UNIVERSITY OF CALIFORNIA,
IRVINE

Essays on Return Insurance and Antitrust Issues

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Phuong Minh Vo

Dissertation Committee:
Associate Professor Michael Choi, Chair
Associate Professor Jiawei Chen
Professor Guillaume Rocheteau

2023

DEDICATION

To my parents.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisor, Michael Choi, for his exceptional guidance and support. He has been a constant source of inspiration for me in developing my research program in Economics. I am deeply grateful for the unmeasurable time he invested in advising and workshopping my papers, from the beginning research ideas to the final drafts. I am also thankful for his immense dedication to mentoring me in all aspects of my professional development in Economics. Under his mentorship, I have gained a wealth of knowledge and valuable lessons that have contributed significantly to my growth as a researcher. I am fortunate to have such an outstanding advisor.

I would also like to thank the other incredible members of my committee, Jiawei Chen and Guillaume Rocheteau, for their guidance. I am particularly grateful to Jiawei Chen for his invaluable support and feedback on my three research projects of this dissertation. His timely guidance and encouragement were a source of inspiration for me during the highs and lows of the PhD program. I am thankful for his support in many aspects that have facilitated the development of my professional career. Additionally, I would like to thank Guillaume Rocheteau for his excellent advice and feedback on all of the chapters in this dissertation. His dedication has significantly contributed to the advancement of my research program and my growth as a researcher. I am grateful for his gracious help and support from the beginning of the PhD program.

I would like to thank Khang Tran for his invaluable assistance during the initial stages of my first PhD research project, which subsequently evolved into the first chapter of this dissertation. His early and gracious support was instrumental in shaping the direction of my first research project. Additionally, I would like to thank Rajeev Tyagi for his excellent comments and feedback on the first chapter of this dissertation.

I would like to thank Joe Nguyen and Harley Vo for their tremendous help and support in writing the third chapter of this dissertation. I am grateful to Joe Nguyen for his tireless effort in workshopping the project from its inception to its final draft. Additionally, I am thankful to Harley Vo for her invaluable moral and financial support throughout the process. Their support was critical to the successful completion of Chapter 3.

I would like to thank the following professors who have imparted their knowledge and expertise during my graduate study: Stergios Skaperdas, Michael McBride, Eric Swanson, Paata Ivanisvili, Damon Clark, William Branch, Fabio Milani, Ying-Ying Lee, Igor Kopylov, Ivan Jeliazkov, Zhaoxia Yu, and Jean-Paul Carvalho.

I would like to thank my undergraduate professors, Jane Sung and Terry Olson, for inspiring and supporting me to pursue a doctoral program in Economics.

I would like to thank Melissa Valdez, Adam Cook, Katie Swindell, John Sommerhauser, and Laura Robledo-Vega for their administrative assistance.

I would like to thank the Department of Economics, the Program in Corporate Welfare, the

School of Social Sciences, and the Graduate Division at the University of California, Irvine, for their financial support, which was essential for the completion of this dissertation.

Lastly, I would like to thank my friends who supported me in various ways during my doctoral program. In particular, Bryan Castro for his help when I started writing my first research paper, Saad Andalib Syed Shah for sharing his knowledge and passion in structural Industrial Organization, Joy Yang and Vu Nguyen for their constant motivation, and many others, including but not limited to Francisco Klapp, Howard Hsu, Avik Sanyal, Abhiprerna Smit, Jaclyn Rosenquist, Mahima Vasishth, Kim Yoonjung, Aria Golestani, Hina Usman, Spencer Knight, Greta Meggiorini, Lucie Lebeau, Patrick Testa, Alexander Luttmann, Molly Maloney, Katie Bollman, Kai Mu, David Wehrly, Daniel Firoozi, Daniel Ladd, Michael Marmorstein, Shelby Kovack, David Phan, Carolina Armenta, Derek Christopher, Vishuddhi Jayawickrema, and Luis Maldonado.

I am forever grateful to the University of California, Irvine for the amazing academic and research opportunity that has shaped my career and life.

VITA

Phuong Minh Vo

EDUCATION

Doctor of Philosophy in Economics **2023**
University of California, Irvine *Irvine, California*
Distinction in Microeconomics

Master of Arts in Economics **2019**
University of California, Irvine *Irvine, California*

Bachelor of Science in Economics **2016**
Truman State University *Kirksville, Missouri*
Minor in Mathematics
Magna Cum Laude

TEACHING EXPERIENCE

Instructor, Basic Economics **2021**
University of California, Irvine *Irvine, California*

Teaching Assistant **2018–2022**
University of California, Irvine *Irvine, California*

Teaching Assistant **2015–2016**
Truman State University *Kirksville, Missouri*

AWARDS

Outstanding Teaching Assistant Award, *UCI Department of Economics* 2022

The Program in Corporate Welfare Fellowship, *UCI School of Social Sciences* 2022

Distinction in Microeconomics, *UCI Department of Economics* 2018

Graduate Dean's Recruitment Fellowship, *UCI Department of Economics* 2017

Summer Research Fellowship, *UCI Department of Economics* 2019–2021

Merit Fellowship in Economics, *UCI School of Social Sciences* 2017–2022

ABSTRACT OF THE DISSERTATION

Essays on Return Insurance and Antitrust Issues

By

Phuong Minh Vo

Doctor of Philosophy in Economics

University of California, Irvine, 2023

Associate Professor Michael Choi, Chair

Chapter 1 introduces a continuous-time monopoly model that considers a return policy allowing consumers to return purchased products within a specified period. The model shows that an easy return policy, allowing for no-questions-asked returns, reduces consumer surplus compared to a stricter policy that only refunds under specific conditions. This decrease in consumer surplus happens when consumers are not highly price-sensitive. The study also explores how product quality affects the return policy and finds that lower quality products result in longer return periods. Additionally, lower quality products can lead to higher market prices if the value of the defective item is high enough.

In Chapter 2, a duopoly model is developed to study competition between an online seller and a local store seller in the presence of a return policy. The model finds that offering returns after sales increases the sellers' market power. The local store targets consumers with lower utility by offering a shorter return period and a lower price. The study also shows that an increase in the information cost at the local store increases market price dispersion and gives the online store more market share.

Chapter 3 uses economic theory and experiments with AI-based pricing algorithms to analyze the impact of consumer search friction on collusion, market prices, and consumer welfare. The chapter develops an oligopoly model where consumers search sequentially for the best product with advertised prices. The study finds that collusion is easier to sustain with lower search costs. However, increasing search costs can reduce the collusive price, but this does not increase consumer surplus if the collusion sustains. The experiments show that simple reinforcement learning algorithms (Q-learning) can adopt a trigger-price strategy to keep prices above the competitive level in a frictional market.

Chapter 1

Product Return Policies, Pricing, and Consumer Welfare

1.1 Introduction

Many retailers accept product returns within a certain time and give refunds. According to "Customer Returns in the Retail Industry 2021" by the National Retail Federation, the annual product returns in the U.S. retail industry totaled an unprecedented amount of 761 billion U.S. dollars.[1] Online returns accounted for more than 218 billion U.S. dollars, which is more than double as much in 2020 and more than five times as much in 2019. The two key factors of a return policy are: (1) how much the consumers receive when they return a product, taking into account deductions, fees, and return hassles, and (2) the length of the return period.[2] The literature has analyzed the role of return policies focusing on the former factor. The latter factor has been treated as exogenous in a common "two-period"

[1] In comparison, the U.S. total spending on national defense in 2020 was 714 billion dollars (https://www.cbo.gov/publication/57170).

[2] Posselt et al. (2008) use these two factors to evaluate return policies.

model, in which the consumers make the purchasing decision in the first period, learn the true valuation, and then decide whether to return the product in the second period (more details in the next section). However, the return periods vary among sellers, ranging from a few days to several months.[3] Longer return period allows the consumers to learn more about the product but potentially results in more product returns. This article studies the return policies by endogenizing the return period.

The two main goals of this article are to explore the relationship between product quality and contract terms and to study the economic implications of different return policies. In our study, the return process serves as a tool to reallocate faulty products where the seller can salvage the defective items while the consumers cannot. We find that lower product quality generally results in a more extended return period. Moreover, suppose the return process is cost-effective enough, meaning that the seller can get a relatively high salvage value from the defective product compared to the cost of the return process. In that case, lower product quality can even increase the market price. Regarding social welfare, we show that an easy return policy that allows consumers to return any product (e.g., no-questions-asked return policy) can lead to a lower consumer surplus compared to a stricter return policy where conditions apply to receive refunds if the consumers have low price sensitivity.

To conduct our study, we present a monopolist model in continuous time in which consumers constantly learn about the quality of their purchased products. The key feature of this model is that each consumer gets an idiosyncratic flow of utility from the product, which can be enjoyed over a span of time. Thus the consumers can enjoy some utility during the return period. Consumers observe their utility flows before purchasing. However, they do not know the exact quality type of the product. Each product can either be a good type, which never breaks down, or a bad type, which can break at a Poisson rate. The utility flow drops when the product breaks. Consumers constantly receive signals about the quality type by

[3] See Posselt et al. (2008) for some evidence.

detecting malfunctions. Unlike the consumers, the seller can extract some salvage value from the defective product because of their recycling or refurbishing ability. However, it is costly for the seller to accept returns because of the refunds and the returns management. We reformulate the dynamic model as a static monopoly problem and provide a solution that fully captures consumer shopping and returning outcomes. This reformulation allows us to study how product quality affects the seller's pricing and return policy and conduct our welfare analysis on different return policies.

We find that if the return process is cost-effective enough in the sense that the salvage value of the returned item is relatively high compared to the cost of the return process, then the market price decreases in product quality, which is opposite the classical result in the literature. If product quality falls, consumers are more price sensitive. The seller resorts to lowering the markup; otherwise, they would lose too many consumers. This mechanism leads to the classical result from the literature. In our model, the relationship between product quality and the market price is more subtle. Given that the return process is cost-effective, the seller accepts returns to reallocate the faulty items and extract the salvage value. Lower product quality makes the consumers more price sensitive and raises the chance of a product being returned to the seller. For each product, it is more costly for the seller to commit to their return policy. This return effect is analogous to the impact of an increase in the marginal production cost, both of which tend to increase the market price.

For the welfare analysis, we find that easy return policies (e.g., no-questions-asked return policy), while most of the time can improve consumer welfare, can also harm consumers under some conditions. More specifically, accepting all returns can decrease the consumer surplus if the consumers have low price sensitivity. Much of this finding is due to the result of our consumer behavior analysis when the no-questions-asked policy is in effect. We find that consumers are more sensitive to the expected expenditure when the return period is longer. The reason is that with an extended return period, using the product strictly in the return period and then sending back the item to get a full refund (i.e., wardrobing or renting) becomes an attractive alternative to consumers. Extending the return period makes the seller lose some potential consumers who were willing to keep the non-broken products if the return period was unchanged. With low price sensitivity, consumers are willing to pay a high markup for the product. Losing consumers is relatively costly to the seller. As a result, the seller shortens the return period and gives up the ability to salvage broken products. By forgoing extra surplus from reallocating faulty products, the seller increases the expected consumer expenditure, reducing the consumer surplus.

When the seller can refund discriminate (i.e., the refund is based on the condition of the returned item), we show that better product quality generally leads to a shorter period. To be more specific, we find that an increase in the probability of getting a good-type product reduces the return period. It is because each product is less likely to be returned. The return process then becomes less cost-effective. Therefore, the seller invests less in the return process by shortening the return period. In contrast, we also find that if the bad type products experience issues at a higher rate, the return period does not necessarily increase. The reason is that if the breaking down rate is high, the bad type products will encounter issues more quickly, and more returns will be made shortly after the purchase, which allows the seller to shorten the return period.

The rest of the article is organized as follows. Section 1.2 reviews the literature and explains how our work contributes to the literature. Section 1.3 introduces the baseline model. Section 1.4 characterizes the market equilibrium and analyzes the optimal return policy. Section 1.5 provides a set of comparative statics results regarding product quality. Section 1.6 introduces the model with an easy return policy and analyzes consumer behavior and welfare. Finally, Section 1.7 offers some concluding remarks. All omitted proofs are in the Appendix.

1.2 Literature review

The literature has focused on experience goods to study the economic rationale for return policies. For an experience good, the utility is difficult to observe before purchase, but this information can be obtained upon consumption. Che (1996) puts forward a model in which consumers are risk averse and do not observe the utilities before purchase. The return policy can insure consumers against *ex-post* loss. He shows sellers can benefit from the return policy if consumers are highly risk averse.

Courty and Li (2000) study a model in which a monopolist of airplane tickets has a menu of refund contracts to screen the consumers.[4] In their model, the consumer privately knows the probability distribution of her valuation before purchasing and subsequently learns her actual valuation before making the return for refund decision. They show that lower return costs will undermine the ability of the firm to screen the consumers.[5] In contrast to Courty and Li (2000), this article restricts the analysis to a single contract and allows the seller to salvage the returned items. We show that lower return cost helps facilitate the reallocation of defective products and therefore increases the profit of the seller and the consumers surplus.[6]

The literature has also studied the effects of return policies under competitive environments. Inderst and Tirosh (2015) study a duopoly model in a vertically differentiated industry. They show that return policies work as "metering devices" to extract more consumer surplus from infra-marginal customers. The market is segmented, where consumers with lower ex-ante expected utility go to the low-quality firm. In contrast, consumers with high ex-ante expected utility go to the high-quality firm. For the high-quality firm, its "marginal" consumers have the highest marginal valuation for a higher refund. Therefore, the high-quality firm can

[4] Escobari and Jindapon (2014) define the menu with two specific contracts: one refundable and one non-refundable. They also show that sellers can use return policies as a tool for price discrimination.

[5] A similar result is obtained by Krähmer and Strausz (2015) and Arya and Mittendorf (2004).

[6] Davis et al. (1995) also allow the seller to extract salvage value from the returned products. They show that accepting returns is more profitable than selling as-is if the seller's salvage value is higher than the buyer's return cost.

increase both the price and the refund to capture more consumer surplus from its infra-marginal consumers without affecting the marginal consumers. The opposite holds for the low-quality firm; a lower price and refund can extract more consumer surplus.

