0.12]	[0.08, 0.42]	0.52
Non-Amazon	[-0.37, 0.15]	[-0.20, 0.06]	[0.00, 0.00]	–	[-0.44, 0.04]	[-0.20, 0.04]	[-0.07, 0.20]	0.44
Amazon Leader Pre-Post	[-0.43, 2.59]	[-0.75, 0.45]	[-0.76, -0.06]	[-0.51, 1.15]	[0.00, 1.30]	[-0.11, 0.15]	[-1.11, 0.28]	0.05
Amazon Leader Pre- Follower Post	[-1.10, -1.10]	[0.12, 0.12]	[-0.55, -0.55]	[0.67, 0.67]	[4.51, 4.51]	[-0.02, -0.02]	[0.69, 0.69]	0.00
Amazon Follower Pre-Post	[-0.44, 0.68]	[-0.25, 0.18]	[-0.43, 0.07]	[0.01, 0.29]	[-0.40, -0.06]	[-0.24, -0.02]	[0.21, 0.67]	0.29
Amazon Follower Pre- Leader Post	[0.37, 1.88]	[-0.40, -0.06]	[-0.28, 0.02]	[-0.25, 0.08]	[-1.64, -1.14]	[-0.22, 0.48]	[-0.14, 0.27]	0.18
FBA	[-0.10, 0.17]	[-0.16, 0.00]	–	–	[-0.22, 0.00]	[-0.01, 0.00]	[-0.08, 0.07]	0.12
FBA Leader Same	[-0.12, 0.09]	[-0.15, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[-0.03, 0.08]	0.11
FBA Leader Different	[-0.03, 1.27]	[-0.52, 0.03]	–	–	[-1.77, -0.17]	[-0.09, 0.00]	[-0.71, 0.00]	0.01
FBM	[-0.53, 0.15]	[-0.26, 0.11]	–	–	[-0.50, 0.07]	[-0.28, 0.05]	[-0.09, 0.27]	0.32
FBM Leader Same	[-0.68, 0.16]	[-0.23, 0.31]	–	–	[-0.16, 0.45]	[-0.03, 0.10]	[-0.14, 0.33]	0.24
FBM Leader Different	[-0.49, 0.55]	[-0.61, -0.14]	–	–	[-1.93, -0.95]	[-1.06, 0.01]	[-0.13, 0.31]	0.08

Table 2.15: *Counterfactual Estimation, Current Rule Less CS Rule 95% Bootstrap CI (Max x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader Lower Bound	Leader Upper Bound	Average Buy Box Price	Fraction of Sample
Full	[-0.12, 0.55]	[0.21, 0.50]	[-0.59, 0.00]	[0.32, 1.12]	[-0.15, 0.40]	[-0.15, 0.40]	[-1.37, -0.75]	1.00
Amazon	[-0.64, 0.72]	[0.17, 0.69]	[-0.57, -0.04]	[0.38, 1.14]	[-0.37, 0.27]	[-0.30, 0.18]	[-1.35, -0.42]	0.52
Non-Amazon	[0.23, 0.62]	[0.08, 0.65]	[0.00, 0.00]	–	[-0.12, 0.84]	[0.09, 0.47]	[-1.75, -0.79]	0.44
Amazon Leader Pre-Post	[-2.49, 2.47]	[0.88, 3.38]	[-2.42, -0.18]	[1.22, 5.63]	[0.28, 2.81]	[-0.12, 0.91]	[-2.08, 0.49]	0.05
Amazon Leader Pre- Follower Post	–	–	–	–	–	–	–	0.00
Amazon Follower Pre-Post	[-1.05, 0.16]	[0.17, 1.06]	[-1.74, 0.49]	[0.11, 2.61]	[-1.46, -0.18]	[-1.55, -0.17]	[-8.58, -2.08]	0.05
Amazon Follower Pre- Leader Post	[-0.62, 0.88]	[-0.06, 0.46]	[-0.42, 0.11]	[0.06, 0.60]	[-0.50, 0.12]	[-0.30, 0.25]	[-0.60, -0.08]	0.41
FBA	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Same	[0.00, 0.00]	[0.00, 0.00]	–	–	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	0.12
FBA Leader Different	–	–	–	–	–	–	–	0.00
FBM	[0.36, 0.84]	[0.13, 0.87]	–	–	[-0.15, 1.15]	[0.13, 0.63]	[-2.32, -1.18]	0.32
FBM Leader Same	[0.44, 0.92]	[0.09, 0.80]	–	–	[-0.48, 0.02]	[0.00, 0.13]	[-2.20, -1.12]	0.29
FBM Leader Different	[-1.16, 0.88]	[-1.12, 3.84]	–	–	[3.19, 9.63]	[2.15, 4.12]	[-3.73, -0.04]	0.03

Table 2.16: *Counterfactual Estimation, Current Rule Less PS Rule 95 % Bootstrap CI (Max x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader Lower Bound	Leader Upper Bound	Average Buy Box Price	Fraction of Sample
Full	[-0.06, 0.42]	[-0.19, 0.00]	[-0.32, -0.02]	[-0.06, 0.19]	[-0.52, -0.22]	[-0.15, 0.04]	[0.06, 0.27]	1.00
Amazon	[0.03, 0.92]	[-0.24, 0.04]	[-0.34, -0.02]	[-0.04, 0.21]	[-0.69, -0.29]	[-0.15, 0.14]	[0.08, 0.40]	0.52
Non-Amazon	[-0.37, 0.13]	[-0.22, 0.06]	[0.00, 0.00]	–	[-0.45, 0.03]	[-0.19, 0.04]	[-0.07, 0.21]	0.44
Amazon Leader Pre-Post	[-0.47, 2.43]	[-0.80, 0.45]	[-0.75, -0.07]	[-0.67, 1.17]	[0.00, 1.23]	[-0.11, 0.15]	[-1.09, 0.24]	0.05
Amazon Leader Pre- Follower Post	[-1.10, -1.10]	[0.12, 0.12]	[-0.55, -0.55]	[0.67, 0.67]	[4.51, 4.51]	[-0.02, -0.02]	[0.69, 0.69]	0.00
Amazon Follower Pre-Post	[-0.44, 0.65]	[-0.22, 0.17]	[-0.43, 0.06]	[0.02, 0.29]	[-0.39, -0.05]	[-0.25, -0.02]	[0.23, 0.66]	0.29
Amazon Follower Pre- Leader Post	[0.32, 1.96]	[-0.38, -0.04]	[-0.28, 0.03]	[-0.24, 0.08]	[-1.66, -1.19]	[-0.19, 0.48]	[-0.16, 0.25]	0.18
FBA	[-0.08, 0.18]	[-0.16, -0.00]	–	–	[-0.22, 0.00]	[-0.01, 0.00]	[-0.09, 0.07]	0.12
FBA Leader Same	[-0.12, 0.10]	[-0.16, 0.01]	–	–	[0.00, 0.00]	[0.00, 0.00]	[-0.03, 0.09]	0.11
FBA Leader Different	[-0.03, 1.27]	[-0.52, 0.03]	–	–	[-1.77, -0.17]	[-0.09, 0.00]	[-0.71, 0.00]	0.01
FBM	[-0.51, 0.16]	[-0.26, 0.11]	–	–	[-0.57, 0.02]	[-0.27, 0.06]	[-0.10, 0.28]	0.32
FBM Leader Same	[-0.68, 0.19]	[-0.23, 0.26]	–	–	[-0.21, 0.47]	[-0.03, 0.10]	[-0.15, 0.34]	0.24
FBM Leader Different	[-0.49, 0.54]	[-0.60, -0.12]	–	–	[-1.97, -0.90]	[-0.99, 0.04]	[-0.13, 0.31]	0.08

Table 2.17: *Counterfactual Estimation, Current Rule Less Non-Amazon PS Rule 95 % Bootstrap CI (Max x).*

Sample	CS	PS	Amazon PS	Non-Amazon PS	Leader Lower Bound	Leader Upper Bound	Average Buy Box Price	Fraction of Sample
Full	[0.21, 0.71]	[0.04, 0.31]	[-0.08, 0.13]	[-0.11, 0.07]	[0.06, 0.61]	[0.07, 0.33]	[-1.18, -0.59]	1.00
Amazon	[0.05, 0.93]	[-0.09, 0.08]	[-0.09, 0.14]	[-0.12, 0.07]	[0.05, 0.61]	[-0.04, 0.35]	[-0.83, -0.19]	0.51
Non-Amazon	[0.27, 0.63]	[0.08, 0.64]	–	–	[-0.07, 0.84]	[0.08, 0.47]	[-1.72, -0.82]	0.44
Amazon Leader Pre-Post	[-0.57, 5.13]	[-0.69, 0.83]	[-0.49, 0.32]	[-0.65, 0.70]	[-1.06, 0.00]	[-2.43, 0.03]	[-2.47, 0.13]	0.02
Amazon Leader Pre- Follower Post	[-1.90, -0.41]	[-0.37, 0.13]	[-0.02, 0.30]	[-0.43, -0.06]	[3.75, 8.50]	[0.28, 4.91]	[0.43, 1.43]	0.03
Amazon Follower Pre-Post	[0.06, 0.98]	[-0.09, 0.09]	[-0.11, 0.14]	[-0.11, 0.08]	[-0.03, 0.09]	[-0.00, 0.12]	[-0.89, -0.23]	0.46
Amazon Follower Pre- Leader Post	–	–	–	–	–	–	–	0.00
FBA	[-0.00, 0.00]	[-0.00, 0.00]	–	–	[-0.00, 0.00]	[-0.00, 0.00]	[-0.00, 0.00]	0.12
FBA Leader Same	[-0.00, 0.00]	[-0.00, 0.00]	–	–	[-0.00, 0.00]	[-0.00, 0.00]	[-0.00, 0.00]	0.12
FBA Leader Different	–	–	–	–	–	–	–	0.00
FBM	[0.37, 0.85]	[0.09, 0.86]	–	–	[-0.12, 1.18]	[0.13, 0.64]	[-2.29, -1.15]	0.33
FBM Leader Same	[0.43, 0.91]	[0.10, 0.77]	–	–	[-0.48, 0.02]	[-0.00, 0.13]	[-2.28, -1.12]	0.29
FBM Leader Different	[-0.79, 1.20]	[-1.24, 3.73]	–	–	[2.70, 9.34]	[2.01, 4.16]	[-4.35, 0.07]	0.03