When consumers shop sequentially and are imperfectly informed about the horizontally differentiated products, the return costs resemble the search costs in consumer search theory, e.g., Weitzman (1979); Choi et al. (2018). Consumers have to pay these costs every time they explore another option. Petrikaitė (2018) studies a duopoly model in which consumers shop sequentially. She shows that higher return costs intensify the price competition, which consequently leads to a lower symmetric equilibrium price.[7] Choi et al. (2018) show that the symmetric equilibrium price decreases in search costs in a market with any number of sellers. In a monopoly market, the result that accepting returns leads to higher prices is intuitive and is obtained by several studies, e.g., Palfrey and Romer (1983), Lutz and Padmanabhan (1998) and Inderst and Ottaviani (2013) for single-product monopolists, or Anderson et al. (2009) for multi-product monopolists.

Lastly, the literature has explored different rationales for return policies in a bilateral trade environment. Inderst and Ottaviani (2013) show that a generous return policy allows the seller to "cheap talk" at the point of sale.[8] The key feature of their model is that before both parties meet at the point of sale, the seller observes some informative signal about the utility the consumer will enjoy. The communication of the seller's observed signal reduces the loss that the seller endures when the buyer returns the product. Hinnosaar and Kawai (2020) instead study a bilateral trade model in which only the buyer knows some private signal about the utility. The key feature of their model is the seller's uncertainty about the buyer's distribution of signals. They show that the seller can use generous return policies as

[7] As opposed to Petrikaitė (2018), Shulman et al. (2011) allow the consumers to learn about their valuation for the competing products after experimenting with the first product. They show that the symmetric equilibrium prices increase when the consumers are less certain about their preferences.

[8] Grossman (1981) and Moorthy and Srinivasan (1995) show that sellers can use the return policies as a costly signal for product quality.

a tool to hedge against uncertainty about the buyer's signal structure.

1.3 Baseline model

Time is continuous with an infinite horizon. The market starts at time $t = 0$. The market consists of one seller and one unit mass of consumers. Each consumer has one unit of demand. In an arbitrarily short period of time from t to $t + \mathrm{d}t$, the consumer gets a utility flow of $v\mathrm{e}^{-\gamma t}\mathrm{d}t$ from using the product.

For each consumer, v is independently and identically drawn from the interval $[\underline{v}, \overline{v}]$ according to a log-concave distribution (distribution function is F, density function is f, and location parameter is μ).[9] We allow $\underline{v}$ and $\overline{v}$ to be infinite. The consumers observe their idiosyncratic v before purchase.

For bounded utility, we assume $\gamma > 0$.[10] We note that γ is not a future discount factor. For simplicity, we assume the seller and consumers do not discount future.[11] The utility from not buying the product is normalized to zero.

There are two types of products: good and bad. Good products never break down, whereas each bad product breaks at a Poisson rate λ. When the product breaks, consumers cannot use it anymore. From that period forward, the utility flow drops to zero. The product's type is not observable unless the product is already broken. At the time of purchase, the probability that a product is good is α.

The seller can determine the length of the return period, $l \geq 0$, in which consumers can

[9] The hazard rate function, $f(v)/[1 - F(v)]$, is then monotone non-decreasing.

[10] One interpretation of γ can be the depreciation rate of the utility flow. It reflects the fact that products constantly suffer wear and tear with use. Older products give less utility than the newer ones.

[11] The time horizon of all payments from buyers and sellers (purchase and refund) is constrained within the return period, which is relatively short in the sense that it is usually bounded by several months. Common payment types (e.g., cash, credit cards, debit cards, digital wallets) typically take less than several days to complete the transactions.

return a product for a refund. The return period begins right after purchasing. To keep the contract simple, we restrict the seller issues full refunds. The seller can observe the condition of the returned product (broken or not) and has the option to decline the non-broken returns.[12] It will later be shown in Section 1.6 that declining non-broken returns is more profitable for the seller. We will momentarily take this result for granted and let the seller only accept broken returns.

The seller incurs a marginal cost of c for each new product. The seller gets a salvage value $se^{-\sigma t}$ for a returned item that has been used up to time t (i.e., salvage value depreciates at rate σ during use). The salvage value is always less than the marginal cost, i.e., $s < c$.

Returns management costs the seller an amount of ml to serve each consumer throughout the return period.[13] The seller commits to their contract and maximizes their profit. The consumers are risk neutral.

Timing. When the market opens at $t = 0$, the seller announces their contract including price and return period (p, l). Consumers observe their v, the contract, and decide whether to buy the product. The seller incurs the production cost and the returns management cost. Consumers can return their broken products in the return period $t \in (0, l]$. The seller issues refunds when consumers return the items.

[12] Because v is observable before purchase, returning a product without any malfunction means that the consumer sends back the item while it is performing at his highest expectation. This behavior is referred to as *moral hazard in consumption* in Davis et al. (1998). In practice, the common names for this behavior are wardrobing or renting. In our model, the seller can prevent this behavior by declining non-broken returns.

[13] Return management is a complex process, especially for e-commerce businesses. This process includes providing a return portal, customer service, logistics, and warehouses. Many direct-to-consumer companies rely on third-party returns solutions and third-party logistics to manage returns. In our model, all fixed costs are captured in the marginal production cost c. The remaining variable costs are captured in the returns management cost ml.

1.4 Equilibrium

When the market opens, let $U_0(v)$ be the expected utility a consumer can get from the product. Thus,

$$U_0(v) = \frac{\alpha v}{\gamma} + \frac{(1-\alpha)v}{\gamma+\lambda}. \tag{1.1}$$

In the first term of $U_0(v)$, α is the probability that the product type is good, and v/γ is the expected utility from a good type product. Similarly, in the second term of $U_0(v)$, $1-\alpha$ is the probability that the product type is bad because each bad product breaks at Poisson rate λ, its expected utility is $v/(\gamma+\lambda)$.

Product Quality: As opposed to the product type (good and bad), we let Q_0 be the ex-ante product quality (or product quality for short). More specifically, we define Q_0 as U_0/v, which means,

$$Q_0 \equiv \frac{\alpha}{\gamma} + \frac{1-\alpha}{\gamma+\lambda}. \tag{1.2}$$

We let Ω_l be the probability that a random product survives the entire return period. Thus

$$\Omega_l \equiv 1 - (1-\alpha)\int_0^l \lambda \mathrm{e}^{-\lambda t}\mathrm{d}t. \tag{1.3}$$

A product breaks only if it is bad, which has a probability of $1-\alpha$. The integration is the probability that a bad product breaks during the return period. Because consumers can get a full refund if the product breaks during l, the expected expenditure (taking into account refunds) for the consumers is given by

$$\hat{p} \equiv \Omega_l p. \tag{1.4}$$

Because of risk neutrality, the consumer purchases the product if and only if her expected utility exceeds the expected expenditure, i.e., $v > \hat{p}/Q_0$.

The seller has to issue a full refund for each return and get a depreciated salvage value. For each sold product at $t = 0$, the seller expected revenue (taking into returns management cost but not the marginal production cost c) is given by

$$\hat{r} \equiv p - (1-\alpha)\int_0^l (p - s\mathrm{e}^{-\sigma t})\lambda \mathrm{e}^{-\lambda t}\mathrm{d}t - ml.$$

The return process creates a difference between the seller's expected revenue and the consumer's expected expenditure for each product. We refer to this difference $(\hat{r} - \hat{p})$ as the *seller's return advantage* and denote it by A_l. Thus

$$A_l = s\lambda(1-\alpha)\int_0^l \mathrm{e}^{-(\sigma+\lambda)t}\mathrm{d}t - ml. \tag{1.5}$$

The salvage value and how often a product is returned due to manufacturing issues, reflected in $\lambda(1-\alpha)$, positively affect this advantage.[14] The seller can control the return advantage through the return period l. This advantage is not affected by the price of the product.

Table 1.1: Table of notations

$U_0(v)$	$\triangleq$	Expected utility formed at time 0
Q_0	$\triangleq$	Product quality (or ex-ante product quality)
Ω_l	$\triangleq$	Probability that a product survives the return period l
$\hat{p}$	$\triangleq$	Consumer expected expenditure (taking into account refunds)
$\hat{r}$	$\triangleq$	Seller expected revenue on each product
A_l	$\triangleq$	Seller return advantage $(\hat{r} - \hat{p})$

Because $\hat{p}$ is strictly increasing in p, the seller's choice variables can technically be changed

[14]Cf. Davis et al. (1995) find that the profitability of a full-refund policy is positively affected by the salvage values of return merchandise and the probability of mismatching the product to consumers.

to $\hat{p}$ and l and the profit can be described as follows,

$$\Pi(\hat{p}, l) = \left[1 - F\left(\frac{\hat{p}}{Q_0}\right)\right](\hat{p} + A_l - c)\,. \tag{1.6}$$

If the seller does not accept returns (i.e., $l = 0$, $\hat{p} = p$, $A_l = 0$), the profit function has the standard monopoly form: $[1 - F(p/Q_0)]\,(p - c)$. However, with a return policy, the seller can control the return period l and enhance the profit by increasing the return advantage A_l. In fact, the optimal return period l^* will maximize the seller's return advantage.

Proposition 1 (Return Period and Price Comparison)**.**

- *If $s\lambda(1-\alpha) > m$, the equilibrium return period is given by*

$$l^* = \frac{\ln\left[\frac{s}{m}\lambda(1-\alpha)\right]}{\sigma + \lambda}$$

and the equilibrium price p^* is higher than in the standard no-return model p'.

- *If $s\lambda(1-\alpha) \leq m$, the equilibrium return period is 0 and $p^* = p'$.*

Proposition 1 implies that the equilibrium return period increases in the ratio s/m, which can be interpreted as the cost-effectiveness of the return process, and decreases in σ, which can be interpreted as the depreciation rate of the salvage value. This result is intuitive. The more meaningful results lie in product quality parameters α and λ. Proposition 1 suggests that a lower product quality in the sense of a lower α leads to a longer return period. If the products are more likely to break, there will be more extra surplus from salvaging the broken products. The seller will extend the return period to collect more of this extra surplus.

The effect of λ on the equilibrium return period is more subtle and is presented in the Corollary below.

Corollary 1. *There exists a unique $\dot{\lambda} > m/[s(1-\alpha)]$ such that the equilibrium return period l^* increases in λ if $\lambda \in [m/[s(1-\alpha)], \dot{\lambda}]$, and decreases in λ if $\lambda > \dot{\lambda}$.*

Similar to α, a higher product quality in the sense of a lower λ can lead to a shorter return period. However, if λ is high enough, the return period can decrease in λ. It is because each bad type tends to break sooner if the breaking rate is high. Therefore, the seller can shorten the return period to save the cost while still does not lose much on extracting salvage value. Figure 1.1 illustrates simulations of the equilibrium return period l^* with different α and λ. The seller accepts the return and has a strictly positive return period if $s\lambda(1-\alpha) > m$. The graph of l^* decreases in α and has a humped shape in λ.

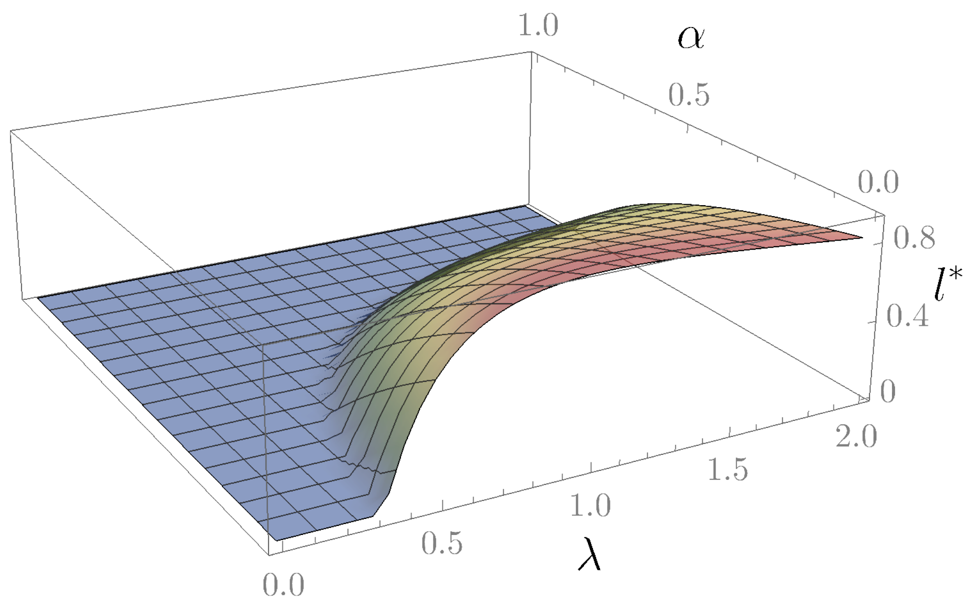

Figure 1.1: Simulations of return period (l^*) with different product quality (α, λ)

Note. Parameters used: $c = 1$, $s = 0.3$, $m = 0.1$, $\gamma = 0.2$, $\sigma = 0.2$, $v \sim \mathcal{N}(2, 1)$.

Proposition 1 also shows that if the return process is cost-effective enough (i.e., $s\lambda(1-\alpha) > m$), the seller accepts returns within a certain period l^*. The price will be higher in this case than in the standard model, where the return channel is unavailable. This result is intuitive because the seller does not have to issue refunds in the standard no-returns model. However, consumers are still better off when the seller accepts returns. The extra surplus from salvaging broken products decreases the consumer expected expenditure.

In our model, the return costs are entirely borne by the seller. If these costs increase, the seller will lose their return advantage, resulting in a lower surplus to both seller and consumers. This result, however, is independent of which party (or both) has to bear the return costs. Suppose the consumer incurs a higher cost to return a product. In that case, the difference between how much the seller expects to get and how much the consumer expects to pay for each product (seller's return advantage) also decreases.