Table 2.18: *Counterfactual Estimation, Current Rule Less Static Equilibrium Bootstrap 95% CI (Max x).*

	Confidence Interval
Full Sample CS difference	[-13.67, -11.46]
Full Sample PS difference	[-5.37, -2.93]
Full Sample Amazon PS difference	[-0.51, 0.70]
Full Sample Non-Amazon PS difference	[-5.49, -2.58]
Full Sample Average Buybox Price difference	[4.23, 5.42]
Amazon Sample CS difference	[-13.43, -10.19]
Amazon Sample PS differences	[-5.24, -2.33]
Amazon Sample Amazon PS difference	[-0.57, 0.70]
Amazon Sample Non-Amazon PS difference	[-5.21, -2.49]
Amazon Sample Average Buy Box Price difference	[3.63, 5.04]
Non-Amazon Sample CS difference	[-15.43, -11.96]
Non-Amazon Sample PS difference	[-6.95, -2.34]
Non-Amazon Sample Average Buy Box Price difference	[4.55, 6.52]

2.7 Discussion

This chapter, the first detailed empirical study of pricing cycles on Amazon, suggests that dynamic pricing strategies that result in cycles come at a significant cost to consumers through substantially higher prices relative to endogeneous prices under static price competition. The costs of these pricing strategies, relative to the static environment, do not include the potential long-run effects of reducing competition against Amazon. Importantly, the impact of preferencing rules on various welfare measures differs significantly from what has been observed in static competition environments, as studied in chapter 1. When Amazon leads pricing cycles, consumers and non-Amazon producers benefit from even *less* self-preferencing than is prevalent in the static-optimal rule. In other cases, when non-Amazon sellers lead cycles, these sellers benefit from preferencing rules that induce *higher* Amazon prices for any fixed own-prices.

Some might wonder why Amazon self-preferences at all, if they obtain benefits from inducing more third-party sellers to sell on their platform. When sellers compete with dynamic pricing strategies, preferencing themselves *more* may benefit third-party sellers without significantly harming Amazon's profits. Amazon appears to preference sufficiently to induce higher profits than they would by forcing themselves to price lower against competitors when they lead pricing cycles. In these dynamic settings, Amazon may simply be choosing a preferencing rule that increases their own profits subject to allowing FBA sellers to obtain sufficiently high profits to continue competing for the Buy Box, at least on some product pages.

There may be many other preferencing mechanisms, outside of the class of linear scoring rules, that improve upon welfare for both consumer and non-Amazon producers, while allowing for the dynamic pricing strategies that are observed on the platform. For example, these preferencing mechanisms might include more dynamic information about sellers' prices. In any case, it is not obvious that any rule would be preferable to the consumer-optimal static setting studied in chapter 1.

A promising policy proposal is that which would allow sellers to update their prices only after the current Buy Box owner exits the market. This would induce a marketplace very similar to the static scoring auction studied in chapter 1, where the auction would only take place when the item *must* be re-assigned. In essence, the platform would auction off the seller's current stock of items to be sold at their Buy Box-winning price. By requiring Amazon to sell its products at a markup proportional to the extent to which consumers prefer their products relative to those of competitors, Amazon may limit product-page entry alongside non-Amazon sellers. Furthermore, it may motivate otherwise similar sellers to invest in practices that improve ratings and reviews, as these factors would influence the markups they can obtain in the auction. Any potential effects of this policy on entry, exit, and market composition warrant further investigation.

Chapter 3

Quantum Games of Complete Information

3.1 Introduction

Quantum game theory expands the study of classical game theory to environments in which players have access to quantum resources, most notably entanglement, typically made possible by a quantum computer. Quantum game theory therefore permits the study of outcomes when players with access to quantum information optimally interact, and how these outcomes may differ– in particular improve upon– their classical counterparts.

Studying these environments is not only relevant for anticipating future scenarios in which quantum technologies are widely accessible, but additionally, contributes to the broader effort of classifying problems that can be more efficiently solved with quantum versus classical computers. Indeed, previous studies of quantum games have been mapped to well-known quantum algorithms, including Grover's search algorithm (Meyer, 1999).

In this chapter, I focus on arguably the simplest and most foundational class of games: 2×2 complete information normal form games. I derive necessary and sufficient conditions under which equilibria in a subclass of corresponding quantum settings yield Pareto improvements relative to all classical correlated equilibria.

3.1.1 Related Literature

The field of quantum game theory was initiated by Meyer (1999), who demonstrated a strategic advantage for a player with access to quantum strategies over a classically restricted opponent in a two-player zero-sum setting.

Since this initial study, most papers have focused on demonstrating existence results of quantum advantage in isolated complete information settings (Eisert et al., 1999; Benjamin and Hayden, 2001; van Enk and Pike, 2002; Ikeda and Aoki, 2022; La Mura, 2005), with the exception of Landsburg (2011), who provides conditions under which *mixed* quantum strategies provide improvements in the particular basis studied in Eisert et al. (1999).

In general, these previous papers study various quantum frameworks without establishing a unified mapping between classical and quantum games—one that would allow systematic variation across both measurement bases and initial quantum states. These choices are critical to the outcomes of the games and are the settings which most realistically, players will face decisions over when confronted with the possibility of deploying a game in a quantum environment.

Instead, I focus on *general* quantum games with a canonical initial state and selected basis. In particular, given that the choice of the initial state and basis may limit the feasible distributions over outcomes and therefore players' deviations given opponent strategies, contrary to the suggestions of previous work, I show that not all equilibria in complete information quantum environments coincide with correlated equilibria, and under the conditions derived in Theorems 7 and 8, Pareto improve upon correlated equilibria. I then discuss the related design questions that players reasonably face when opting into a quantum game framework, and provide a participation constraint along with a selection criterion of an initial state and basis for a quantum game.

Other previous work has expanded the initial studies of quantum games to settings of incomplete information (Brunner and Linden, 2013; Brandenburger and La Mura, 2016; Milchtaich, 2024). I relegate a more detailed study of incomplete information settings, including necessary conditions for advantage in quantum settings relative to a classical baseline, to a further publication. Indeed, these necessary conditions for advantage in non-degenerate incomplete information settings may not precisely coincide with those studied in the complete information settings of this chapter.

Essentially, quantum games of incomplete information study payoff advantages when outcome distributions violate the existence of non-local hidden variables. In these settings, quantum formulations may allow for more flexible distributions *across* states of Nature relative to those of restricted classes of correlated distributions, when signals depend upon private information. This chapter instead studies what entanglement alone affords players relative to classical correlation in particular quantum environments, whereby restricting the feasible distributions over outcomes is *advantageous* to players within a single degenerate state of Nature.

3.2 Quantum Games of Complete Information

This chapter studies the quantum formulations of normal form complete information games, defined in the section that follows.

3.2.1 Normal Form Games

A game is given by a tuple $\Gamma = \{N, (S_i)_{i \in N}, (\pi_i)_{i \in N}\}$, where N is a finite set of players, a set of strategies $S_1 \times S_2 ... \times S_N$ where S_i is a finite set of strategies for player i, and a set of payoff functions $\pi_1 \times \pi_2 ... \times \pi_N$, where $\pi_i : S \to \mathbb{R}$ denotes player i's payoff function.

In this chapter, I will be concerned with 2×2 games of complete information that can be represented by the following matrix.

	s_2	s_2'
s_1	$\pi_1(s_1, s_2), \pi_2(s_1, s_2)$	$\pi_1(s_1, s_2'), \pi_2(s_1, s_2')$
s_1'	$\pi_1(s_1', s_2), \pi_2(s_1', s_2)$	$\pi_1(s_1', s_2'), \pi_2(s_1', s_2')$

Figure 3.1: 2×2 *complete information game matrix.*

Equilibrium payoff results in quantum settings will be compared with classical correlated equilibria (CE), inclusive of mixed strategy Nash equilibria (MSNE) of the corresponding normal form game of complete information.

3.2.2 Corresponding Quantum Games

I next define a mapping from classical $N \times N$ normal-form games of complete information—where each of the N players has N pure strategies—to a corresponding quantum game. For relevant mathematical background, see B.0.1.

The quantum game is additionally defined by a fixed orthonormal basis

$$\{|e_1\rangle, \dots, |e_N\rangle\},$$

for an $N \times N$-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^{N \otimes N}$. Each basis element $\left|e_j\right\rangle$ can be expressed as a tensor product:

$$\left|e_j\right\rangle = \left|e_j^1\right\rangle \otimes \cdots \otimes \left|e_j^N\right\rangle.$$

Each basis vector $\left|e_j\right\rangle$ corresponds with a classical pure strategy profile $(s_1, \dots, s_N) \in S_1 \times \cdots \times S_N$, where player i's corresponding pure strategy will be denoted by $s_i(\left|e_j^i\right\rangle)$.

The game is additionally defined by an initial shared quantum state $|\psi\rangle \in \mathcal{H}$, written as a superposition over basis states:

$$|\psi\rangle = \sum_j \alpha_j \left|e_j\right\rangle,$$

with complex coefficients $\alpha_j \in \mathbb{C}$ satisfying $\sum_j |\alpha_j|^2 = 1$.

We will be concerned with bases in which each basis element $\left|e_j\right\rangle$ is composed of qubits, and therefore $\left|e_j^i\right\rangle \in \{0, 1\}\ \forall i$.

Each player $i \in \{1, \dots, N\}$ is space-like separated and has access only to their own subsystem, or qubits, given by the i-th register of the tensor product. For instance, if $|\psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \in \mathbb{C}^2 \otimes \mathbb{C}^2$, then player 1 holds the first qubit of the entangled state, and player 2 holds the second qubit. That is, player 1 accesses the first register in the tensor product, corresponding to the left-hand qubit in each basis element $|00\rangle, |11\rangle$.