1.5 Comparative statics

This section presents a set of comparative statics results regarding product quality and market price. We analyze the effect of product quality, particularly in respect of α, the probability that the product is good, and λ, the breaking rate of a bad product.

By equation (1.6), the equilibrium expected expenditure must satisfy

$$\hat{p}^* = \frac{1 - F(\frac{\hat{p}^*}{Q_0})}{f(\frac{\hat{p}^*}{Q_0})} Q_0 + c - A_{l^*}. \tag{1.7}$$

A decrease in α or an increase in λ not only raises the seller's return advantage for each sold product (A_l increases) but also has a negative impact on the consumers' expected utilities (Q_0 decreases). These effects together will decrease the equilibrium expected expenditure (technical proof is in the Appendix). Figure 1.2a represents simulations of equilibrium expected expenditure under different product quality.

However, the market price does not necessarily go down with lower product quality (see figure

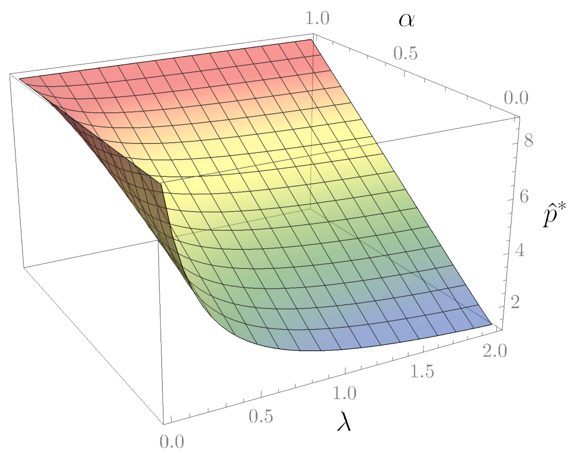

(a) Expected expenditure

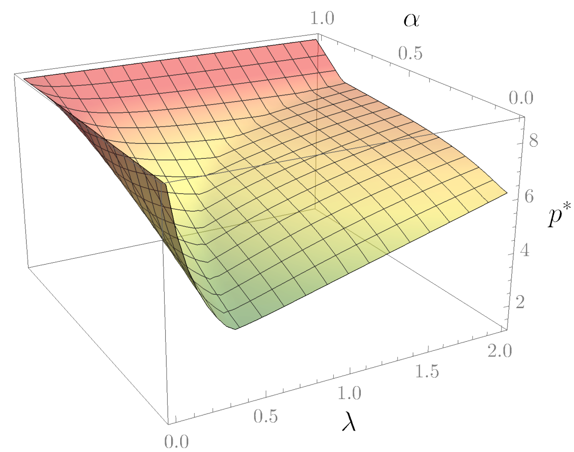

(b) Market price

Figure 1.2: Simulations of expected expenditure ($\hat{p}^*$) and market price (p^*) with different product quality (α, λ)

Note. Parameters used: $c = 1$, $s = 0.3$, $m = 0.1$, $\gamma = 0.2$, $\sigma = 0.2$, $v \sim \mathcal{N}(2,\, 1)$.

1.2b). When the return period is 0, i.e., $s\lambda(1-\alpha) \leq m$, the expected expenditure is the same as the market price. They both increase in α and decrease in λ. When the return period is strictly positive, the market price is bent upward relative to the expected expenditure. We find that if the salvage value is relatively higher than the returns management cost, a lower α can actually increase the market price (see figure 1.3).

Proposition 2 (Effect of Product Quality Regarding α on Market Price)**.** *Lower product quality in the sense of a lower α can lead to a higher price if the salvage value is relatively higher than the returns management cost. More specifically, for any $\tilde{\alpha} \in (0, 1)$ then the equilibrium price decreases in α for $\alpha \in (0, \tilde{\alpha})$ if s/m is higher than some finite threshold.*

Proposition 2 says that if the salvage value of the broken product is relatively higher than the return management cost, lower product quality in the sense that a product is more likely to break down leads to a higher market price. To understand this outcome, it is helpful to see when the seller increases their price. By raising the price, the benefit of extracting more surplus from the infra-marginal consumers comes with the cost of losing the marginal

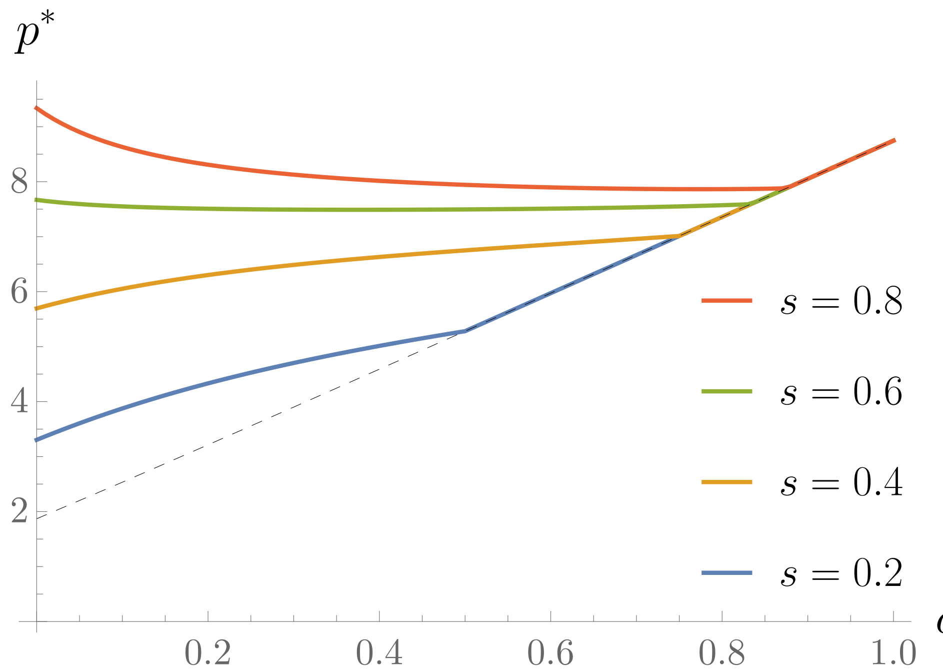

Figure 1.3: Simulations of the market price (p^*) with different salvage values (s)

Note. Parameters: $c = 1$, $\lambda = 1$, $m = 0.1$, $r = 0.2$, $v \sim \mathcal{N}(2, 1)$. The thin dashed line is the market price in the standard monopoly model where returns are not accepted.

consumers. This cost is determined by the markup on each product and the price sensitivity of the marginal consumers.

Without the return channel and given a fixed contract, lower product quality does not affect the markup but makes the consumers more price sensitive. Raising or keeping the same price is not optimal for the seller because the cost of losing marginal consumers exceeds the benefit of extracting surplus from the infra-marginal consumers. As a result, the seller lowers their price. This mechanism leads to the classical product quality result.

Now, suppose the return channel is available. In that case, lower product quality makes the consumers more price sensitive and impacts the expected profit from each product. More specifically, a higher probability that a product is a good type leads to a higher probability that a product survives the return period (i.e., Ω_l increases in α). As a result, if α decreases, the expected profit from each product will decrease because it will be more likely to be returned. We refer to this effect as *the return effect*.

If the return process is cost-effective enough in the sense that s/m is high enough, the return period will be long, which intensifies the return effect. If the return effect is dominant compared to the price sensitivity effect, raising the price can benefit the seller because losing marginal consumers becomes less costly. In other words, a decrease in α makes each product more costly for the seller to commit to its return policy. This return effect is analogous to an increase in the marginal cost, which leads to a higher market price.

Now we look at the other dimension of the product quality: the breaking rate λ of each bad product. If the product quality drops due to an increase in λ, the market price can also increase. As shown in Figure 1.2b, the market price increases in λ when λ is slightly higher than the threshold $m/[s(1-\alpha)]$, that is when the equilibrium return period is short but not zero.

Proposition 3 (Effect of Product Quality Regarding λ on Market Price)**.** *Suppose $\gamma \geq \sigma$. There exists a $\tilde{\lambda} > m/[s(1-\alpha)]$ such that the equilibrium price increases in λ for $\lambda \in [m/[s(1-\alpha)], \tilde{\lambda}]$.*

An increase in λ also has a return effect and a quality effect by making each bad-type product tend to break sooner. Unlike a decrease in α, an increase in λ only has a minor return effect when the return period is long. Suppose the return period is long; then, most of the breakdowns occur during the return period, so an increase in λ will only increase returns by a minimal amount. Therefore, a shorter return period intensifies the return effect. When the return period is short enough (i.e., λ is close to $m/[s(1-\alpha)]$ from above), the market price increases in λ.

1.6 Easy return model

In this section, we analyze the model where the consumers can return any product during the return period. This policy is commonly known as a no-questions-asked return policy. The seller commits to accepting all returns, no matter the product's condition, during the return period. Recall that in the baseline model, we assume that the seller can observe the condition of the returned items so the seller can prevent wardrobing by declining non-broken returns. If we relax this assumption, the seller can only implement a no-questions-asked return policy. In this section, we ask whether this easy return policy improves the consumer surplus and the conditions under which this policy might harm the consumers. We will first analyze the effect of this policy on consumer behavior. Then we will derive the market equilibrium and study the effects of the policy on social welfare.

1.6.1 Consumer belief

In this subsection, we first analyze how a consumer's subjective belief about their product's quality type evolves. Then we derive the consumer's optimal strategy and provide the shopping and returning outcomes. We refer to the consumer's subjective belief about the probability that her product is a good type as the *consumer's belief*. All consumers share the same prior belief α before purchase. If the product breaks down, this belief immediately drops to zero. If the product works normally, the consumer updates her belief in a Bayesian fashion.

Let the belief of a typical consumer with a functioning product at time t be $\pi_t \in [0, 1]$. The consumer updates her belief after every short period of time $\mathrm{d}t$. If the product works normally at time $t + \mathrm{d}t$, the consumer will update her belief accordingly to $\pi_{t+\mathrm{d}t}$, and so on.[15] Because $\mathrm{d}t$ is arbitrarily small, the probability that a bad-type product breaks during

[15]This is similar to the learning process in Mason and Välimäki (2011). They develop a model in which

dt is λdt. Let $\mathrm{d}\pi_t \equiv \pi_{t+\mathrm{d}t} - \pi_t$. By Bayes' rule,

$$\mathrm{d}\pi_t = \frac{\pi_t}{\pi_t + (1-\pi_t)(1-\lambda \mathrm{d}t)} - \pi_t.$$

The rate of change of the belief with respect to time as $\mathrm{d}t \to 0$, converges to the following Riccati equation:

$$\frac{\mathrm{d}\pi_t}{\mathrm{d}t} = \pi_t(1-\pi_t)\lambda. \tag{1.8}$$

When the consumer is unsure about her product type (i.e., $0 < \pi_t < 1$), her belief will increase in the following short period if the product works normally. The consumer's belief will most rapidly increase when she is most uncertain about the product type (i.e., $\pi_t = 0.5$). A signal of having no malfunction in a short amount of time will make the consumer's belief lean quickly toward the good type. Because all consumers have the same prior belief before they purchase, π_t can be described explicitly as a function of time.

Lemma 1. *The belief of a typical consumer with a functioning product is given by*

$$\pi_t = \frac{\alpha}{(1-\alpha)\mathrm{e}^{-\lambda t} + \alpha}.$$

All consumers with a non-broken product share the same belief. It is because the information they get is the same. The belief converges to 1 (i.e., the product is a good type) if the product keeps working normally over time.

---

the seller learns about the arrival rate of buyers.

1.6.2 Shopping and returning outcomes

Each consumer faces a sequential problem in which she has to make buying and returning decisions. The consumers first decide whether to buy the product. If they buy it, they can return it at any time during the return period. It is clear that the consumers will buy the product if and only if they have $v > 0$ because they can return the product for a full refund. If the product has broken, the consumer's optimal strategy is to return the item before the return period expires. If the product is still working fine, it is not optimal for them to return it immediately if they still have some time before the return period expires. By postponing the return, they can enjoy more utility and still get the full refund.

At the last moment of the return period ($t = l$), the consumers have to decide whether to return for a refund or to keep and accept the risk of breaking in the future. The expected utility from keeping the product depends on the depreciated utility $v\mathrm{e}^{-rl}$ and the belief at that moment π_l. More specifically, at time l, the expected utility from keeping the product is $Q_l v$ where

$$Q_l \equiv \mathrm{e}^{-\gamma l}\left(\frac{1-\pi_l}{\gamma+\lambda} + \frac{\pi_l}{\gamma}\right). \tag{1.9}$$

Therefore, the consumer keeps her product if and only if v exceeds the keeping threshold v_k where $Q_l v_k = p$. The consumer's optimal strategy and the shopping/returning outcomes are described in the following Proposition.

Lemma 2. *Suppose consumers can return any product. The consumer's optimal strategy is as follows:*

(i) at time $t = 0$, buys if and only if $v > 0$;

(ii) at time $t \in (0, l)$, returns if the product breaks and keeps otherwise;

(iii) at time $t = l$, returns if the product breaks or if $v < v_k$ and keeps otherwise.

Although the consumer problem is cast in continuous time, the solution can be described in a static fashion by the two thresholds 0 and v_k. For further analysis, we refer to consumers with $v \geq v_k$ as *potential consumers*, and consumers with $v \in (0, v_k)$ as *non-potential consumers.* Potential consumers will keep their product if it survives the return period. Non-potential consumers will return the item eventually.

Because $\Omega_l p = \hat{p}$, at the keeping threshold we must have

$$\Omega_l Q_l v_k = \hat{p}. \tag{1.10}$$

Note that Ω_l is the probability that the product survives the return period. Therefore, $\Omega_l Q_l v$ represents the expected utility of the product from time l and forward. Loosely speaking, $\Omega_l Q_l v$ is the expected utility the consumer enjoys if she forgoes all the utility during the return period. Consumers with v_k are indifferent between paying the price and keeping the product that survives the return period. Therefore, these consumers will also be indifferent between not buying the product and this following contract: (a) deposit the money when the market opens, (b) forgo all the utility during the return period, (c) receive a refund if the product breaks during the return period, and (d) keep the product if it survives the return period. Under this hypothetical contract, the consumer expected utility is $\Omega_l Q_l v_k$, and the expected expenditure is the same as in the actual contract, $\hat{p}$, which explains equation (1.10). Clearly, a longer return period will increase the expected forgone utility and affect the potential consumer sensitivity to the expected expenditure. The following Proposition provides the effect of the return period on consumer behavior.