Each player selects a quantum strategy given by a unitary operator $U_i \in \mathcal{U}(N)$, which acts on their respective subsystem. The joint strategy profile is represented by the joint unitary:

$$U = U_1 \otimes \cdots \otimes U_N.$$

This operator is applied to the initial state, resulting in the quantum state:

$$|\sigma\rangle = U |\psi\rangle = (U_1 \otimes \cdots \otimes U_N) |\psi\rangle.$$

A projective measurement is then performed in the given basis $\{|e_1\rangle, \ldots, |e_N\rangle\}$, collapsing the state $|\sigma\rangle$ to basis element $|e_j\rangle$ with probability:

$$P(|e_j\rangle) = |\langle e_j|\sigma\rangle|^2.$$

Players receive payoffs $\pi_i(s_i(|e_j^i\rangle), s_{-i}(|e_j^{-i}\rangle))$ according to the measured basis element $|e_j\rangle$[1].

This permits the following equilibrium definition:

Definition 2. *A* Nash Equilibrium in Quantum Strategies (QNE) *is a profile of unitary operators* $(U_1, \ldots, U_N) \in \mathcal{U}(N)^N$ *such that, for each player i and for all alternative unitaries* $U_i' \in \mathcal{U}(N)$,

$$\sum_j P(|e_j\rangle)\, \pi_i(s_i(|e_j^i\rangle), s_{-i}(|e_j^{-i}\rangle)) \geqslant \sum_j P'(|e_j\rangle)\, \pi_i(s_i(|e_j^i\rangle), s_{-i}(|e_j^{-i}\rangle)),$$

given initial state $|\psi\rangle$ *and basis* $\{|e_1\rangle, ..., |e_N\rangle\}$ *where:*

- $P(|e_j\rangle) = \left|\langle e_j|(U_1 \otimes \cdots \otimes U_N) |\psi\rangle\rangle\right|^2$,
- $P'(|e_j\rangle) = \left|\langle e_j|(U_1 \otimes \cdots \otimes U_{i-1} \otimes U_i' \otimes U_{i+1} \otimes \cdots \otimes U_N) |\psi\rangle\rangle\right|^2$,
- *and* $s_i(|e_j^i\rangle) \in S_i$ *is the pure strategy derived from* $|e_j^i\rangle$.

[1] A discussion of how the measurement outcome practically maps to *implemented* pure strategies is relegated to section 3.5, along with a discussion of the generation of the initial state $|\psi\rangle$.

3.2.3 Computational Basis

In this chapter, I will focus on measurements that occur in the *computational basis*, which in 2×2 games is given by $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$, and therefore will map to corresponding pure strategy profiles as follows:

$$(s_1, s_2) \leftrightarrow |00\rangle$$

$$(s_1, s_2') \leftrightarrow |01\rangle$$

$$(s_1', s_2) \leftrightarrow |10\rangle$$

$$(s_1', s_2') \leftrightarrow |11\rangle$$

I focus on this basis for several reasons. First, it is simple in the respect that each player can perform the relevant measurements on her own qubits, eliminating the need for an additional third party or referee to perform measurements on behalf of the players. Games in this basis may therefore be lower-cost relative to games in alternative bases.

Secondly, as I will demonstrate in the sections that follow, in this basis, for particular choices of initial states $|\psi\rangle$, players can effectively restrict the sets of feasible distributions over corresponding pure strategies to their advantage, given *any* choice of unitary operators. In other bases, including those chosen in the well-known Eisert et al. (1999) game, these same advantages relative to CE only arise if players' decisions over unitaries are additionally restricted– which is less reasonable in realistic settings where players are free to choose any unitary operators to apply to their respective qubits.

3.2.4 Outcome Distributions in the Computational Basis

In general, a unitary operator in $SU(2)$, the special unitary operators of degree 2 may be written as

$$U_i = \begin{bmatrix} e^{i\alpha_i}\cos\theta_i & e^{i\beta_i}\sin\theta_i \\ -e^{-i\beta_i}\sin\theta_i & e^{-i\alpha_i}\cos\theta_i \end{bmatrix}$$

for $\theta \in \mathbf{R}$, $\alpha, \beta \in \mathbf{R}$, where θ is a *rotation* and α, β are relative *phases*.

Any initial state $|\psi\rangle$ in $\mathcal{H} = \mathbb{C}^{2\otimes 2}$ can be written as

$$|\psi\rangle = a\,|00\rangle + b\,|01\rangle + c\,|10\rangle + d\,|11\rangle$$

$\exists a, b, c, d, \in \mathbb{C}, |a|^2 + |b|^2 + |c|^2 + |d|^2 = 1.$

Therefore, the state resulting from both player's unitaries acting upon their respective qubits in the initial state is given by

$$U_a \otimes U_b(a\,|00\rangle + b\,|01\rangle + c\,|10\rangle + d\,|11\rangle) =$$

$$\begin{aligned}
&\Big(ae^{i(\alpha_1+\alpha_2)}\cos\theta_1\cos\theta_2 + be^{i(\alpha_1+\beta_2)}\cos\theta_1\sin\theta_2 + ce^{i(\beta_1+\alpha_2)}\sin\theta_1\cos\theta_2 + de^{i(\beta_1+\beta_2)}\sin\theta_1\sin\theta_2\Big)\,|00\rangle \\
&+\Big(-ae^{i(\alpha_1-\beta_2)}\cos\theta_1\sin\theta_2 + be^{i(\alpha_1-\alpha_2)}\cos\theta_1\cos\theta_2 - ce^{i(\beta_1-\beta_2)}\sin\theta_1\sin\theta_2 + de^{i(\beta_1-\alpha_2)}\sin\theta_1\cos\theta_2\Big)\,|01\rangle \\
&+\Big(-ae^{i(\alpha_2-\beta_1)}\sin\theta_1\cos\theta_2 - be^{i(\beta_2-\beta_1)}\sin\theta_1\sin\theta_2 + ce^{i(\alpha_2-\alpha_1)}\cos\theta_1\cos\theta_2 + de^{i(\beta_2-\alpha_1)}\cos\theta_1\sin\theta_2\Big)\,|10\rangle \\
&+\Big(ae^{-i(\beta_1+\beta_2)}\sin\theta_1\sin\theta_2 - be^{-i(\beta_1+\alpha_2)}\sin\theta_1\cos\theta_2 - ce^{-i(\alpha_1+\beta_2)}\cos\theta_1\sin\theta_2 + de^{-i(\alpha_1+\alpha_2)}\cos\theta_1\cos\theta_2\Big)\,|11\rangle
\end{aligned}$$

where the probability of measuring a given basis element is given by the amplitude squared of the corresponding coefficients above.

It is known that independent unitaries acting on an initial pure, separable product state, in other words, in which $ad - bc = 0$, yield separable probabilities over measurement outcomes. Therefore, an initial pure state must be entangled and not a product state for quantum strategies to result in distributions over corresponding pure strategies not obtainable by classical MSNE.[2]

Given an initial, pure entangled state, the sets of feasible measurement outcomes, and therefore possible *deviations* given opponent strategies are restricted. For example, for an initial state in which $a = d = 1/2, b = 1/\sqrt{2}, c = 0$, and $\theta_2 = 0$, for no choices of $\theta_1, \alpha_1, \beta_1$ can player 1 ensure a zero probability measurement outcome on both basis elements $|10\rangle, |11\rangle$. Essentially, these parameters eliminate her ability to guarantee an outcome from playing a corresponding pure strategy with probability one, which is always feasible in any CE. In

[2]Note that players can generate measurement outcome distributions that correspond with any mixed strategy distribution by choosing one of $a, b, c, d, = 1$ and choosing θ_1, θ_2 appropriately.

this way, the initial entangled state acts as a commitment device.[3]

It is for this reason that a QNE may not be guaranteed to exist for any initial choice of entangled state, and any distributions over measurement outcomes given by a QNE may not correspond with a CE. If we consider settings in which players must explicitly agree to play a corresponding quantum game, the decision over the basis and initial state are *design* choices, chosen by the players themselves.

I am therefore interested in classifying games of complete information in the computational basis in which there is a QNE that outperforms any CE inclusive of MSNE, and therefore whose payoffs cannot be generated classically.

3.2.5 Maximally-Entangled Initial States $|\psi\rangle$

I focus on games with a *maximally entangled* initial state. Up to local unitaries, any such state can be written in the form:

$$|\psi\rangle = a\,|00\rangle + d\,|11\rangle, \quad \text{with } a = \pm\tfrac{1}{\sqrt{2}},\ d = \pm\tfrac{1}{\sqrt{2}},\ b = c = 0,$$

or

$$|\psi\rangle = b\,|01\rangle + c\,|10\rangle, \quad \text{with } b = \pm\tfrac{1}{\sqrt{2}},\ c = \pm\tfrac{1}{\sqrt{2}},\ a = d = 0.$$

These maximally entangled states are also known as Bell states.

When local unitaries are applied to $|\psi\rangle$, the computational basis measurement probabilities depend only on the sum of the players' rotation angles θ_1 and θ_2, for any $\alpha_1, \beta_1, \alpha_2, \beta_2$.

Specifically, for any maximally entangled initial state, the resulting measurement probabilities are either:

$$P(|00\rangle) = P(|11\rangle) = \tfrac{1}{2}\cos^2(\theta_1 + \theta_2),$$
$$P(|01\rangle) = P(|10\rangle) = \tfrac{1}{2}\sin^2(\theta_1 + \theta_2),$$

[3]Note that this differs from other setups like the Eisert et al. (1999) protocol which relies on restricted unitary operations. In the Prisoner's Dilemma example used in Eisert et al. (1999), players cannot commit to using these restricted unitaries as they have an incentive to deviate to those that admit higher probabilities on strictly dominant strategies.

or

$$P(|00\rangle) = P(|11\rangle) = \tfrac{1}{2}\sin^2(\theta_1 + \theta_2),$$
$$P(|01\rangle) = P(|10\rangle) = \tfrac{1}{2}\cos^2(\theta_1 + \theta_2),$$

depending on the choice of initial state.