Proposition 4 (Consumer Behavior with Easy Return)**.** *Suppose the seller accepts all returns. A longer return period makes the potential consumers more sensitive to the expected expenditure (taking into account refunds) in the sense that the critical expected expenditure decreases in the length of the return period.*

- *Potential consumer: a consumer who keeps the product if it does not break during the return period.*
- *Critical expected expenditure: the maximum expected expenditure for a consumer to remain being a potential consumer.*

Even though Q_l can increase in l because consumer belief that their product is a good type increases with l. However, $\Omega_l Q_l$ is proved to be decreasing in l. The intuition is that in the actual contract, consumers also have an option of using the product strictly in the return period while paying nothing. Consumers with v_k are indifferent between this free ride option and the keeping option. As the return period extends, the free ride option becomes more attractive, thus making the potential consumers more sensitive to their expected expenditure with the keeping option.

1.6.3 Welfare

Similar to the baseline model, the expected profit from selling a product to a potential consumer is $\hat{p} + A_l - c$. Let S_l^n be the expected salvage value the seller can collect from selling a product to a non-potential consumer.[16] It is clear that if the salvage value does not depreciate (i.e., $\sigma = 0$), $S_l^n = s$. However, when the salvage value depreciates over time, a longer return period will decrease S_l^n because the consumers postpone their return as late as possible. The seller's profit can be described as follows,

$$\Pi(\hat{p}, l) = \left[1 - F\left(\frac{\hat{p}}{\Omega_l Q_l}\right)\right](\hat{p} + A_l - c) + \left[F\left(\frac{\hat{p}}{\Omega_l Q_l}\right) - F(0)\right](S_l^n - c). \tag{1.12}$$

[16] The explicit function of S_l^n is given by

$$\Omega_l s\mathrm{e}^{-\sigma l} + (1-\alpha)\int_0^l s\lambda \mathrm{e}^{-(\sigma+\lambda)t}. \tag{1.11}$$

Because the salvage value is always less than the marginal cost, the second part of the right-hand side of (1.12) is negative. When the seller is required to accept all returns, they have to endure the sure loss from selling the product to consumers with a value lower than the keeping threshold. From Proposition 4, we know that $\Omega_l Q_l$ is decreasing in l. This implies $\Omega_l Q_l < Q_0$ for any $l \leq 0$. Compared to the seller's profit of the baseline model, equation (1.6), it is clear that the seller is better off in the baseline model where they can decline the non-broken returns.

Given any $(\hat{p}, l)$, the consumer surplus can be derived similarly. The potential consumers only return the product when it breaks during the return period. Therefore, the expected consumer surplus for a potential consumer is $U_0(v) - \hat{p}$. The non-potential consumers will return the product. Their expected surplus strictly comes from the return period. We denote this expected surplus for a non-potential consumer as $R_l(v)$.[17] Therefore, the total consumer surplus is given by

$$\int_0^{v_k} R_l(v)\mathrm{d}F(v) + \int_{v_k}^{\overline{v}} \left[U_0(v) - \hat{p}\right] \mathrm{d}F(v). \tag{1.14}$$

The right-hand side of (1.14) is the surplus to non-potential consumers, and the right-hand side is the surplus to the potential consumers. Because of larger consumer participation, the total consumer surplus can potentially be higher than the baseline model. However, it is not always the case. The following Proposition provides the condition under which the consumer surplus decreases if the seller has a no-questions-asked return policy.

Proposition 5 (Consumer Welfare Comparison)**.** *Suppose $s\lambda(1-\alpha) > m$. There exist an $\tilde{s} < c$ and a finite $\tilde{\mu}$ such that accepting all returns decreases the total consumer surplus if $s \geq \tilde{s}$ and $\mu \geq \tilde{\mu}$ (i.e., consumers are less price sensitive enough).*

[17] The explicit function of $R_l(v)$ is

$$(1-\alpha)\int_0^l \left(\int_0^t v\mathrm{e}^{-r\tau}\mathrm{d}\tau\right) \lambda\mathrm{e}^{-\lambda t}\mathrm{d}t + \Omega_l \int_0^l v\mathrm{e}^{-rt}\mathrm{d}t. \tag{1.13}$$

We now give an intuition for why the easy return policy can result in a lower consumer surplus compared to a stricter return policy in the baseline model. When the no-questions-asked return policy is implemented, the seller faces a trade-off related to the return period. On the one hand, extending the return period might increase the surplus on each trade with a potential consumer (i.e., increase the seller's return advantage) because each product will be more likely to be salvaged when it breaks. On the other hand, by Proposition 4, extending the return period makes the potential consumers more sensitive to the expected expenditure, which eventually becomes the seller's income. In other words, the seller will lose some potential consumers if they lengthen the return period while placing the same expected expenditure (by increasing the market price).

If the consumers are willing to pay a high markup for the product (i.e., they have low price sensitivity), losing potential consumers will become too costly compared to the gain in salvaging more broken products. Therefore, the seller will give up the ability to salvage broken products by choosing a short return period to retain their potential consumers. As a result, the non-potential consumers will enjoy less surplus. The potential consumer surplus makes up most of the total consumer surplus. These potential consumers, however, have to endure a high expected cost because the seller gives up the ability to salvage broken products. Whereas if the seller can decline non-broken returns, they can extend the return period to salvage more broken products and offer a contract with lower expected expenditure to the consumers.

1.7 Conclusion

We develop a continuous time model in which consumers constantly learn about the quality type of their products during use and can return a product during a certain period. By deriving a solution that fully captures consumers' shopping and returning outcomes, we characterize the seller's optimal pricing and return policies, including the return period's length. Endogenizing the return period allows us to explore the novel relationship between product quality and pricing. The model also endorses us to study the economic implications when the sellers are required to accept returns.

In summary, we find that extending the return period makes the potential consumers more sensitive to their expected expenditure, provided that they can return any product. Accepting all returns can decrease consumer surplus if consumers are less price sensitive. The after-sales returns also reveal some insights about the impacts of product quality on pricing. For instance, a higher market price can imply a lower product quality if the sellers can extract a high salvage value from a broken product. Product quality also helps explain the heterogeneity in the return periods among sellers. We show that an increased probability of getting a good quality product results in a shorter return period. However, the effect of the breaking down rate on the return period is not monotone.

Chapter 2

Product Return Policies under Duopoly: Online and Local Stores

2.1 Introduction

Return policy plays a role in consumers' purchasing decisions. According to a survey of 1660(US) and 1783(UK) people by the Global Web Index in November 2018, almost 80% of people check the returns policy on a retailer's website before purchasing. I develop a duopoly model in continuous time to study the effect of the return policy under competition. In this model, an online seller and a local store sell the same products. The setting for each consumer is based on the previous chapter. When the consumer purchases, she will be matched with a product. Each consumer gets an idiosyncratic flow of utility. Consumers observe this information before purchasing but are unsure if they will encounter conflicts in the future. If the match is reliable, the consumer will not face any conflicts. However, if the match is unreliable, the match is subject to conflicts at any time. Consumer utility will drop when the conflict happens. The consumers do not know the type of their match. Consumers can return the product during the return period, which the seller determines. The seller issues a full refund when the consumer sends back the item. The seller can extract some value from the returned product because of their recycling or refurbishing ability. Compared to the online store, the local store offers the product hands-on experience to the consumers. Consumers will receive informative signals about their match type when they visit the local store. However, obtaining such information is costly, such as traveling to the local store.

Using the reformulation in the previous chapter, I find the closed-form solution for the return periods when both sellers co-exist and fully characterize the coexistence equilibrium. I find that accepting returns can provide market power to both firms under price competition. The local store can strategically target consumers with lower utility by offering informative signals. Since returning a product is a hassle[1], these consumers are willing to pay the information cost at the local store to save them from returning the products when they encounter conflicts. On the other hand, the online store has an advantage over the consumers with higher utility. Those consumers are more likely to buy the product, no matter what signal they get. Therefore, they prefer buying from the online store to the local store. These two groups of consumers are not only different in utility but also in the probability of getting a well-matched product. I find that the online store offers a more extended return period and higher prices than the local store. The local store's return period decreases in the precision of the signal they provide.

Among other findings, I show that costly informative signals at the local store are necessary to maintain the market power for both stores. An increase in information costs will disperse the market prices of the two sellers. This change reduces the local store's market power but relatively increases the online store's competitiveness. The local store has to reduce the price to keep their consumers, while the online store can charge a higher price. This pricing

[1]Davis et al. (1998) show that sellers can control the moral hazard return behavior by imposing hassles on returns. Hess et al. (1996) show that sellers can also impose non-refundable charges, such as shipping or handling fees, to decrease this behavior.

strategy also gives the online store more market share while reducing the local store's market share.

Under duopoly, Inderst and Tirosh (2015) find that high-quality sellers offer a more generous return policy. This result was also found by Grossman (1981) and Moorthy and Srinivasan (1995). In the previous chapter, I introduced a monopolistic model and studied the effect of product quality on return policy in a more general dimension. I find that sellers can offer more generous terms with lower product quality. I also find that lower product quality can lead to a higher price if the consumers are less patient. Petrikaitė (2018) studies a single-product duopoly where consumers shop sequentially. She finds that a decrease in the cost of return leads to a higher symmetric equilibrium price. The results that accepting returns leads to a higher price can also be found in the previous chapter for single-product monopolists or Anderson et al. (2009) for the multi-product monopolist.

The literature has studied other effects of accepting returns. Che (1996) shows sellers can use return policies as insurance for risk-averse consumers. Escobari and Jindapon (2014) demonstrate that return policy can be served as a tool for price discrimination. Hinnosaar and Kawai (2020) show that return policy can be used to reduce seller uncertainty about the buyer's signal structure.

The rest of the paper is organized as follows. Section 2.2 introduces the model. Section 2.3 provides consumer shopping and returning outcomes. Section 2.4 characterizes the equilibrium where both sellers are active. Section 2.5 analyzes each seller's optimal return period and price. Section 2.6 provides a set of comparative statics results regarding match quality, signal precision, and information cost. Finally, Section 2.7 offers some concluding remarks. All omitted proofs are in the Appendix.

2.2 Model

The market consists of local and online sellers and one unit mass of consumers. Time is continuous with an infinite horizon. The market starts at time $t = 0$. Each consumer has one unit of product demand and gets an idiosyncratic utility flow v from using it. The utility flow depreciates at rate δ, universal to all consumers. For each consumer, v is independently and identically drawn from the interval $[\underline{v}, \overline{v}]$ according to a distribution function F with a density function f. Consumers observe their v before purchase.

Each consumer gets matched with a product when she purchases. The match can be reliable (R) or unreliable (U). Consumers do not know their match type before purchasing. I denote θ as the *state of nature* for each consumer, $\theta \in \{\theta_R, \theta_U\}$. If the match is reliable, consumers get the utility flow of $v\mathrm{e}^{-\delta t}$ at any time t without conflicts. Unreliable match experiences conflict at a Poisson rate λ. When the conflict occurs, the utility flow drops to zero from that period forward. The probability that a match is reliable is α. The utility of not buying the product is zero. Consumers are risk-neutral.

Each consumer receives a binary signal σ that can be R or U when they visit the local store. There is a cost of T to obtain such information (e.g., traveling cost). The signal σ is the realization of a random variable that depends on the true state. Table 2.1 illustrates the signal structure. If the match is reliable, the probability of receiving signal R is q_R. If

States \ Signal	$\sigma = R$	$\sigma = U$
$\theta = \theta_R$	q_R	$1 - q_R$
$\theta = \theta_U$	$1 - q_U$	q_U

Table 2.1: Signal structure

the match is unreliable, the probability of receiving signal U is q_U. The signal structure is informative if and only if $q_R \neq 1 - q_U$, or $q_R + q_U \neq 1$. Without loss of generality, I assume $q_R > 1 - q_U$, or $q_R + q_U > 1$, so that signal R increases the belief in state R and signal U

increases the belief in state U.

Consumers can return the product for a full refund within the return period l, which the seller determines. However, there is a return hassle h for consumers to return their products. The sellers can verify if the consumer has encountered any conflicts or not. Therefore, the sellers can impose extra hassle h_e if the consumer returns a product without evidence of conflict. The sellers incur a production cost of c for each product and get a salvage value of $s < c$ for each returned item. To serve each unit of consumers throughout the return period l, the customer service costs the sellers an amount of ml.

Timing. When the market opens at $t = 0$, the sellers announces their contracts (p, l, h_e). Consumers see the contracts and decide whether to visit a local store or buy from an online store or not to buy the product. They can decide where to buy the product if they visit the local store. The seller incurs the retail cost and customer service cost. The buyers can return the product anytime during the return period $(0 < t \leq l)$. Seller issues full refunds when consumers return the items.

2.3 Consumer behavior

I refer to the consumer's subjective probability that her match is reliable as the *consumer's belief.* As a direct result of the previous chapter, the belief of a typical consumer who has not experienced any conflict is

$$\pi_t = \frac{\alpha}{(1-\alpha)\mathrm{e}^{-\lambda t} + \alpha}. \tag{2.1}$$

When the market opens $(t = 0)$, each consumer has to decide whether to get the signal from the local store or go straight to the online option or not to buy the product. Once they buy the product, they must decide whether to keep it or keep using it during the return period. If they encounter a conflict, they should return the product before the return period expires.

The previous chapter provides the solution to the consumers' dynamic problem in a monopoly setting. The solution can be described in a static fashion by the two thresholds v_b and v_k. A consumer buys the product if and only if $v > v_b$ and commits moral hazard return if and only if $v < v_k$. The contract determines these two thresholds.