Since all maximally entangled two-qubit pure states yield equivalent outcome distributions under local unitary transformations, I henceforth restrict attention—without loss of generality—to the canonical Bell state:

$$|\psi\rangle = \tfrac{1}{\sqrt{2}}(|00\rangle + |11\rangle).$$

Accordingly, I limit each player's strategy to choosing a unitary of the form

$$U_i = \begin{bmatrix} \cos\theta_i & -\sin\theta_i \\ \sin\theta_i & \cos\theta_i \end{bmatrix}$$

for $\theta_i \in [\frac{-\pi}{2}, \frac{\pi}{2}]$.

Under this model, the expected payoffs for each player i are given by:

$$\begin{aligned}\mathbb{E}[\pi_i(\theta_i, \theta_{-i})] = {} & \tfrac{1}{2}\cos^2(\theta_i + \theta_{-i})\pi_i(s_i(|0\rangle), s_{-i}(|0\rangle)) + \tfrac{1}{2}\sin^2(\theta_i + \theta_{-i})\pi_i(s_i(|0\rangle), s_{-i}(|1\rangle)) \\ & + \tfrac{1}{2}\sin^2(\theta_i + \theta_{-i})\pi_i(s_i(|1\rangle), s_{-i}(|0\rangle)) + \tfrac{1}{2}\cos^2(\theta_i + \theta_{-i})\pi_i(s_i(|1\rangle), s_{-i}(|1\rangle))\end{aligned}$$

Clearly, the resulting distributions over corresponding pure strategies cannot be generated by any MSNE. Given the commitment to such an initial state, players can *never* guarantee a resulting pure strategy distribution that would correspond with any MSNE, or any CE that does not follow the above form.

Note also that players could choose *mixed* unitary operators, density matrices applied to the corresponding density matrix given by the initial state $|\psi\rangle$. In this special case, since the resulting distributions are convex combinations of the distribution described above, any distribution generated by mixed unitaries can be generated by pure unitaries. Therefore, I restrict attention to pure unitary operators only.

3.3 Results on Maximally Entangled Initial States

All of the following results pertain to 2×2 games of complete information with a corresponding quantum game in the computational basis and a maximally entangled initial state.

Let $B_1 = \{|00\rangle, |11\rangle\}$ and $B_2 = \{|01\rangle, |10\rangle\}$ denote the two basis pairs that appear with equal probability under measurement of a maximally entangled state.

We define player i's preferred basis pair as:

$$\mathcal{B}_i = \begin{cases} B_1, & \text{if } \pi_i(s_i(|0\rangle), s_{-i}(|0\rangle)) + \pi_i(s_i(|1\rangle), s_{-i}(|1\rangle)) \geqslant \pi_i(s_i(|0\rangle), s_{-i}(|1\rangle)) + \pi_i(s_i(|1\rangle), s_{-i}(|0\rangle)), \\ B_2, & \text{otherwise.} \end{cases}$$

Theorem 4. *A Nash Equilibrium in Quantum Strategies (QNE) exists if and only if all players share the same preferred basis pair:*

$$\mathcal{B}_1 = \mathcal{B}_2 = \cdots = \mathcal{B}_N.$$

In other words, the pair of basis outcomes that maximizes the expected payoff must coincide across players.

Proof. First we show the necessary condition. Without loss of generality suppose that $\mathcal{B}_1 = B_1$ and $\mathcal{B}_2 = B_2$. For any θ_1, player 2 can maximize her payoff by choosing θ_2 appropriately to maximize $\sin(\theta_1 + \theta_2)^2$. In particular, for $\theta_1 \in [0, \pi/2]$, she can choose $\theta_2 = \pi/2 - \theta_1$, and for $\theta_1 \in [-\pi/2, 0)$, she can choose $\theta_2 = -\pi/2 - \theta_1$. An analogous argument holds for player 1. Some player has a profitable deviation for any distribution implied by θ_1, θ_2. Next we show the sufficient condition. Clearly if $\mathcal{B}_1 = \mathcal{B}_2$, then any choice of θ_1, θ_2 that implies a total probability of 1 on $\mathcal{B}_1 = \mathcal{B}_2$ is a QNE. $\square$

This does not require the game to be symmetric, but only that players' basis pair preferences agree. Since $\theta \in [\frac{-\pi}{2}, \frac{\pi}{2}]$, there are infinitely many unitaries that may coincide with a QNE, conditional on a QNE existing.

Theorem 5. *For any classical game, if there is a correlated equilibrium (CE) where the joint signal* *distribution is equally weighted on all elements corresponding with either B_1 or B_2, then there is a QNE of the corresponding quantum game where the distribution over measurement outcomes coincides with the signal distribution of the CE.*

Proof. Without loss of generality, suppose that there is a CE whose joint signal distribution places probability $\frac{1}{2}$ on each of (s_1, s_2) and (s_1', s_2'). Then it must be that $\pi_1(s_1, s_2) \geqslant \pi_1(s_1', s_2), \pi_2(s_1, s_2) \geqslant \pi_2(s_1, s_2'), \pi_1(s_1', s_2') \geqslant \pi_1(s_1, s_2'), \pi_2(s_1', s_2') \geqslant \pi_2(s_1', s_2)$. Taking the summation implies the sufficient condition for the corresponding QNE. $\square$

	P_h	P_l
P_h	$(2,2)$	$(0,3)$
P_l	$(3,0)$	$(0,0)$

Figure 3.2: *Bertrand Pricing Game.*

Consider, for example, the Bertrand Pricing game in Figure 3.2. The only MSNE of the game is one in which both players play P_l with probability one. The CE with a signal distribution that places probability $\frac{1}{2}$ on each of $(P_l, P_h), (P_h, P_l)$ also coincides with the distribution implied by a QNE. Therefore, the quantum version of the game can be viewed as an environment to induce CE without the use of a *public* randomizing device, when such CE exist. Rather, in the quantum game, the distributions over measurement outcomes are only known to those who directly observe them.

Theorem 6. *A Nash Equilibrium in Quantum Strategies (QNE) may generate probability distributions over pure strategies that do not correspond with any classical correlated equilibrium (CE).*

Proof. Consider the Prisoner's Dilemma shown in Figure 3.3. $\theta_1 = \theta_2 = \frac{\pi}{2}$ is a QNE in the computational basis with a maximally entangled initial state. This yields a distribution over corresponding pure strategies of $P(s_1', s_2) = P(s_1, s_2') = \frac{1}{2}$ which yield expected payoffs of 2.5 for both players. Clearly such a distribution cannot correspond with any CE, since in no CE will players place positive probability on strictly dominated strategies. $\square$

	s_2	s_2'
s_1	$(3,3)$	$(0,5)$
s_1'	$(5,0)$	$(1,1)$

Figure 3.3: *Prisoner's Dilemma.*

Indeed, such a QNE may exist even when all CE of the classical game are MSNE. In particular, by restricting the set of feasible deviations, players may place positive probability on strictly dominated strategies in a QNE, which may outperform any CE.

I now derive necessary and sufficient conditions under which a QNE Pareto improves upon any CE, inclusive of MSNE.

Theorem 7. *If a QNE yields expected payoffs which Pareto improves upon all CE inclusive of MSNE, then at least one player must have a strictly dominant strategy.*

Proof. Suppose that no player has a strictly dominant strategy. First, suppose instead that both players have weakly dominant strategies. Without loss of generality, suppose that player 1's weakly dominant strategy is s_1, $\pi_1(s_1, s_2) > \pi_1(s_1', s_2)$. Given this, the only feasible QNE results in probability $\frac{1}{2}$ over each of the measurement outcomes corresponding with $(s_1, s_2), (s_1', s_2')$. Given player 2 also has a weakly dominant strategy, the proposed measurement outcome distribution corresponds with an equilibrium only if either, s_2 is weakly dominant $\pi_2(s_1, s_2) > \pi_2(s_1, s_2')$, or s_2' is weakly dominant and $\pi_2(s_1', s_2') > \pi_2(s_1', s_2')$. In either case, a signal distribution with probability $\frac{1}{2}$ over each of $(s_1, s_2), (s_1', s_2')$ coincides with a CE.

If player 2 does not have a weakly dominant strategy, then under player 1's weakly dominant strategy, unitaries that imply probability $\frac{1}{2}$ over each of measurement outcomes corresponding with $(s_1, s_2), (s_1', s_2')$ are a QNE only if $\pi_2(s_1, s_2) \geqslant \pi_2(s_1, s_2')$, $\pi_2(s_1', s_2') \geqslant \pi_2(s_1', s_2)$. But then this distribution also coincides with a CE.

Lastly if no players have weakly dominant strategies, then unitaries that imply probability $\frac{1}{2}$ over each of the measurement outcomes corresponding with $(s_1, s_2), (s_1', s_2')$ are a QNE only if $\pi_1(s_1, s_2) \geqslant \pi_1(s_1', s_2)$, $\pi_1(s_1', s_2') \geqslant \pi_1(s_1, s_2')$, $\pi_2(s_1, s_2) \geqslant \pi_2(s_1, s_2')$, $\pi_2(s_1', s_2') \geqslant$

$\pi_2(s'_1, s_2)$. But then this distribution also coincides with a CE. The same logic applies for any distribution with probability $\frac{1}{2}$ over each of measurement outcomes corresponding with $(s_1, s'_2), (s'_1, s_2)$. $\square$

Theorem 8. *If at least one player has a strictly dominant strategy, and players' payoffs satisfy one of the following conditions, then there is a QNE that Pareto improves upon CE inclusive of MSNE. If player j's strictly dominant strategy is s_j*

(A) $\pi_{-j}(s_j, s_{-j}) = \pi_{-j}(s_j, s'_{-j})$, *and one of the following holds*

- **(A1)** $\frac{\pi_j(s_j,s_{-j})+\pi_j(s'_j,s'_{-j})}{2} \geqslant p\pi_j(s_j, s_{-j}) + (1-p)\pi_j(s_j, s'_{-j}) \forall p \in [0,1]$, *and* $\frac{\pi_{-j}(s_j,s_{-j})+\pi_{-j}(s'_1,s'_2)}{2} \geqslant \pi_{-j}(s_j, s_{-j})$, *where at least one equality is strict, and* $\pi_i(s_i, s_{-i}) + \pi_i(s'_i, s'_{-i}) \geqslant \pi_i(s_i, s'_{-i}) + \pi_i(s'_i, s_{-i})$ *for all i.*
- **(A2)** $\frac{\pi_j(s_j,s'_{-j})+\pi_j(s'_j,s_{-j})}{2} \geqslant p\pi_j(s_j, s_j) + (1-p)\pi_j(s_j, s'_{-j}) \forall p \in [0,1]$ *and* $\frac{\pi_{-j}(s'_j,s_{-j})+\pi_{-j}(s_j,s'_{-j})}{2} \geqslant \pi_{-j}(s_j, s_{-j})$ *where at least one equality is strict, and* $\pi_i(s_i, s'_{-i}) + \pi_i(s'_i, s_{-i}) \geqslant \pi_i(s_i, s_{-i}) + \pi_i(s'_i, s'_{-i})$ *for all i.*

(B) $\pi_{-j}(s_j, s_{-j}) > \pi_{-j}(s_j, s'_{-j})$, *and one of the following holds*

- **(B1)** $\pi_i(s'_i, s'_{-i}) \geqslant \pi_i(s_i, s_{-i})$ *and* $\pi_i(s_i, s_{-i}) + \pi_i(s'_i, s'_{-i}) \geqslant \pi_i(s_i, s'_{-i}) + \pi_i(s'_i, s_{-i})$, *for all i and the first inequality is strict for some i.*
- **(B2)** $\frac{\pi_i(s'_1,s_2)+\pi_i(s_1,s'_2)}{2} \geqslant \pi_i(s_i, s_{-i})$ *for all i and* $\pi_i(s_i, s'_{-i}) + \pi_i(s'_i, s_{-i}) \geqslant \pi_i(s_i, s_{-i}) + \pi_i(s'_i, s'_{-i})$, *for all i and the first inequality is strict for some i.*

(C) $\pi_{-j}(s_j, s_{-j}) < \pi_{-j}(s_j, s'_{-j})$, *and one of the following holds*

- **(C1)** $\pi_i(s'_i, s_{-i}) \geqslant \pi_i(s'_i, s'_{-i})$ *and* $\pi_i(s_i, s'_{-i}) + \pi_i(s'_i, s_{-i}) \geqslant \pi_i(s_i, s_{-i}) + \pi_i(s'_i, s'_{-i})$, *for all i and the first inequality is strict for some i.*
- **(C2)** $\frac{\pi_i(s_i,s_{-i})+\pi_i(s'_i,s'_{-i})}{2} \geqslant \pi_i(s_i, s'_{-i})$ *for all i and* $\pi_i(s_i, s_{-i}) + \pi_i(s'_i, s'_{-i}) \geqslant \pi_i(s'_i, s_{-i}) + \pi_i(s_i, s'_{-i})$, *for all i and the first inequality is strict for some i.*

Proof. Without loss of generality, suppose player 1 has a strictly dominant strategy, s_1. Then the only feasible CE are also MSNE, since player 1 will never place positive probability on

s_1' under any signal distribution. First suppose $\pi_2(s_1, s_2) = \pi_2(s_1, s_2')$. Then any distribution over $(s_1, s_2), (s_1, s_2')$ is a MSNE. Under condition $(A1)$ any unitaries resulting in probability $\frac{1}{2}$ over each of the measurement outcomes corresponding with $(s_1, s_2), (s_1', s_2')$ are a QNE. Under condition $(A2)$ any unitaries resulting in probability $\frac{1}{2}$ over each of the measurement outcomes corresponding with $(s_1, s_2'), (s_1', s_2)$ are a QNE.

If $\pi_2(s_1, s_2) > \pi_2(s_1, s_2')$ then the only MSNE is (s_1, s_2). Under condition $(B1)$ any unitaries resulting in probability $\frac{1}{2}$ over each of the measurement outcomes corresponding with $(s_1, s_2), (s_1', s_2')$ are a QNE. Under condition $(B2)$, any unitaries resulting in probability $\frac{1}{2}$ over each of the measurement outcomes corresponding with $(s_1', s_2), (s_1, s_2')$ are a QNE.

Finally if $\pi_2(s_1, s_2) < \pi_2(s_1, s_2')$ then the only MSNE is (s_1, s_2'). Under condition $(C1)$ any unitaries resulting in probability $\frac{1}{2}$ over each of the measurement outcomes corresponding with $(s_1', s_2), (s_1, s_2')$ are a QNE. Under condition $(C2)$, any unitaries resulting in probability $\frac{1}{2}$ over each of the measurement outcomes corresponding with $(s_1, s_2), (s_1', s_2')$ are a QNE.

In all cases, the conditions ensure that all QNE Pareto improve upon MSNE. □

3.4 Choosing a Quantum Environment

If we consider settings in which players must *agree* to engage in a corresponding quantum game, then one can consider the choice of the basis $\{|e_1\rangle, ..., |e_N\rangle\}$ and initial state $|\psi\rangle$ as a *design* decision that players must jointly agree upon.

This naturally yields a participation constraint for a corresponding quantum game.

Definition 3. *A quantum game with a basis $\{|e_1\rangle, ..., |e_N\rangle\}$ and initial state $|\psi\rangle$ is individually rational (IR) if and only if there is a QNE $(U_1, \ldots, U_N) \in \mathcal{U}(N)^N$ such that*

$$\sum_j P(|e_j\rangle)\, \pi_i(s_i(\left|e_j^i\right\rangle), s_{-i}(\left|e_j^{-i}\right\rangle)) - c_{i\{|e_1\rangle,...,|e_N\rangle\}|\psi\rangle} \geqslant max_{(\omega_1,...,\omega_N),P(\omega_1,...,\omega_N)} \sum_{\omega_{-i}} \pi_i(s_i(\omega_i), s_{-i}(\omega_{-i}))P(\omega_i, \omega_{-i})$$

for all signals $(\omega_1, ..., \omega_N)$ under joint distributions $P(\omega_1, ..., \omega_N)$ which satisfy

$$\sum_{\omega_{-i}} \pi_i(s_i(\omega_i), s_{-i}(\omega_{-i}))P(\omega_i, \omega_{-i}) \geqslant \sum_{\omega_{-i}} \pi_i(s_i'(\omega_i), s_{-i}(\omega_{-i}))P(\omega_i, \omega_{-i}) \forall i, \omega_i, s_i'$$

and where $c_{i\{|e_1\rangle,...,|e_N\rangle\}|\psi\rangle}$ are i's costs for the quantum environment under $\{|e_1\rangle, ..., |e_N\rangle\}, |\psi\rangle$, relative to the classical environment.

In other words, a quantum game is IR if and only if there is a QNE that weakly outperforms all CE, accounting for additional costs incurred for engaging in the quantum game.

Conditional on there existing an IR quantum game, players require a condition under which to *select* a particular quantum game corresponding to a basis and initial state decision.

Definition 4. *A QNE corresponding with a basis $\{|e_1\rangle, ..., |e_N\rangle\}$, initial state $|\psi\rangle$ and unitaries $(U_1, \ldots, U_N)$ is renegotiation-proof if and only if, for all i*

$$\sum_j P(|e_j\rangle)\,\pi_i(s_i(\left|e_j^i\right\rangle), s_{-i}(\left|e_j^{-i}\right\rangle)) - c_{i\{|e_1\rangle,...,|e_N\rangle\}|\psi\rangle} \geqslant \sum_j \tilde{P}(|\tilde{e}_j\rangle)\,\pi_i(s_i(\left|\tilde{e}_j{}^i\right\rangle), s_{-i}(\left|\tilde{e}_j{}^{-i}\right\rangle)) - c_{i\{|\tilde{e}_1\rangle,...,|\tilde{e}_N\rangle\}|\tilde{\psi}\rangle}$$

for any QNE corresponding with any alternative basis, $\{|\tilde{e}_1\rangle, ..., |\tilde{e}_N\rangle\}$, initial state $|\tilde{\psi}\rangle$, and unitaries $(\tilde{U}_1, \ldots, \tilde{U}_N)$.

The renegotation-proof criterion requires that a QNE weakly outperforms any other QNE for all players, given the costs associated with quantum environments.

Claim 1. *When a classical game satisfies the conditions of Theorem 7 and Theorem 8, and $c_{i\{|e_1\rangle,...,|e_N\rangle\}|\psi\rangle} = 0$ for all i under the computational basis and a maximally entangled state $|\psi\rangle$, there is an individually rational, renegotiation-proof QNE.*

Proof. The proof follows directly from Theorems 7 and 8. □

For games that satisfy the necessary and sufficient conditions studied in section 3.3, there is a QNE in the computational basis with a maximally entangled state that strictly outperforms any CE, and therefore guarantees the existence of an IR, renegotiation-proof QNE when the costs associated with the computational basis and maximally entangled environment are sufficiently low.

Conjecture 9. *For the Prisoner's Dilemma, the QNE in the computational basis with a maximally entangled state is the IR, renegotiation-proof QNE among all QNE in the computational basis.*

3.5 Practical Implementations

3.5.1 Enforceability of Measurement Outcomes

Critical to the necessary and sufficient conditions of Theorems 7 and 8 for a QNE in the computational basis to Pareto improve upon CE is the enforceability of measurement outcomes. Meaning, that if players measure individual qubits locally, they cannot implement pure strategies that differ from the corresponding realized measurement outcomes. If this were possible, there would clearly be no QNE that would not correspond with CE.

Enforceability can be implemented in one of two ways. First, as in other settings (Eisert et al., 1999), players can delegate both measurements and pure strategy implementation to a third-party referee. Essentially, the referee must be trusted to correctly conduct measurements and play the corresponding pure strategies on the player's behalves. Ensuring referee confidence may require additional protocols (Fitzsimons and Kashefi, 2017).