In the duopoly setting, each seller has a pair of thresholds v_b and v_k if their contracts differ. For each seller, v_b and v_k do not determine the entire consumer shopping and returning outcome. However, if a consumer buys the product from the seller, her utility v has to be higher than v_b. Otherwise, this buying option is dominated by the outside option. Also, she will keep the product with no conflict at the end of the return period if and only if $v > v_k$.

For those who decide to buy the product, if their $v > v_k$, their expected utility, and net refund will not be affected if the seller increases the extra hassle h_e. Their decision to buy from this seller will remain the same if the seller increases h_2. It is because they will not exercise the moral hazard return. The number of consumers who buy and keep the products after the return period will not decrease. An increase in extra hassle does not make these consumers leave the seller for another seller or opt out.

If $v_b < v_k$, some consumers buy the product with v falls into this range. An increase in extra hassle can make these consumers give up the moral hazard return option. The following proposition provides the optimal level of extra hassle for each seller.

Proposition 6. *For each seller, it is optimal for them to set the extra hassle cost h_e such that their corresponding buying threshold v_b is greater than or equal to their corresponding keeping threshold v_k.*

In equilibrium, consumers will not have the incentive to return a product without a conflict.

Their expected utility from buying the product is

$$
\begin{aligned}
U(v,\pi) &= (1-\pi)\int_0^\infty \left[\int_0^y v\mathrm{e}^{-\delta x}\mathrm{d}x\right]\lambda\mathrm{e}^{-\lambda y}\mathrm{d}y + \pi\int_0^\infty v\mathrm{e}^{-\delta x}\mathrm{d}x \\
&= \frac{(1-\pi)v}{\delta+\lambda} + \frac{\pi v}{\delta} \\
&= \frac{v(\delta+\pi\lambda)}{\delta(\delta+\lambda)}.
\end{aligned} \tag{2.2}
$$

If the consumer buys the product from the online store without visiting the local store, her expected utility will be $U_\alpha = \frac{v(\delta+\alpha\lambda)}{\delta(\delta+\lambda)}$.

The consumer only buys the product if the expected utility and net refund are higher than the price. In other words, opting out is dominated if the expected utility is higher than the effective price $\hat{p}$, which is the price minus the expected net refund,

$$
\hat{p} := p - (p-h)(1-\pi)\int_0^l \lambda\mathrm{e}^{-\lambda t}\mathrm{d}t. \tag{2.3}
$$

The consumer will buy from the seller that has the highest $U - \hat{p}$. For further analysis, I denote subscript "o" for the online store and "l" for the local store. For example, $\hat{p}_o$ is the online store effective price, and $\hat{p}_l$ is the local store effective price.

Suppose the consumer visits the local store and decides where to buy the product. In this case, the consumer must pay T upfront. If the consumer receives the "good" signal (i.e., $\sigma = R$), her belief about the match type will then increase from α to

$$
\pi_R = \frac{\alpha q_R}{\alpha(q_R+q_U-1)+1-q_U}. \tag{2.4}
$$

Her expected utility from using the product now increases to $U_R = \frac{v(\delta+\pi_R\lambda)}{\delta(\delta+\lambda)}$. She will choose to buy from a seller that offers the lowest expected price or opt out if both expected prices are higher than the expected utility. In this case, the online effective price $\hat{p}_o(\alpha)$ will increase to $\hat{p}_o(\pi_R)$. It is because the match is more likely to be reliable, and the chance of return will decrease. It makes the expected refund decrease and leads to a higher effective price.

Similarly, if the consumer receives the "bad" signal (i.e., $\sigma = U$), her belief about the match type will decrease from α to

$$\pi_U = \frac{\alpha(1-q_R)}{q_U - \alpha(q_R + q_U - 1)}. \tag{2.5}$$

Her expected utility from using the product now decreases to $U_U = \frac{v(\delta+\pi_U\lambda)}{\delta(\delta+\lambda)}$. The online effective price will decrease from $\hat{p}_o(\alpha)$ to $\hat{p}_o(\pi_U)$. In the following analysis, I will go through all of the possible decisions after the consumer visits the local store:

Case 1: Good signal - local store; Bad signal - local store

If a consumer chooses this option, no consumer buys from the online store. It is because the consumer buys from the local store no matter what signal she gets. The expected utility is the same if the consumer buys from the online store or visits a local store and buys from it, which is U_α. Since this consumer buys from the local store, it must be true that

$$[\alpha q_R + (1-\alpha)(1-q_U)]\,\hat{p}_l(\pi_R) + [\alpha(1-q_R) + (1-\alpha)q_U]\,\hat{p}_l(\pi_U) + T < \hat{p}_o. \tag{2.6}$$

The expected cost they pay from the local store is lower than the effective price of the online store. It means buying from the online store is strictly dominated by going to the local store and buying from them no matter what signal is received.

After receiving the signal, the expected utility from the product will be the same between the online store and the local store, which is either U_R or U_U. Since there exists a consumer that chooses the local store over the online store, it must be true that $\hat{p}_l(\pi_R) < \hat{p}_o(\pi_R)$ and $\hat{p}_l(\pi_U) < \hat{p}_o(\pi_U)$. It means consumers will never choose to buy from the online store after visiting the local store, no matter what signal they get. Combined with the previous result, in this case, no consumer will end up buying from the online store.

Case 2: Good signal - local store; Bad signal - online store

The expected utility from visiting a local store and following this option or from buying straight from the online store are the same, which is U_α. If a consumer chooses to visit the store and follow this strategy, it must be true that

$$[\alpha q_R + (1-\alpha)(1-q_U)]\,\hat{p}_l(\pi_R) + [\alpha(1-q_R) + (1-\alpha)q_U]\,\hat{p}_o(\pi_U) + T < \hat{p}_o. \tag{2.7}$$

The left-hand side of the inequality is the expected cost of visiting the local store and following the given strategy. It means buying straight from the online store is strictly dominated by this strategy. Therefore, all of the consumers who purchase choose to pay for the cost T to get the information. The following equilibrium is that both sellers set the lowest prices possible and get zero profit. However, this is not an equilibrium because the inequality (2.7) will not hold.

Case 3: Good signal - online store; Bad signal - online store

This strategy is strictly dominated by buying straight from the online store. Because by doing that, consumers can save the traveling cost of T to the local store.

Case 4: Good signal - online store; Bad signal - local store

Similar to case 2, all consumers who purchase will get the signal information, which is not an equilibrium.

Case 5: Good signal - online store; Bad signal - opt out

If a consumer chooses to buy from the online store after receiving the good signal, then $\hat{p}_o(\pi_R)$ has to be lower than $\hat{p}_l(\pi_R)$. It means consumers will never buy from local stores if they get a good signal. Since the strategies in case 3 and case 4 cannot happen in equilibrium, if a consumer chooses to buy from the online store after receiving the good signal, no consumer will buy from the local store. The following proposition sums up the findings.

Proposition 7. *In the equilibrium where both sellers make a positive profit, consumers who visit the local store buy from the local store if they receive a good signal (i.e., $\sigma = R$) and opt out if they receive a bad signal (i.e., $\sigma = U$).*

2.4 Market equilibrium

In the equilibrium, each seller sets the extra hassle cost high enough for $v_b > v_k$. To find the equilibrium price and return period for each term, it is sufficient to consider p and l as the only two control variables for each seller. Let $\Pi_l(p_l, l_l)$ and $\Pi_o(p_o, l_o)$ be the local and online store profits, respectively. The combination $(p_l^*, l_l^*, p_o^*, l_o^*)$ is an equilibrium if $\Pi_l(p_l^*, l_l^*) \geq \Pi_l(p_l, l_l)$ for all $p_l, l_l \geq 0$ given $p_o = p_o^*$, $l_o = l_o^*$; and $\Pi_o(p_o^*, l_o^*) \geq \Pi_o(p_o, l_o)$ for all $p_o, l_o \geq 0$ given $p_l = p_l^*$, $l_l = l_l^*$.

Since in the equilibrium, consumers who visit the local store will only buy from the local store if they receive a good signal and opt out otherwise, I denote $\hat{p}_l$ be the effective local price corresponding with $\sigma = R$. Therefore,

$$\hat{p}_l(p_l, l_l) = p_l - (p_l - h)(1 - \pi_R) \int_0^{l_l} \lambda \mathrm{e}^{-\lambda t} \mathrm{d}t. \tag{2.8}$$

The expected payoff of going to the local store (denoted by B_l) will then become

$$B_l = \left[\alpha q_R + (1 - \alpha)(1 - q_U)\right] (U_R - \hat{p}_l) - T. \tag{2.9}$$

If consumers buy from the online store, their belief will remain at $\pi = \alpha$, and the effective online price will be

$$\hat{p}_o(p_o, l_o) = p_o - (p_o - h)(1-\alpha)\int_0^{l_o} \lambda \mathrm{e}^{-\lambda t}\mathrm{d}t. \tag{2.10}$$

Therefore, the expected payoff of buying from the online store is $B_o = U_\alpha - \hat{p}_o$.

The consumer will choose to go to the store that has the highest expected payoff or not go to any store if both of the expected payoffs are less than zero. Figure 2.1 illustrates the consumer buying outcome.

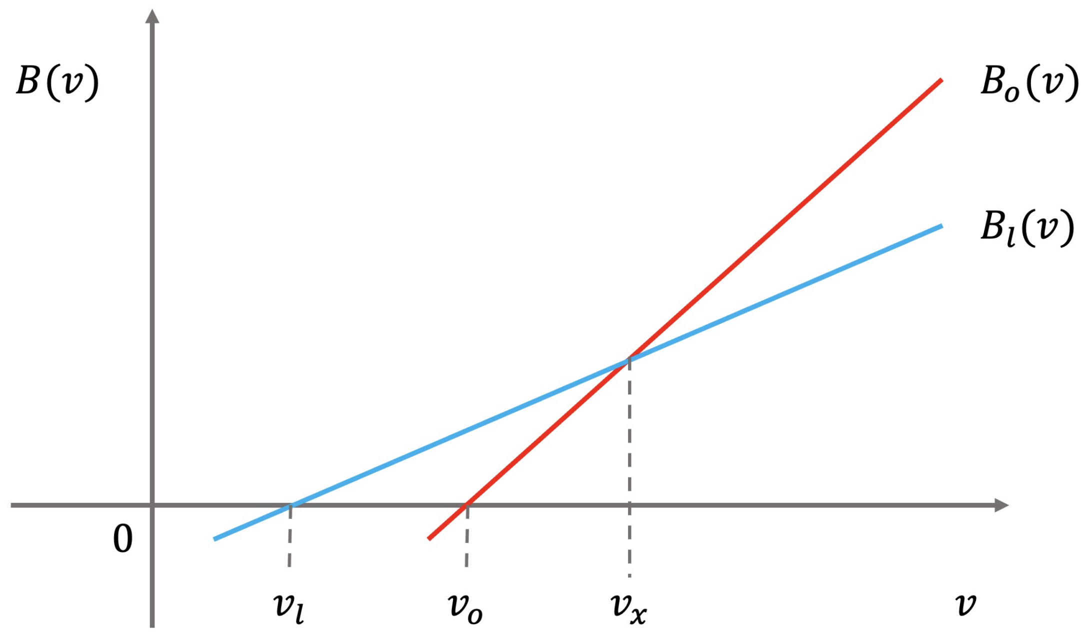

Figure 2.1: Consumer expected payoff

I denote v_l, v_o,v_x be the values such that $B_l(v_l) = 0$, $B_o(v_o) = 0$, and $B_l(v_x) = B_o(v_x)$. In Figure 2.1, the consumer will buy from the online store without visiting the local store if $v > v_x$ and visit the local store if $v \in (v_l, v_x)$.

The slope of $B_l(v)$ is smaller than $B_o(v)$ because the expected utility of using the product from visiting the local store is smaller than the online store. The consumer only buys the product at the local store if she receives a good signal. However, the consumer will purchase the product if she chooses the online option. It is the key factor that raises a coexistence equilibrium, where both sellers can get positive profit.

The local store can only sell the products if they can make the consumers who receive good signals buy them. It means their effective price $\hat{p}_l$ has to be lower than the effective online price with good signal (denoted as $\hat{p}_{os}$). Otherwise, consumers will just showroom the local store for more information and make the purchase at the online store. In this case, the expected payoff from visiting the local store becomes

$$B_{showroom} = [\alpha q_R + (1-\alpha)(1-q_U)]\,(U_R - \hat{p}_{os}) - T, \tag{2.11}$$

which is strictly greater than B_l. Figure 2.2 illustrates this case where $\hat{p}_l > \hat{p}_{os}$. When

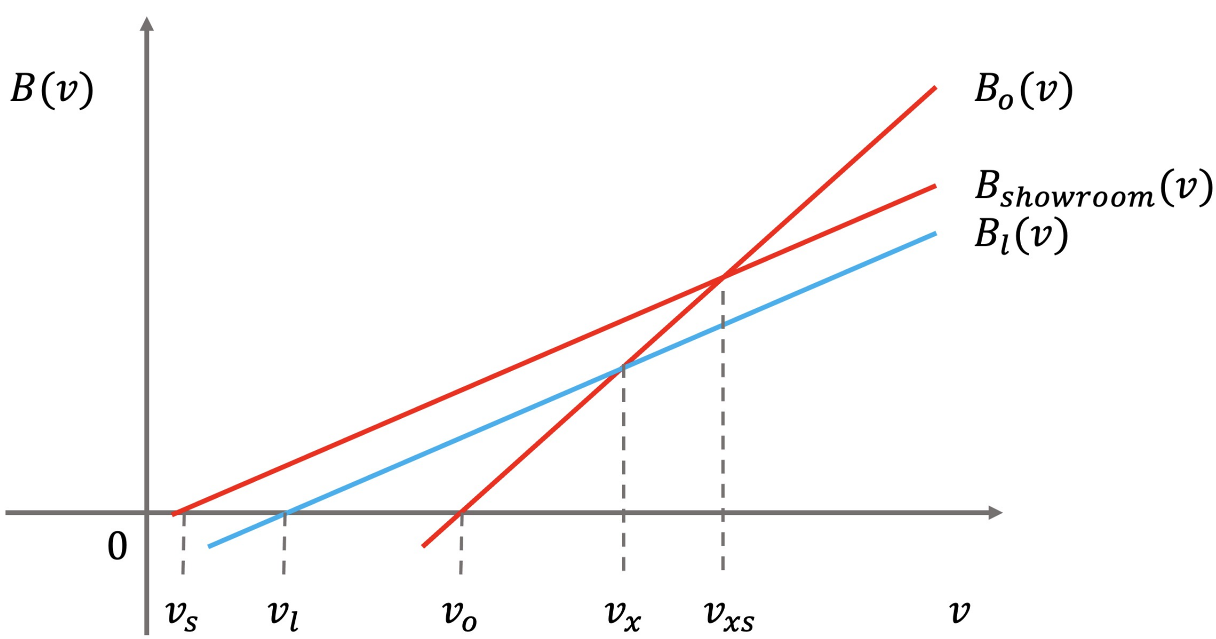

Figure 2.2: Consumer expected payoff with showrooming

$\hat{p}_l > \hat{p}_{os}$, no consumer will end up buying from the local store. I denote v_s and v_{xs} be the values such that $B_{showroom}(v_s) = 0$ and $B_{showroom}(v_{xs}) = B_o(v_{xs})$. Consumers with $v > v_{xs}$

will buy from the online store without visiting the local store. Consumers with $v \in (v_s, v_{xs})$ will visit the local store but not buy the products there. They will buy the products from the online store if they receive a good signal.