Secondly, players can delegate the *implementation* of their measurement outcomes to corresponding strategies, while performing measurements themselves. Since players can observe their own measurement outcomes and the implemented strategies, they will observe when the implementation is untrusted. This would require a hybrid workflow designed to enable users to perform quantum measurements locally while delegating classical implementations in a verifiable, observable manner. This requirement is very reasonable in practice, as it is common for programs and platforms to enable quantum measurement outcomes—whose irreversibility is foundational (Nielsen and Chuang, 2010)—to directly map to classical feed-forward automation routines (Walther et al., 2005; IBM Quantum, 2025).

Enforceability has other important implications. These conditions ensure that payoffs typically achievable in repeated settings (i.e. under the Folk Theorem) may be sustained in static settings without the requirement of dynamic punishment strategies. In classes of games in which there are Pareto improvements in quantum environments relative to CE, more *collusive* equilibria may be sustainable with the absence of reputation effects.

3.5.2 Generating Initial States

Additionally, upon jointly committing to an initial state $|\psi\rangle$, players must be able to reasonably verify that the initial state generated coincides with the jointly selected state. If not, resulting strategy distributions may not coincide with a QNE, or players may possibly improve their own payoffs at the expense of their opponent by generating an alternative initial state.

In scenarios where players generate the entangled state themselves in a one-shot game, full verification via repeated measurement is infeasible, as is common in procedures that validate entanglement distributions via Bell tests or quantum state tomography (Nielsen and Chuang, 2010). Nonetheless, players can reasonably certify correctness by agreeing upon and executing a known, symmetric preparation procedure, similar to mechanisms used in quantum bit commitment and secure multi-party protocols (Kent, 2012; Crépeau et al., 2002), which deter strategic mispreparation and ensure consistency without requiring trust or repeated measurements. For example, each player can prepare an entangled state and send one half to their opponent, along with a cryptographic commitment to the description or measurement basis of their prepared state. A shared random coin flip determines which state is used for play, and the other player reveals their committed information for verification.

3.6 Discussion

The future availability of quantum devices has substantial implications for marketplace design and welfare. For classes of games with quantum counterparts that outperform all correlated equilibria, players can achieve payoffs achievable by repeated interaction, but without any long-term commitments or punishment strategies. Detecting these potentially collusive outcomes may be particularly challenging if players are able to deploy these strategies in one-shot settings.

In other cases where QNE coincide with CE, quantum game formulations provide an alternative function, in allowing players to correlate on strategies *privately* in lieu of the assistance of a public randomizing device in a classical counterpart. A much more detailed discussion of coordination under private information in quantum settings is relegated to a further publication on quantum games of incomplete information.

This chapter is merely a very brief introduction to a much broader frontier of research at the intersection of multiple fields. In general, bridging the fields of game theory, quantum physics, and computing opens an entire spectrum of new research that sheds light on not only the welfare implications of allowing for any interaction that could reasonably be facilitated by a quantum computer, and the classes of problems more efficiently solved by quantum computer relative to classical computers, but also questions within foundational physics that may benefit from the study of equilibria in game theory, multi-agent optimization, and structural estimation.

References

Aguiar, L., J. Waldfogel, and S. Waldfogel (2021). Playlisting favorites: Measuring platform bias in the music industry. *International Journal of Industrial Organization 78*, 102765.

Amazon Seller University (2024). How to automate pricing on amazon. `https://www.youtube.com/watch?v=otWnDx2UdwY`. Accessed: December 11, 2024.

Anderson, S. P. and Ö. Bedre-Defolie (2021). Hybrid platform model. Technical report, CEPR discussion paper no. DP16243.

Armstrong, M. and J. Zhou (2011). Paying for prominence. *The Economic Journal 121*(556), F368–F395.

Asker, J. and E. Cantillon (2008). Properties of scoring auctions. *The RAND Journal of Economics 39*(1), 69–85.

Asker, J., C. Fershtman, and A. Pakes (2021). Artificial intelligence and pricing: The impact of algorithm design. Technical report, National Bureau of Economic Research.

Athey, S. and G. Ellison (2011). Position auctions with consumer search. *The Quarterly Journal of Economics 126*(3), 1213–1270.

Aura (2024). Repricing interface demo. `https://www.youtube.com/watch?v=qKLwKIgiUEA`. Accessed: December 11, 2024.

Banchio, M., A. Mehta, and A. Perlroth (2024). Auctions with dynamic scoring. *arXiv preprint arXiv:2403.11022*.

Banchio, M. and A. Skrzypacz (2022). Artificial intelligence and auction design. *arXiv preprint*.

Benjamin, S. C. and P. M. Hayden (2001). Multiplayer quantum games. *Physical Review A 64*(3), 030301.

Benner, M. J. and J. Waldfogel (2016). The song remains the same? technological change and positioning in the recorded music industry. *Strategy Science 1*(3), 129–147.

Bergemann, D. and A. Bonatti (2011). Targeting in advertising markets: implications for offline versus online media. *The RAND Journal of Economics 42*(3), 417–443.

Bergemann, D. and M. Said (2010). Dynamic auctions: A survey.

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics 25*(2), 242–262.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, 841–890.

Bourreau, M. and G. Gaudin (2022). Streaming platform and strategic recommendation bias. *Journal of Economics & Management Strategy 31*(1), 25–47.

Brandenburger, A. and P. La Mura (2016). Team decision problems with classical and quantum signals. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 374*(2058), 20150096.

Brown, Z. Y. and A. MacKay (2021). Competition in pricing algorithms. Technical report, National Bureau of Economic Research.

Brunner, N. and N. Linden (2013). Connection between bell nonlocality and bayesian game theory. *Nature communications 4*(1), 2057.

Brynjolfsson, E., Y. Hu, and M. D. Smith (2003). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management science 49*(11), 1580–1596.

Calvano, E., G. Calzolari, V. Denicolo, and S. Pastorello (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review 110*(10), 3267–97.

Calvano, E., G. Calzolari, V. Denicoló, and S. Pastorello (2021). Algorithmic collusion with imperfect monitoring. *International Journal of Industrial Organization 79*, 102712.

Che, Y.-K. (1993). Design competition through multidimensional auctions. *The RAND Journal of Economics*, 668–680.

Chevalier, J. and A. Goolsbee (2003). Measuring prices and price competition online: Amazon. com and barnesandnoble.com. *Quantitative marketing and Economics 1*, 203–222.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public choice*, 17–33.

Conlon, C. and J. Gortmaker (2020). Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics 51*(4), 1108–1161.

Crépeau, C., D. Gottesman, and A. Smith (2002). Secure multi-party quantum computation. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pp. 643–652.

de Cornière, A. and G. Taylor (2019). A model of biased intermediation. *The RAND Journal of Economics 50*(4), 854–882.

Eckert, A. (2003). Retail price cycles and the presence of small firms. *International journal of industrial organization 21*(2), 151–170.

Edelman, B. and M. Ostrovsky (2007). Strategic bidder behavior in sponsored search auctions. *Decision support systems 43*(1), 192–198.

Edelman, B., M. Ostrovsky, and M. Schwarz (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review 97*(1), 242–259.

Eisert, J., M. Wilkens, and M. Lewenstein (1999). Quantum games and quantum strategies. *Physical Review Letters 83*(15), 3077.

Elliott, J. T., G. V. Houngbonon, M. Ivaldi, and P. Scott (2023). Market structure, investment, and technical efficiencies in mobile telecommunications. *Investment, and Technical Efficiencies in Mobile Telecommunications (April 25, 2023)*.

European Union (2022). Regulation (eu) 2022/1925 of the european parliament and of the council of 14 september 2022 on contestable and fair markets in the digital sector (digital markets act). `https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022R1925`. Accessed: 2025-01-04.

Farronato, C., A. Fradkin, and A. MacKay (2023). Self-preferencing at amazon: Evidence from search rankings. Technical report, National Bureau of Economic Research.

Feedvisor (2021). Amazon playbook for sellers. Accessed: 2022-08-10.

Fitzsimons, J. F. and E. Kashefi (2017). Unconditionally verifiable blind quantum computation. *Physical Review A 96*(1), 012303.

Groves, T. (1973). Incentives in teams. *Econometrica: Journal of the Econometric Society*, 617–631.

Gur, Y., G. Macnamara, I. Morgenstern, and D. Saban (2023). Information disclosure and promotion policy design for platforms. *Management Science 69*(10), 5883–5903.

Gutiérrez, G. (2021). The welfare consequences of regulating amazon. *Job Market Paper, New York University*.

Hagiu, A. and B. Jullien (2011). Why do intermediaries divert search? *The RAND Journal of Economics 42*(2), 337–362.

Hagiu, A. and B. Jullien (2014). Search diversion and platform competition. *International Journal of Industrial Organization 33*, 48–60.

Hagiu, A., T.-H. Teh, and J. Wright (2022). Should platforms be allowed to sell on their own marketplaces? *The RAND Journal of Economics 53*(2), 297–327.

IBM Quantum (2025). Classical feedforward and control flow in qiskit. Online Documentation. Measurement outcomes drive classical control via `QuantumCircuit.if_test()` and other constructs.

Ikeda, K. and S. Aoki (2022). Theory of quantum games and quantum economic behavior. *Quantum Information Processing 21*(1), 27.

Johnson, J., A. Rhodes, and M. R. Wildenbeest (2020). Platform design when sellers use pricing algorithms. *Available at SSRN 3753903*.

Jürgensmeier, L. and B. Skiera (2023). Measuring self-preferencing on digital platforms. *arXiv preprint arXiv:2303.14947*.

Ke, T. T., S. Lin, and M. Y. Lu (2022). Information design of online platforms. *HKUST Business School Research Paper* (2022-070).

Kent, A. (2012). Unconditionally secure bit commitment with flying qudits. *New Journal of Physics 13*, 113015.

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

La Mura, P. (2005). Correlated equilibria of classical strategic games with quantum signals. *International Journal of Quantum Information 3*(01), 183–188.

Lam, H. T. (2021). Platform search design and market power.

Landsburg, S. (2011). Nash equilibria in quantum games. *Proceedings of the American Mathematical Society 139*(12), 4423–4434.

Leisten, M. (2024). Algorithmic competition, with humans. *Available at SSRN 4733318*.

Levin, J. and P. Milgrom (2010). Online advertising: Heterogeneity and conflation in market design. *American Economic Review 100*(2), 603–607.