For each product, the consumer is expected to pay the effective price. However, this is different from what the seller expects to get. Together with the service cost, for each product, the seller is expected to get

$$p - (p - s)(1 - \alpha) \int_0^l \lambda \mathrm{e}^{-\lambda t} \mathrm{d}t - ml.$$

The difference between what the firm expects to get and the effective price is mentioned in the previous chapter as the *seller's return advantage* and denoted as

$$A(l) := (s - h)(1 - \alpha) \int_0^l \lambda \mathrm{e}^{-\lambda t} \mathrm{d}t - ml. \tag{2.12}$$

The local store's profit can be described as follows,

$$\Pi_l(p_l, l_l) = \begin{cases} \beta \left[F\left(\max\{v_o, v_x\}\right) - F\left(\min\{v_l, v_o\}\right) \right] (\hat{p}_l + A_l - c) & \text{if } \hat{p}_l < \hat{p}_{os} \\ 0 & \text{if } \hat{p}_l \geq \hat{p}_{os}, \end{cases}$$

where $\beta := [\alpha q_R + (1 - \alpha)(1 - q_U)]$ is the probability of getting a good signal at the local store.

If $\hat{p}_l < \hat{p}_{os}$, consumers will not showroom. The local store only gets some consumers if and only if $v_l < v_0$, or equivalently $v_o < v_x$. Otherwise, consumers will only visit the local store. Therefore, $F\left(\max\{v_o, v_x\}\right) - F\left(\min\{v_l, v_o\}\right)$ represents how many consumers visit the local store. Among them, only a fraction of β receive a good signal and make the purchase. The expected profit of each product sold upfront is $\hat{p}_l + A_l - c$. If $\hat{p}_l \geq \hat{p}_{os}$, all consumers will end up showrooming, and the local store will not have any sales.

The online store's profit can be described as follows,

$$\Pi_o(p_o, l_o) = \begin{cases} [1 - F(\max\{v_o, v_x\})](\hat{p}_o + A_o - c) & \text{if } \hat{p}_l < \hat{p}_{os} \\ \\ [1 - F(\max\{v_o, v_{xs}\})](\hat{p}_o + A_o - c) + & \\ \beta [F(\max\{v_o, v_{xs}\}) - F(\min\{v_s, v_o\})](\hat{p}_{os} + A_{os} - c) & \text{if } \hat{p}_l \geq \hat{p}_{os}. \end{cases}$$

If $\hat{p}_l < \hat{p}_{os}$, there will be no showrooming. Buying from the online store strictly dominates visiting the local store if $v > v_x$, but is strictly dominated by not buying from anywhere if $v < v_o$. Therefore, $1 - F(\max\{v_o, v_x\})$ represents how many consumers buy from the online store.

If $\hat{p}_l \geq \hat{p}_{os}$, there will be two groups of consumers that buy the product online. The first group is the consumers that do not visit the local store. For these consumers, showrooming is strictly dominated by buying straight from the online store. Therefore, the number of these consumers are $1 - F(\max\{v_o, v_{xs}\})$. The other group of consumers is from the local store. The number of consumers that visit the local store with the intention of showrooming is $F(\max\{v_o, v_{xs}\}) - F(\min\{v_s, v_o\})$. Only those with a good signal will purchase at the online store among these consumers.

In the coexistence equilibrium, the local store has to be able to make the consumers who visit their store receive a good signal to buy from them. The following proposition describes the coexistence equilibrium.

Proposition 8. $(p_o^*, l_o^*, p_s^*, l_s^*)$ *is a coexistence equilibrium if and only if it satisfies all of the following*

- $l_o^* = \underset{l_o \geq 0}{\operatorname{argmax}} \, A_o(l_o)$

- $l_l^* = \underset{l_l \geq 0}{\operatorname{argmax}} A_l(l_l)$

- $\hat{p}_o^* = \underset{\hat{p}_o > c - A_o(l_o^*)}{\operatorname{argmax}} \left[1 - F\left(\max\{v_o, v_x\}\right)\right]\left(\hat{p}_o + A_o(l_o^*) - c\right)$

- $\hat{p}_l^* = \underset{\hat{p}_l > c - A_l(l_l^*)}{\operatorname{argmax}} \left[F\left(\max\{v_o, v_x\}\right) - F\left(\min\{v_l, v_o\}\right)\right]\left(\hat{p}_l + A_l(l_l^*) - c\right)$

- $\hat{p}_s^* < \hat{p}_{os}^*$

It is not a unique equilibrium but a unique coexistence equilibrium.

2.5 Optimal return periods and prices

In the monopoly setting, the previous chapter finds that the return period will maximize the seller's return advantage. I show that this result still holds under a duopoly between online and local sellers. The reason is that by offering more information, such as hands-on experience at the local store, the local store can obtain some loyal customers with lower utility. Returning a product is relatively costly to these consumers. They appreciate the option of knowing that they might get an unreliable match so that they will not buy the product and endure the return hassle. They are more willing to pay for the information to reduce the risk of return. The return process allows the local seller to attract this type of consumer. On the other hand, consumers will higher utility will find the online seller a better option. These consumers are more likely to buy the product no matter what signal they get.

The after-sales returns provide the two sellers with two different types of consumers. The return policy and the local store's ability to provide more information are the keys to market power. Therefore, the return period maximizes the return advantage for each seller. The following Corollary provides each seller's closed-form solution for the return period.

Corollary 2. *In the coexistence equilibrium, the return periods of the online store and local store are:*

$$l_o^* = \begin{cases} \frac{1}{\lambda}\ln\left(\frac{(s-h)(1-\alpha)\lambda}{m}\right) & \textit{if } (s-h)(1-\alpha)\lambda > m \\ 0 & \textit{if } (s-h)(1-\alpha)\lambda \leq m \end{cases}$$

$$l_l^* = \begin{cases} \frac{1}{\lambda}\ln\left(\frac{(s-h)(1-\pi_R)\lambda}{m}\right) & \textit{if } (s-h)(1-\pi_R)\lambda > m \\ 0 & \textit{if } (s-h)(1-\pi_R)\lambda \leq m. \end{cases}$$

In the previous chapter setting, the online store return period will remain the same if it is the only store in the market. However, if consumers showroom the local store in the equilibrium, the online store becomes the only store to sell their product, and their return period might differ. Besides the consumers with the belief α, another group of consumers with a higher chance of getting a reliable match to make the online purchase.

Corollary 3. *In the coexistence equilibrium, if $(s-h)(1-\alpha)\lambda > m$, the online store offers a longer return period than the local store.*

For the local store, their targeted consumers are the ones that receive the good signal. Therefore, their consumers are more likely to get a reliable match. Their marginal return advantage from each product is lower than the online store. The local store will invest less in the return process and offer a shorter return period compared to the online store.

Corollary 4. *In the coexistence equilibrium, the local store offers a lower price than the online store.*

The local store targets consumers with lower utility, and they also need to keep these consumers from showrooming. Given that their optimal return period is shorter than the online store, the local store has to offer a lower price to make the consumers buy from them after visiting.

2.6 Comparative statics

2.6.1 Match quality

Match quality is measured in two dimensions: probability α that a match is reliable and the Poisson rate λ that an unreliable match encounters conflicts.

Proposition 9. *Suppose that $(s-h)(1-\pi_R)\lambda > m$ and the coexistence equilibrium exists. Then both the return periods of online and local stores decrease in α.*

This result is intuitive. When the match is more reliable, the consumers who buy from the online store are less likely to return their products. The consumers who visit the local store will be even more sure that their match is reliable if they receive a good signal. Therefore, both sellers will invest less in the return process and offer a shorter return period.

Since the return period maximizes the seller's return advantage, a result regarding λ from the previous chapter will still hold in this model.

Proposition 10. *Effect of λ on the return periods in the coexistence equilibrium.*

- *Suppose that $(s-h)(1-\pi_R)\lambda > m$. Then the local store's return period l_l^* increases in λ if $\lambda < \frac{em}{(1-\pi_R)(s-h)}$, and decreases in λ if $\lambda > \frac{em}{(1-\pi_R)(s-h)}$.*
- *Suppose that $(s-h)(1-\alpha)\lambda > m$. Then the online store's return period l_o^* increases in λ if $\lambda < \frac{em}{(1-\alpha)(s-h)}$, and decreases in λ if $\lambda > \frac{em}{(1-\alpha)(s-h)}$.*

An increase in λ has two effects on the marginal return advantage from each product. The first effect is that it increases the chance of return of any unreliable match in any given period. However, the second effect is that conflicts tend to happen quicker. The number of remaining unreliable matches that have not encountered any conflicts decreases at any time.

This reduction goes at an exponential rate λ. Therefore, if λ is large, an increase in λ will reduce each product's marginal return advantage and result in a shorter return period.

2.6.2 Signal precision

Corollary 2 shows that the online return period is not affected by the signal structure of the local store. However, by affecting the posterior belief of the consumers who visit the local store, the signal structure impacts the return period of this store.

Proposition 11. *Suppose that $(s-h)(1-\pi_R)\lambda > m$ and the coexistence equilibrium exists. Then the local store's return period l_l^* decreases in q_R and q_U.*

Proposition 11 shows that if the signal is more accurate, the local store will invest less in the return process. If the signal is a symmetric binary ($q_R = q_U$), the local store will reduce the return period if there is an increase in the precision of the signal.

2.6.3 Local store's information cost

Theoretically, without the information cost, all consumers will visit the local store before making decisions. It will eliminate the market power of both sellers, and no firm will get a positive profit in the equilibrium. However, with a positive information cost, the local store will become less competitive.

Proposition 12. *Suppose that $(s-h)(1-\pi_R)\lambda > m$, the coexistence equilibrium exists and*

$$\underline{v} = \underset{v}{\operatorname{argmax}}\ [1 - F(v)]\left[\hat{p}_l(v) + A_l(l^*) - c\right].$$

If $\frac{f}{1-F}$ is monotone non-decreasing and $\frac{f}{F}$ is monotone non-increasing, then local store price p_l^* and local store's market share decrease in traveling cost T, while online store price p_o^* their *market share increase in T.*

An increase in the information/traveling cost will make the local store lose its competitiveness and give the online store more market power. Therefore, the local store has to reduce their price to keep their consumers while the online store can charge a higher price. A higher information/traveling cost will increase the price dispersion and reduce the market share of the local store while improving the online store's market share.

2.7 Conclusion

I developed a duopoly model of online and local store sellers. The setting is based on the previous chapter. Both sellers have the advantage of extracting value from returned products. This advantage can be the ability to recycle or refurbish the products. I find the closed-form solution for the return periods when both sellers co-exist and fully characterize the coexistence equilibrium. I show that the comparative statics results regarding the return period in the previous chapter are robust under competition.

By considering the after-sales returns, I find that the return policy can provide market power to both firms under price competition. By offering an informative signal, the local store will target consumers with lower utility. These consumers are willing to pay the information cost at the local store to save them from returning the products when they encounter conflicts. On the other hand, online store consumers have higher utility. Those consumers are more likely to buy the product, no matter what signal they get. Therefore, they prefer buying from the online store to the local store. These two groups of consumers are not only different in utility but also in the probability of getting a well-matched product. I find that the online store offers a more extended return period and higher prices than the local store. The higher precision of the signal, the lower the local store's return period.

I show that an increase in the information cost, such as the traveling cost to the local store, will make the market prices of the two sellers more dispersed. This change reduces the local store's market power but relatively increases the online store. The local store has to reduce the price to keep their consumers, while the online store can charge a higher price. It also gives the online store more market share while reducing the local store's market share.

Chapter 3

Consumer Search, Collusion, and Artificial Intelligence

3.1 Introduction

Internet and technological advancements have been transforming our lives. Besides the apparent aspects that enhance the quality of life, new technology can affect market competitiveness in a less transparent way. This paper will look at how technological advances impact consumer shopping behavior and eventually affect the firms' incentive to collude and raise markups.[1] Our research involves two factors that are transforming many markets. (1) A reduction in search friction on both the demand and supply sides. For example, in e-commerce, consumers can easily find product information, such as prices, product details, and ratings. It is common nowadays for consumers to check and compare prices online before purchasing. Sellers can also monitor their competitors' prices and adjust their actions

[1] In the past four decades, there has been a rising in markups in the U.S. economy. De Loecker, Eeckhout, and Unger (2020) find that aggregate markups start to rise from 21% above marginal cost in 1980 to 61% in 2020.

quickly. Less friction usually relates to higher market competition. However, its effect on sellers' incentive to collude is not translucent. (2) The use of algorithmic pricing is becoming common.[2] These pricing software help sellers react quickly to competitors' actions (e.g., changing prices) and market dynamics. It is a natural choice for sellers in high-frequency markets such as e-commerce, trading, and transportation. Algorithmic pricing powered by AI can autonomously monitor the market and develop its pricing strategy with little initial input. At first glance, adopting algorithmic pricing may improve price competition because it creates high-frequency price matching. However, scholars and policymakers have raised concerns that algorithmic pricing can foster collusion and raise prices.[3]

This paper asks how search friction reduction affects collusion; and what implications it has on market prices and consumer welfare. We tackle this question both theoretically and experimentally. We conduct our experiments by simulating sellers' interactions with AI-powered pricing algorithms. We then investigate if the algorithms learn to collude in the frictional market and compare the theoretical and experimental results.

To lay out the framework, we present an oligopoly model in which consumers sequentially search for the best product with advertised prices. A consumer's payoff from purchasing a product from seller i is $z_i - p_i$ where p_i is the price, and z_i is only revealed when she incurs a search cost. Firms can influence consumer search behavior by coordinating prices. Through this mechanism, we study the sustainability of collusion among firms. The consumer search framework is similar to Choi, Dai, and Kim (2018). Our paper analyses collusion outcomes in repeated settings when sellers use a price-trigger strategy (Green and Porter, 1984) as

[2] Chen, Mislove, and Wilson (2016) estimated that in 2015, there were over 500 among 1,641 most popular products on Amazon involving algorithmic pricing.

[3] In their empirical study about Germany's retail gasoline market, Assad, Clark, Ershov, and Xu (2020) find that adoption of algorithmic pricing increases margins by 9%. Brown and MacKay (2022) conduct an empirical analysis of the online allergy drugs market in the U.S. and estimate that algorithmic pricing increases average prices by 5.2% relative to the simulated counterfactual Bertrand competition. Calvano, Calzolari, Denicolò, and Pastorello (2020) and Klein (2021) provide experimental evidence that pricing algorithms can maintain high prices by collusive strategies. Concern about collusion when sellers use algorithmic pricing has been raised in OECD (2017) and a remark from the U.S. Department of Justice (Delrahim, 2018).

opposed to static Bertrand-Nash competition in Choi, Dai, and Kim (2018).

Our model is most related to e-commerce because it captures some essential features of online shopping behavior and the mechanism by which firms can collude. Online shopping typically involves typing the product name or a keyword into the search box of an e-commerce platform. The website will then show the list of related products or sellers with prices and brief descriptions for each option. Consumers can click on any option to learn more detailed information about that product. Our model can speak to this shopping behavior well. Building from this demand structure, we show how the cartel can coordinate their prices to maximize the joint profit and how a firm can deviate from collusion by manipulating the consumer search order.

Weitzman (1979) provides an optimal search strategy for each consumer, including the search order and the stopping rule. Utilizing this elegant solution, we find that the lower search costs make it easier for firms to collude. This finding is distinctive because it differs from existing consumer search literature results and raises attention to antitrust policies, especially for e-commerce marketplaces.

We show that lower search costs make undercutting the cartel price less profitable. This outcome is opposite to classical results in the literature, which stem from Wolinsky (1986) sequential search framework with unobservable prices. More consumers happen to visit the deviating firm as the search costs decrease. The profit obtained by the undercutting firm will then increase as consumers are more willing to search. In our model, with observable prices, the deviating firm can manipulate consumers to visit them first by undercutting the collusive price. As the value of searching for another option increases, consumers are more likely to leave the seller to explore other options and eventually reduce the demand for the deviating firm. Lower search costs make deviation less tempting and collusion relatively more sustainable.

On the other hand, higher search costs can reduce the price even when sellers collude. This result aligns with Choi, Dai, and Kim (2018), which suggests price falls with search costs in static Bertrand competition. The reason is that higher search cost makes being visited first more valuable and increases price competition among sellers. Even if sellers collude, the cartel still needs to compete with consumer non-search options (outside options), leading the lower price with higher search costs. However, we find that this price reduction is insufficient to offset the search burden on consumers if the collusion sustains. Therefore, higher search costs decrease consumer surplus when sellers collude.

For our experimental approach, we follow Calvano, Calzolari, Denicolò, and Pastorello (2020) to simulate sellers' interactions with reinforcement learning algorithms. Agents in Calvano, Calzolari, Denicolò, and Pastorello (2020) are symmetric and engage in a repeated Bertrand competition. Each period, agents simultaneously place prices, and consumers purchase their best products with no friction. We take a similar setup to Calvano, Calzolari, Denicolò, and Pastorello (2020) and plug the frictional demand structure on top. We find that Q-learning, a simple AI algorithm, can adopt a punishment scheme similar to the trigger-price strategy to maintain prices above the competitive level in the frictional market. However, AI price and consumer surplus both decrease with higher search costs. These experimental results are consistent with our theoretical results, which suggest that higher consumer search costs can lower the price and consumer surplus if collusion sustains.

The above finding raises an important question on regulating consumer search friction. Lowering search costs can increase consumer surplus per se; however, the sellers can collude more easily, and the cartelized market becomes more sustainable, eventually harming the consumers. It is crucial to characterize the collusion equilibrium that realizes in practice to find an optimal policy that prevents cartel formation. Theoretically, this task is difficult because collusion models typically have a large equilibrium set. However, with our experiment results, we can refine the equilibrium and make a tighter prediction on market outcomes.

We find that markets with AI pricing algorithms have a similar outcome to half-collusive equilibrium, where sellers randomize between fully collusive pricing and competitive pricing. This result can be especially beneficial in determining the optimal policy for markets where sellers prominently use AI-powered pricing software. For example, it can be used to find the optimal level of search friction that maximizes consumer surplus.

Our experimental results also amplify recent policy concerns about AI pricing algorithms and collusion (Mehra, 2016 and Ezrachi and Stucke, 2017). The issue is that these algorithms may autonomously learn to collude without having any explicit communication or agreement to establish a violation of Section 1 of the Sherman Act.[4] To address this policy issue, we must look at collusion from different angles and discuss what unlawful collusion is when sellers use algorithmic pricing.

The rest of the article is organized as follows. Section 2 reviews the literature and explains how our work contributes to the literature. We then introduce the baseline model in Section 3 and characterize the fully collusive equilibrium in Section 4. Symmetric collusive equilibrium is characterized in Section 5. We introduce our experiment design in Section 6 and present our experimental results in Section 7. Finally, we conduct the robustness check in Section 8 and offer some concluding remarks in Section 9. All omitted proofs are in the Appendix.

3.2 Literature review

On the theoretical aspect, our paper is related to a rich set of literature that studies the roles of market transparency on collusion. In particular, the literature focuses on market transparency from the sellers' viewpoint, such as Abreu, Pearce, and Stacchetti (1986), Compte (2002), Green and Porter (1984), Matsushima and Kandori (1998). Typically, it involves

[4] Harrington (2018) proposes a legal approach to prosecuting collusion by algorithms.

imperfect public monitoring where firms cannot observe each other actions perfectly.[5] The literature has shown that collusion is less sustainable when sellers' actions are not directly observable. As the sellers' actions become less transparent, the punishment for a deviation from collusion is milder.

Our paper is more related to market transparency from the consumer's point of view and how it affects the sustainability of collusion. The literature on this topic is much less than on the supply side. We are only aware of four articles that are related to this topic: Nilsson (1999), Campbell, Ray, and Muhanna (2005), Schultz (2005), and Petrikaitė (2016). These four papers focus on price transparency for consumers. In their models, some consumers either cannot observe the price or have to incur a positive search cost to explore it. Instead, our paper studies a different notion of transparency. In our model, consumers can observe the price but are uninformed about the match values.

Building on the non-sequential consumer search model of Burdett and Judd (1983), Nilsson (1999) studies the roles of consumer search costs for prices on the sellers' incentive to collude. In his model, two firms are selling homogeneous products. Some consumers perfectly observe the prices, while the rest have to incur a search cost to learn the prices. Nilsson (1999) shows that collusion becomes more stable as the search cost for the price decreases. In a similar setting, when consumers are paired with sellers in a certain way, Campbell, Ray, and Muhanna (2005) also show that lower search costs for prices can facilitate collusion.

By developing a duopoly model in the framework of Hotelling (1929), Schultz (2005) also studies how price transparency on consumers affects collusion. In his model, a fraction of consumers can observe both prices, and the rest are uninformed. The ratio of informed fraction measures price transparency. When the products are differentiated, Schultz (2005) finds collusion is less sustainable with more price transparency because a higher share of price-

[5]See Abreu, Pearce, and Stacchetti (1990) and Fudenberg, Levine, and Maskin (1994) for a comprehensive set of techniques for this class of problems.

informed consumers makes undercutting the collusive price more profitable. This finding differs from our results. In our model, consumers can observe the price but are uninformed about the match value. We show that sellers have less incentive to deviate from collusion when the search cost for match value decreases.

Among the abovementioned papers, our paper is most related to Petrikaitė (2016). She presents two oligopoly models with different consumer search frameworks and studies the roles of search costs on collusion. She first analyses collusion in the sequential search framework of Wolinsky (1986) with horizontally differentiated products. This framework is closest to ours because consumers in our models sequentially search for the best products. The difference is that in Wolinsky's framework, prices are not observable to consumers. Whereas in our model, consumers observe the price but have to search for the match values. Petrikaitė (2016) finds that in Wolinsky's framework, the cartel is more stable with higher search costs. The mechanism behind this result is that as search costs increase, fewer consumers happen to visit a deviating firm, leading to lower deviation profit. This finding is opposite to our result, which suggests deviation is more profitable with higher search costs. The reason is that consumers can observe the price in our model and search sequentially based on the prices they see, whereas, in Wolinsky's framework, consumers search randomly without knowing the price. The deviating firm in our model will get the consumer's attention and be visited first; in Petrikaitė (2016), consumers visit the deviating firm by chance.

Petrikaitė (2016) also studies collusion in the consumer search framework of Stahl (1989) with homogeneous products. In her model, a positive fraction of consumers is fully informed about the prices. The rest of the consumers observe one price and incur a search cost to learn other prices. Petrikaitė (2016) shows that search costs do not influence the deviation profit. Therefore, higher search costs make collusion less easy to sustain. This finding is along the line with our results even though our models are built on different consumer search frameworks and aspects of search.

The price-directed search framework that we use in our model has been studied extensively in the context of static price competition (e.g., Armstrong and Zhou, 2011, Haan, Moraga-Gonzalez, and Petrikaite, 2018, Choi, Dai, and Kim, 2018, and Obradovits and Plaickner, 2022). The literature has shown that market prices fall with search costs. This result is similar to what we find in the repeated setting, which suggests that the collusive price also decreases with higher search costs. The reason for these results has been discussed in the previous section. A higher search cost generally increases the price competition whether the sellers collude or not.

Besides the theoretical literature, our paper contributes to the ongoing discussion about algorithmic pricing and collusion. This strand of literature draws the attention of scientists from different disciplines, legal scholars, and policymakers. Notably, recent studies have conducted experiments with AI-powered pricing algorithms to study their collusive strategies. Waltman and Kaymak (2008) put forward a study of repeated Cournot competition with reinforcement learning algorithms. They find that algorithms manage to reduce output, leading to a price higher than the Nash equilibrium price in a one-shot game. Calvano, Calzolari, Denicolò, and Pastorello (2020) study a Bertrand competition where pricing algorithms (Q-learning) interact repeatedly. They find that algorithms adopt a punishment scheme to keep the price above the competitive level. The collusive scheme is similar to the price-trigger strategy. Punishment is followed by a gradual return to the price before deviation. They also find that algorithms adopt this collusive scheme under imperfect monitoring (Calvano et al., 2021).

In our paper, we take a similar setup to Calvano et al. (2020) and incorporate search friction into the model. We also find that algorithms maintain non-competitive prices by collusive strategies in the frictional market. However, the reward-punishment scheme differs from Calvano et al. (2020), where the market has no friction. In both markets, punishment is followed by a gradual increase in prices. As opposed to a frictionless market, where algorithms return to the previous collusive price, in a frictional market, they converge to a lower price than before the deviation.

With a similar experimental design to Calvano et al. (2020), Johnson, Rhodes, and Wildenbeest (2022) study the effects of different platform designs on the market price and consumer surplus. Using the same reinforcement learning algorithms, Klein (2021) analyses the model of Maskin and Tirole (1988), where agents take turns to price sequentially. He analyses how algorithms coordinate with different the number of discrete prices. He finds that when the number of discrete prices is limited, Q-learning algorithms coordinate on a stable high price; when the number of discrete prices increases, Q-learning algorithms converge to a supra-competitive Edgeworth price cycling pattern.

3.3 Model

There are n sellers, each indexed by $i \in \{1, \ldots, n\}$. Each seller sells one differentiated product with the same constant marginal cost c. In each period $t = 0, 1, \ldots,$ each seller simultaneously announces a price p_i^t. The price history is observable to all sellers. The sellers interact repeatedly and discount the future at rate $\delta \in (0, 1)$.

In each period t, one unit mass of consumers enter the market and exit before the next period starts. New consumers come in the next period. Each consumer has one unit demand of the product and obtains utility v_i from seller i's product. For each consumer, v_i is an independent random variable of type I extreme value distribution with scale parameter σ and location parameter μ.

Consumers learn their v_i only if they visit seller i. They must visit seller i to be able to purchase product i. They can leave the market and take the outside option at any point. The outside option has the value u_0 for all consumers.

Each consumer has to incur a search cost of $s \geq 0$ when they visit a seller. The recall is costless (i.e., consumers can purchase from any previously visited seller without incurring the search cost again).

The maximum number of sellers a consumer can visit is $k < n$. The ex-post utility for a representative consumer who buys product i after visiting m sellers is equal to

$$U(v_i, p_i, m) = v_i - p_i - sm.$$

If she takes the outside option after visiting m sellers, then her ex-post utility is $U(m) = u_0 - sm$. Consumers are risk neutral and follow a strategy that maximizes their expected utility.