Lewis, M. and M. Noel (2011). The speed of gasoline price response in markets with and without edgeworth cycles. *Review of Economics and Statistics 93*(2), 672–682.

Liebowitz, S. J. and A. Zentner (2024). The challenges of using ranks to estimate sales. *Journal of Economics & Management Strategy 33*(1), 3–24.

Maskin, E. and J. Tirole (1988). A theory of dynamic oligopoly, ii: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica: Journal of the Econometric Society*, 571–599.

Mermin, N. D. (2007). Quantum computer science: An introduction. `https://doi.org/10.1017/CBO9780511976667` or `http://www.physics.rutgers.edu/~mermin/qcomp/`. Accessed: 2025-07-13.

Meyer, D. A. (1999). Quantum strategies. *Physical Review Letters 82*(5), 1052.

Milchtaich, I. (2024). Quantum advantage in bayesian games.

Musolff, L. (2022). Algorithmic pricing facilitates tacit collusion: Evidence from e-commerce. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 32–33.

Nielsen, M. A. and I. L. Chuang (2010). *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press.

Nishimura, T. (2015). Optimal design of scoring auctions with multidimensional quality. *Review of Economic Design 19*(2), 117–143.

Noel, M. D. (2007). Edgeworth price cycles: Evidence from the toronto retail gasoline market. *The Journal of Industrial Economics 55*(1), 69–92.

Noel, M. D. (2008). Edgeworth price cycles and focal prices: Computational dynamic markov equilibria. *Journal of Economics & Management Strategy 17*(2), 345–377.

Noel, M. D. (2012). Edgeworth price cycles and intertemporal price discrimination. *Energy Economics 34*(4), 942–954.

Raval, D. (2022). Steering in one click: Platform self-preferencing in the amazon buy box. *Harvard Law School*.

Reimers, I. and J. Waldfogel (2023). A framework for detection, measurement, and welfare analysis of platform bias. Technical report, National Bureau of Economic Research.

Seller Assistant App (2024). Free amazon automate pricing tool. `https://www.sellerassistant.app/blog/free-amazon-automate-pricing-tool`. Accessed: December 11, 2024.

SellerEngine (2024). Repricing interface demo. `https://www.youtube.com/watch?v=Lzay6iybqT4`. Accessed: December 11, 2024.

Shor, P. W. (2009). Lecture 4: Quantum measurement and observables. `https://math.mit.edu/~shor/435-LN/Lecture_04.pdf`. Accessed: 2025-07-13.

United States Federal Trade Commission (2023). Federal trade commission v. amazon.com, inc. `https://www.ftc.gov/legal-library/browse/cases-proceedings/231-0032`. Accessed: 2025-01-04.

U.S. Senate (2021). American innovation and choice online act. `https://www.congress.gov/bill/117th-congress/senate-bill/2992`. Accessed: 2025-01-07.

van Enk, S. J. and R. Pike (2002). Classical rules in quantum games. *Physical Review A 66*(2), 024306.

Varian, H. R. (2007). Position auctions. *international Journal of industrial Organization 25*(6), 1163–1178.

Varian, H. R. and C. Harris (2014). The vcg auction in theory and practice. *American Economic Review 104*(5), 442–445.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance 16*(1), 8–37.

Walther, P., K. J. Resch, T. Rudolph, E. Schenck, H. Weinfurter, V. Vedral, M. Aspelmeyer, and A. Zeilinger (2005). Experimental one-way quantum computing. *Nature 434*(7030), 169–176.

Yu, C. (2024). The welfare effects of sponsored product advertising. *Available at SSRN 4817542*.

Zhang, X. and J. Feng (2005). Price cycles in online advertising auctions. *ICIS 2005 Proceedings*, 61.

Zou, T. and B. Zhou (2023). Self-preferencing and search neutrality in online retail platforms. *Available at SSRN 3987361*.

Appendix A

Appendix to Chapter 1

A.1 Imputing Daily-Level Product Page Sales

As previous literature suggests, Amazon's reports of product-category level sales ranks are excellent measures of aggregate sales on the product-page level (Musolff, 2022; Yu, 2024; Chevalier and Goolsbee, 2003; Brynjolfsson et al., 2003; Benner and Waldfogel, 2016). In this section, we provide a detailed explanation of how we use information on sales ranks to impute daily-level aggregate sales on the product-page level.

While Amazon does not disclose the formula that they use to convert product page sales in category-level sales ranks, previous work has shown that a polynomial transformation of the log of sales ranks fits the log of sales quantities very well (Liebowitz and Zentner, 2024).

Sales ranks are believed to depend on moving averages of sales on a product page over time, and information on sales ranks are updated by Amazon across product pages at different rates, although the platform claims to update sales rank information on pages between the hourly and daily level. Thus, there is some care needed in using information on ranks to impute aggregate product page sales on a daily level.

Although data from our seller includes information on sales ranks, we only observe aggregate daily sales on pages on which our seller was the only seller. Most of these instances occur on pages that have relatively low sales (log sales ranks quantities of 11 or higher) or on pages on which all other sellers temporarily exited and our seller ran out of stock on the relevant day relatively quickly. In order to be able to get accurate measures of sales ranks on quantities for pages on which there are a substantial number of daily sales, we use an estimator obtained from JungleScout, a third party service provider to Amazon sellers. As noted in other papers that use estimates of sales quantities from JungleScout, this service provider has access to substantial aggregate sales data from their customers and uses these data to estimate the relationship between sales ranks and aggregate sales. We obtained estimates for sales ranks and daily level sales from Junglescout on the product category level for products in the Arts, Crafts and Sewing, Baby, Beauty and Personal Care, Grocery and Gourmet Food, Health and Household, Office Products, and Pet Supplies categories. We then fit a regression using $\log(\text{rank}), \log(\text{rank})^5$ as outlined in Liebowitz and Zentner (2024) to impute $\log(\text{sales})$. The observed and fitted values are shown in Figure A.1.

Yu (2024) reports that these estimates perform quite well on data on observed aggregate daily shares that we obtained from a large-scale seller on Amazon. While we cannot verify the extent to which this imputation estimates daily sales for pages which take on the majority of rank values, we select pages on which to use this imputation method carefully.

First, since sales rank values are based on moving averages of sales over time, average daily ranks likely better predict total daily sales during time periods on which average sales, and therefore sales ranks, are relatively stable. We select pages and days on which the ratio of weakly average rank and the daily average rank is between 1.2 and .8, in order to eliminate imputing sales on days which may deviate substantially from long-run averages. Since it is likely more difficult to impute sales on pages that either have an extremely high volume of sales (and very low ranks) and we do not wish to estimate demand on pages with extremely low volumes of sales (and very high ranks), we limit our estimation to pages on which average daily log ranks are within 4 and 10.

A.2 Additional Theoretical Results

A.2.1 Consumer Welfare Under Independent Profits

In any locally envy-free profile studied in 1.2.1, the seller awarded rank i will set the minimum possible score such that her opponent at rank $i+1$ above does not wish to deviate and obtain rank i. This corresponds to the condition

$$f_i(p_i) = \frac{(\alpha_i - \alpha_{i+1}) f_{i+1}(p^*_{i+1}) + f_{i+1}(p_{i+1})}{\alpha_i}$$

Then, under any welfare function that is decreasing in score markups, for any fixed ranking of sellers, any scoring rule f that awards i relative to $i+1$ moreso than is consumer-optimal, is unambiguously worse for consumers.

As in 1.2.3, if the designer had information over $\mathbf{p}^*$ (or the distribution F), then she could solve for the optimal scoring rule (or ex-ante optimal scoring rule) f under any particular locally envy-free profile.

A.2.2 The Corresponding Designer's Problem

We consider the corresponding designer's problem of choosing an incentive compatible, individually rational mechanism which consists of an assignment, or ranking rule, r and a corresponding price p, where $r_j \in [0,1]$ for all $j \in J$ such that $\sum_j r_j = 1$.

Each seller j has a type $\theta_j = X_j'\beta - p_j^*$. p_j^* as in 1.2.2, is common knowledge among sellers, but is not observed by the designer. The vector X_j is common knowledge among all sellers and the designer.

Seller j's profit is given by $\pi_j^W(\mathbf{p})$ which is strictly increasing in $p_j > p_j^*$. We assume that whenever $r_j \neq 1$, j obtains her outside option $\max_{p_j} \pi_j^L(p_j)$.

The designer's objective is given by

$$\max_{r(\theta), p(\theta)} X_j'\beta - p_j(\theta)$$

for j such that $r_j(\theta) = 1$, subject to

$$p_j(\theta) \geqslant p_j^* \; \forall j \text{ s.t. } r_j = 1$$
$$\pi_j(\mathbf{p}(\theta)) \geqslant \pi_j(p_j(\theta_j', \theta_{-j}), \mathbf{p}_{-j}(\theta_j', \theta_{-j})) \; \forall j, \theta_j' \neq \theta_j$$

First, we characterize the optimal payment rule $p_j(\theta)$, under all incentive compatible mechanisms in which $r_j(\theta) = 1$ if and only if $j = \arg\max_k X_k'\beta - p_k^*$. Note that $p_j(\theta)$ depends on both $\mathbf{p}^*, \mathbf{X}$.

Let $i = \arg\max_k X_k'\beta - p_k^*$. Under any mechanism in which all types $j \neq i$, report truthfully, i can still be assigned $r_i = 1$ and receive a higher payment by reporting any

$$\hat{p}_i^* \leqslant X_i'\beta - X_j'\beta + p_j^*$$

for $j = \arg\max_{k \neq i} X_k'\beta - p_k^*$. Therefore any incentive compatible payment rule, must satisfy $p_i(\theta) = \sup\{p_i : p_i \leqslant X_i'\beta - (\arg\max_{k \neq i} X_k'\beta - p_k^*)\}$ for the highest-reported type, and $p_j(\theta) = \max_{p_j} \pi_j^L(p_j)$ for all other types.

Under such a payment rule, all lower-types j report truthfully, since when i reports truthfully, any other type j must submit

$$p_j < X_j'\beta - X_i'\beta + p_i^* < p_j^*$$

in order to be awarded $r_j = 1$, given $i = \arg\max_k X_k'\beta - p_k^*$.

Equilibrium welfare under the given assignment and payment rule is then proportional to

$$\arg\max_{j \neq i} X_j'\beta - p_j^*$$

Next, we show that it is optimal for the designer to choose such a mechanism over any other mechanism in which $r_k(\theta) = 1$ for $k \neq \arg\max_j X_j'\beta - p_j^*$. If the mechanism results in such an assignment, individual rationality requires $p_k(\theta) \geqslant p_k^*$. This implies that equilibrium welfare

$$X_k'\beta - p_k(\theta) \leqslant \arg\max_{j \neq i} X_j'\beta - p_j^*$$

Equilibrium welfare, under the optimal mechanism, is proportional to the second highest type, exactly as in the perfect information scoring setting. Any ex-post incentive compatible mechanism must ensure a price that corresponds with that described above, regardless of whether the designer has additional information over the distribution of types $\mathbf{p}^*, \mathbf{X}$. When the designer has perfect information over types, first-best welfare is given by $\arg\max_i X_i'\beta - p_i^*$, which may not be achievable in a Nash equilibrium nor a locally envy-free profile of the corresponding preferencing game. Equilibrium welfare in this perfect information setting corresponds exactly with equilibrium welfare under the same payment rule in which players do not observe opponents' types, but rather types are drawn from a distribution that is common knowledge. It is easily shown that it is weakly dominant under the above payment and assignment rules to report truthfully, by the same reasoning that it is weakly dominant to bid truthfully in a second price auction.

A.3 Supplementary Figures

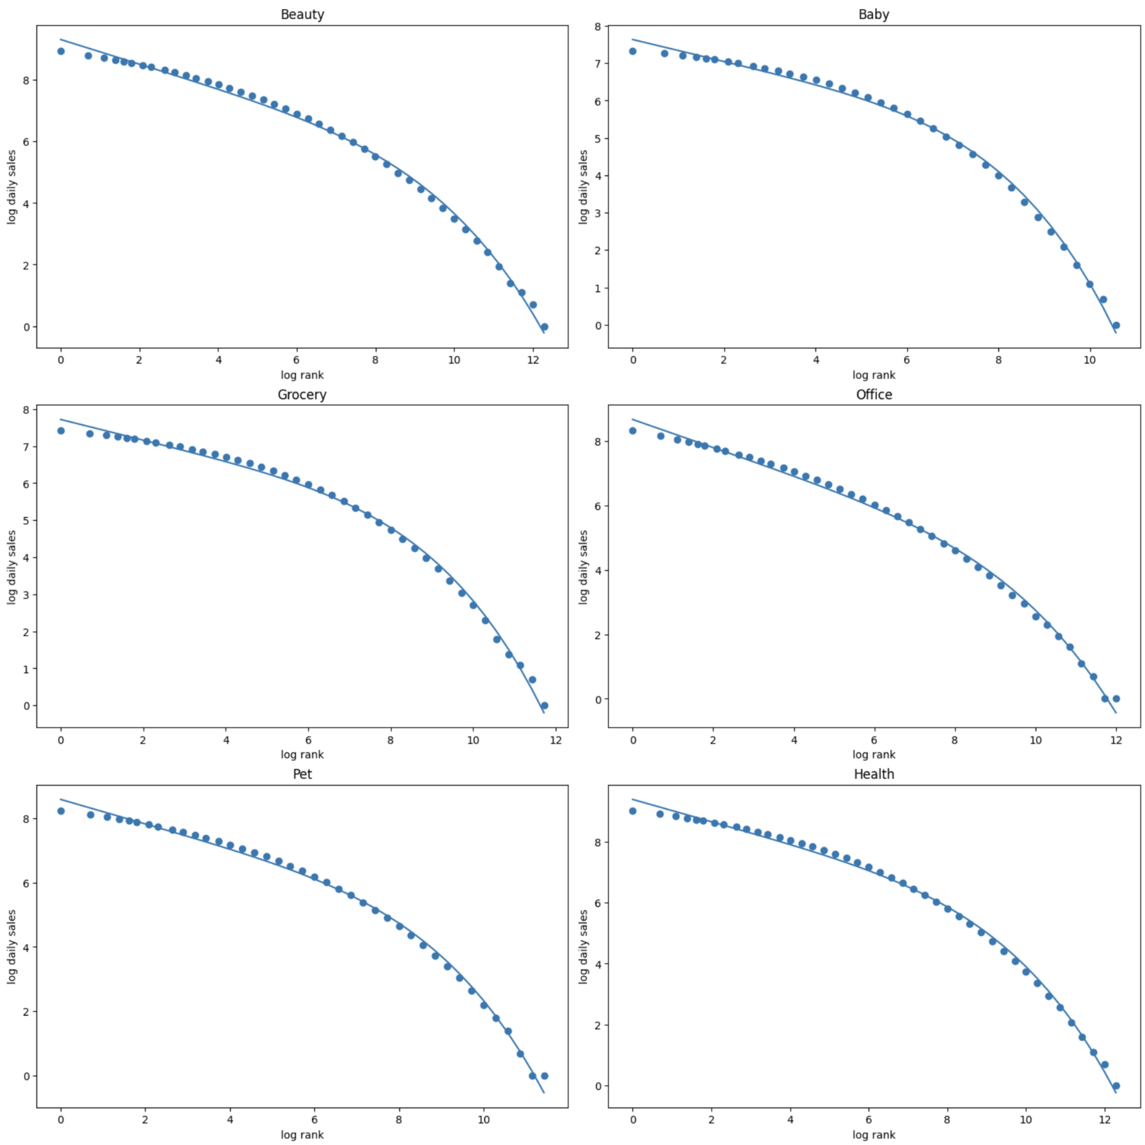

Figure A.1: *Sales rank regression fit based on Junglescout.*

Appendix B

Appendix to Chapter 3

B.0.1 Mathematical Background

Definition B.5 (Hilbert Space)**.** *A* Hilbert space $\mathcal{H}$ *is an n-dimensional complex vector space equipped with an inner product* $\langle V|W\rangle$ *for vectors* $|V\rangle, |W\rangle \in \mathcal{H}$.

Definition B.6 (Quantum State)**.** *A quantum state is a unit vector* $|\psi\rangle \in \mathcal{H}$ *such that* $\sqrt{\langle\psi|\psi\rangle} = 1$.

Definition B.7 (Unitary Operator)**.** *A* unitary operator $U : \mathcal{H} \to \mathcal{H}$ *is an* $N \times N$ *matrix that satisfies* $U^\dagger U = UU^\dagger = I$. *Any unitary operator* U *applied to* $|\psi\rangle$ *will rotate* $|\psi\rangle$ *in the unit-norm complex plane.*

Definition B.8 (Projective Measurement)**.** *A complete projective measurement corresponds to an orthonormal basis* $\{|e_1\rangle, ..., |e_N\rangle\}$. *Suppose* $\{|e_1\rangle, ..., |e_N\rangle\}$ *is an orthonormal basis for a quantum system. Then, if the system is in state* $|\psi\rangle$, *the von Neumann measurement corresponding to this basis yields* $|e_i\rangle$ *with probability* $|\langle e_i|\psi\rangle|^2$.[1]

Once the system is measured, it is projected onto one of the basis vectors $|e_i\rangle$, ensuring that if it were measured again instantaneously, the measurement would return the same result.

[1] There are infinitely different bases available, and a single basis vector in one basis will generically be a linear combination of basis vectors according to a different basis, so having a definite value for one kind of measurement will not be compatible with having a definite value for a different kind of measurement. This is the general statement of the uncertainty principle.

Definition B.9 (Computational Basis Vectors)**.** *Two orthonormal vectors spanning the two-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^2$ are given by:*

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

These vectors form the computational basis *for a single qubit.*

Definition B.10 (Qubit)**.** *A* qubit *is the quantum analogue of a classical bit. It is any unit vector in the two-dimensional complex Hilbert space $\mathcal{H}$, and can be expressed as a linear combination of the computational basis vectors:*

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle, \quad \textit{where } \alpha, \beta \in \mathbb{C} \textit{ and } |\alpha|^2 + |\beta|^2 = 1.$$

Definition B.11 (Multi-Qubit Systems and Tensor Products)**.** *Given two qubits in states*

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle, \quad |\phi\rangle = \delta\,|0\rangle + \gamma\,|1\rangle,$$

the combined two-qubit state is represented by the tensor product

$$|\Phi\rangle = |\psi\rangle \otimes |\phi\rangle = (\alpha\,|0\rangle + \beta\,|1\rangle) \otimes (\delta\,|0\rangle + \gamma\,|1\rangle).$$

Expanding, this equals

$$\alpha\delta\,|00\rangle + \alpha\gamma\,|01\rangle + \beta\delta\,|10\rangle + \beta\gamma\,|11\rangle,$$

where $|ij\rangle = |i\rangle \otimes |j\rangle$ for $i, j \in \{0, 1\}$.

More generally, systems of n qubits reside in the Hilbert space

$$\mathcal{H} = (\mathbb{C}^2)^{\otimes n}.$$

Definition B.12 (Computational Basis for Multi-Qubit Systems)**.** *The* computational basis *for an n-qubit system is the set of all 2^n classical states, the set of all possible products of individual qubit states $|0\rangle, |1\rangle$:*

For example, in a two-qubit (4-dimensional) system, the computational basis is

$$\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}.$$

Definition B.13 (Entangled Qubits)**.** *A general two-qubit pure state*

$$|\Psi\rangle = \alpha_{00}\,|00\rangle + \alpha_{01}\,|01\rangle + \alpha_{10}\,|10\rangle + \alpha_{11}\,|11\rangle$$

is said to be separable *(or a* product state*) if it can be written as a tensor product of two single-qubit states. This is the case if and only if the following condition holds:*

$$\alpha_{00}\alpha_{11} = \alpha_{01}\alpha_{10}.$$

The state is said to be entangled if it is not a product state.

For additional background see Shor (2009); Mermin (2007).