Denote the prices of all sellers in period t by $\overrightarrow{p^t} \equiv \{p_1^t, p_2^t, \ldots, p_n^t\}$. The demand for seller i in period t is denoted by $D_i(\overrightarrow{p^t})$ Seller i's profit in period t is then given by $\pi_i^t = (p_i^t - c)D_i(\overrightarrow{p^t})$. Each seller maximizes their discounted expected flow of profit:

$$E\left[\sum_{t=0}^{\infty} \delta^t \pi_i(\overrightarrow{p^t})\right].$$

The equilibrium concept that we focus on is the subgame-perfect Nash equilibrium.

3.3.1 Consumer behavior

After observing the prices, each consumer faces a classic sequential search problem. Before she reaches the maximum observation, she has to decide which seller to visit next or to stop searching and take the best available option. Weitzman (1979) provides a general solution for this problem with $k \geq n$. The solution consists of a search order and a stopping rule. Each seller has a "reservation" value z_i based on the prior distribution of their product utility realization. Weitzman (1979) shows that for consumers, it is optimal to visit sellers with a descending reservation value order and stop searching when the best available option exceeds the reservation values of unvisited sellers.

With $k < n$, F, and f being the c.d.f. and p.d.f. of random variable v, the optimal strategy is presented as follows.

Proposition 13 (Optimal search strategy)**.** *Let v_s be the value such that*

$$s = \int_{v_s}^{\infty} (1 - F(v))\, dv. \tag{3.1}$$

Define the reservation value for the seller i by $z_i \equiv v_s - p_i$ and the reward of choosing product i by $x_i \equiv v_i - p_i$. Consumer optimal strategy is characterized as follows.

1. *Selection Rule: rank all sellers in a non-ascending reservation value order and visit them by this order. If there is more than one non-ascending reservation value order, randomly select one of them.*

2. *Stopping Rule: stop searching and take the best available option if the maximum reward among visited sellers or u_0 exceeds the reservation value of every unvisited seller or if k sellers have been visited.*

Proposition 1 applies even when products are ex-ante different from sellers (i.e., sellers have different cumulative distribution functions $F(v)$. In our model, sellers are ex-ante symmetric and share the same v_s. The search order, therefore, only depends on prices. Sellers will be visited in a non-descending price order.

If the maximum number of observations is non-binding (i.e., $k \leq n$), an elegant result of the Weitzman solution is that we can express the consumer's final decision in a discrete choice fashion. By assigning a new index $w_i \equiv \min\{z_i, x_i\}$ to each seller and $w_0 = u_0$ to the outside option, it can be shown that the chosen option has to have the highest index (see Theorem 1 in Choi et al. (2018) and Armstrong (2017)).

If the maximum number of observations is binding (i.e., $k \leq n$), the eventually chosen option cannot be pinned down entirely by the index w. However, the index w can still help narrow down the chosen seller, as shown in the following Corollary.

Corollary 5. *For a given optimal search order, let K be the set of k first sellers and $w_i \equiv \min\{z_i, x_i\}$ $\forall i \in K$. A consumer buys from seller $i \in K$ if and only if:*

- $w_i \geq w_j$ $\forall j \in K_{-i}$ *and* $w_i \geq u_0$ *and*
- *seller i is ranked before seller j in the search order if $w_j = w_i$.*

Proof. Let $W \equiv \{w_1, w_2, \ldots, w_k\}$. The set of all realizations of W can be divided into two subsets: S_1, where only one seller has the highest w, and S_2, where two or more sellers share the same highest w. The probability that a realization of W is in S_2 is strictly positive if a group of sellers (more than one) in K has the same price. Suppose $W \in S_1$, corollary 1 holds by the eventual purchase theorem in Choi et al. (2018).

Suppose $W \in S_2$.

$\rightarrow$: Suppose $w_i \geq w_j$ $\forall j \in K_{-i}$ and $w_i \geq u_0$ and seller i is ranked before seller j if $w_j = w_i$. Because the realization is in S_2, besides seller i, another seller j has the same $w_j = w_i$. Because the probability that $x_i = x_j$ or $x_i = z_j$ or $x_j = z_i$ is zero, sellers i and j have to have the same reservation values (i.e., $z_i = z_j$) and the rewards are higher than the reservation values (i.e., $x_i > z_i$ and $x_j > z_j$). Because seller i is ranked before seller j, all other sellers before i have to have w lower than w_i. By Proposition 1, seller i will be visited. Because $x_i > z_i$, the consumer will stop and buy the product from seller i.

$\leftarrow$: Suppose the consumer buys the product from seller i. Therefore, $w_i = \max W$. Because

$W \in S_2$, there exists $j \neq i$ and $j \in K$ such that $w_j = \max W$. If by contradiction, seller j is ranked before seller i in the search order. By Proposition 1, the probability that the consumer buys from product i is zero. □

3.3.2 Demand structure

Denote by $p_{(k)}$ the k^{th} order statistic price among all sellers. By Proposition 1, all sellers with prices $> p_{(k)}$ will get demand of zero. Divide all sellers with prices $\leq p_{(k)}$ into groups such that sellers have the same price in the same group. Denote by G be the set of all groups. For each group $i \in G$, denote by

- P_i, N_i the price and number of sellers in that group, respectively;
- $|N_i|$ the number of sellers in group i that will get visited if a consumer visits all k sellers, meaning $|N_i| = N_i$ if $P_i < p_{(k)}$, and $|N_i| = k - \sum_{j \in G_{-i}} N_j$ if otherwise;
- $F_{w_i}(x)$ the CDF of w_i, meaning $F_{w_i}(x) = F_{x_i}(x)$ for $x < v_s - P_i$ and $F_{w_i}(x) = 1$ for $x \geq v_s - P_i$.

By Corollary 1, the eventual demand for each seller can be derived by discrete choice fashion and is presented in the following Corollary.

Corollary 6. *The total demand for each group $i \in G$, denoted by D_{G_i}, is*

$$|N_i| \int_{-\infty}^{v_s - P_i} F_{w_i}^{|N_i|-1}(x) \prod_{j \in G_{-i}} F_{w_j}^{|N_j|}(x) dF_{w_i}(x) + \sum_{i'=1}^{|N_i|} (1 - F(v_s)) F^{i'-1}(v_s) \prod_{j \in G_{-i}} F_{w_j}^{|N_j|}(v_s - P_i). \tag{3.2}$$

All sellers in group i have the same demand of D_{G_i}/N_i.

We define the competitive market equilibrium as an equilibrium in which no seller can gain profit in any period by a one-shot deviation. In each period, sellers choose prices the same as the one-shot Bertrand-Nash equilibrium price. If the one-shot Bertrand-Nash equilibrium price exists at all periods, this price also is the subgame-perfect Nash equilibrium price. It is simple to show that under a competitive market equilibrium, k lowest prices equal the marginal cost.

3.3.3 Collusive strategy

We consider the price-trigger strategy similar to Green and Porter (1984) as follows.

(a) Collusion phase: sellers place collusive price according to $\overrightarrow{p_c}$ for 1 period.

(b) Punishment phase: sellers place punishment price at marginal cost (i.e., competitive price level) for L periods, which can be infinite.

The collusion phase is assigned initially and proceeds with the following rules:

1. If a phase ends with no deviation, go to a new collusion phase in the next period.
2. If a deviation happens during any phase, go to a new punishment phase in the next period.

The collusion is considered sustainable if no seller has the incentive to make a one-shot deviation. To illustrate the sustainable condition more formally, we denote the profit per period for each firm as follows:

- π_c is the profit when all sellers collude,

- π_d is the maximum profit to the deviating firm when all other sellers place the collusive price,
- π_p is the profit when all sellers place punishment price ($\pi_p = 0$),
- π_{dp} is the maximum profit to the deviating firm when other sellers place the punishment price ($\pi_{dp} = 0$).

In the collusive phase, there is no profitable one-shot deviation if

$$\sum_{t=0}^{\infty} \delta^t \pi_c \geq \pi_d + \sum_{t=1}^{L} \delta^t \pi_p + \sum_{t=L+1}^{\infty} \delta^t \pi_c. \tag{3.3}$$

In the punishment phase, if deviate, it is best for the seller to deviate at the beginning period. There is no profitable one-shot deviation if, in the initial period of the punishment phase, no seller has the incentive to deviate,

$$\sum_{t=0}^{L-1} \delta^t \pi_p + \sum_{t=L}^{\infty} \delta^t \pi_c \geq \pi_{dp} + \sum_{t=1}^{L} \delta^t \pi_p + \sum_{t=L+1}^{\infty} \delta^t \pi_c. \tag{3.4}$$

The collusion is sustainable if both conditions (3.3) and (3.4) are satisfied.

3.4 Full collusion

In this section, we consider the case where all firms choose the price as if they are a monopoly with n product lines.

3.4.1 Market outcomes

Thus sellers will maximize the joint profit in each period. Sellers can obtain the highest joint profit if its k lowest prices are similar to prices of a k-product monopoly. For simplicity, the following analysis is about the k-product monopoly. Let $\pi(\vec{p})$ be the profit to the k-product monopoly where $\vec{p} = \{p_1, p_2, ..., p_k\}$. Let $D_i(\vec{p})$ be the demand for seller i for all $i \in \{1, 2, ..., k\}$. The k-product monopoly is then given by

$$\pi(\vec{p}) = \sum_{i=1}^{k}(p_i - c)D_i(\vec{p}).$$

When $s = 0$, the reservation value of each seller will be arbitrarily large, and each consumer will visit all k sellers. Consumers buy from seller i if and only if the reward from i exceeds all other sellers and the outside option. If the reward from seller i exceeds the outside option, v_i has to be larger than $u_0 + p_i$. If the reward from seller i exceeds seller j, v_j has to be smaller than $p_j - p_i + v_i$. Therefore, the demand for seller i is given by

$$D_i(\vec{p})\big|_{s=0} = \int_{u_0+p_i}^{\infty} \left(\prod_{j \neq i}^{k} F(p_j - p_i + v) \right) f(v)dv.$$

Lemma 3. *Suppose $s = 0$. For any $\vec{p}$ with $p_i < p_j$ then*

$$\frac{\partial \pi}{\partial p_i} > \frac{\partial \pi}{\partial p_j}$$

A direct result from Lemma 3 is that if there is no search cost, the k-product monopolist will set all prices the same. This result significantly reduces the complexity of finding the k-monopolist prices by focusing on symmetric prices only. When all k sellers are visited and have the same price, a consumer will buy a product if at least one seller has a reward higher than the outside option value. Therefore, given the symmetric price p, the joint profit is

$$\pi(p) = \left(1 - F^k(u_0 + p)\right)(p - c). \tag{3.5}$$

Because F is log-concave, by Prekopa Theorem, $\pi(p)$ has a unique maximizer p^* (technical details are in the Appendix). Therefore, for $s = 0$, k-product monopoly would set all prices equal to the unique p^*.

When $s > 0$, we cannot use Lemma 3 to simplify the profit-maximizing problem. We also have to consider the case where k prices are different. It is now helpful to introduce two new notations as follows. When considering between visiting a seller with the advertised price p and taking the outside option, it is convenient to introduce a new value $\overline{s}_p$ as the *maximum search cost* that the consumer is willing to visit that seller. That is to say, $\overline{s}_p = E\max\{v - p - u_0, 0\}$.[6] Therefore, seller i will not get visited if $s > \overline{s}_{p_i}$. Similarly, let $\overline{p}_s$ be the *maximum price* for a given search cost. That means if $p_i > \overline{p}$, visiting seller i is always dominated by taking the outside option, and seller i will not get visited.

Lemma 4. *Let $p_{max} = \max\{p_1, \ldots, p_k\}$ and $p_{min} = \min\{p_1, \ldots, p_k\}$.*

- *If $p_{max} > p_{min}$, then π is strictly decreasing in s for $s \leq \overline{s}_{p_{max}}$ and is monotone non-increasing in s for $s \geq \overline{s}_{p_{max}}$.*
- *If $p_{max} = p_{min}$, then π is constant in s for $s \leq \overline{s}_{p_{max}}$ and is zero for $s > \overline{s}_{p_{max}}$.*

The first part of Lemma 4 states that if the prices are not uniform, the joint profit will decrease with a higher search cost. The intuition is that consumers will search for fewer

[6] We can express $\overline{s}_p$ more explicitly as

$$\overline{s}_p = \int_{u_0+p}^{\infty} (1 - F(v))\, dv$$

where $F(v)$ is the cumulative distribution function of the random variable v_i.

sellers when the search cost is high. High-price sellers will be less likely to be visited as they are at the last of the search order. Therefore, the demand shifts toward low-price sellers, leading to a lower joint profit.

However, the second part of Lemma 4 shows that the monopoly can hedge against losing profit by placing a uniform price profile, provided that the search cost is not too high. In this case, higher search cost reduces the consumer surplus but does not affect the cartel profit. The loss in consumer surplus is partly due to the increase in search cost per se and partly due to the inability to explore more options.

Previously, we show that in the most favorable situation (i.e., no search cost), the monopoly will choose a uniform price (p^*) to obtain the maximum profit. This profit is also the upper limit profit that the monopoly can obtain when the search cost is positive. The reason is that with a nonuniform price, Lemma 2 shows that joint profit is maximized with no search cost. On top of that, Lemma 1 shows that under no search cost, no nonuniform price can yield higher collective profit than the uniform price p^*.

Lemma 4 indicates that if the search cost becomes positive but sufficiently low ($s \leq \overline{s}_{p^*}$), the upper limit profit is still obtainable if all prices remain equal to p^*. When the search cost is sufficiently high ($s > \overline{s}_{p^*}$), the monopoly would choose the closest price to p^* with the constraint that at their price, the reservation value is not dominated by the outside option. That means the monopoly would choose a price such that it is equivalent to $v_s - u_0$.

To sum up, the monopoly would place uniform prices for all sellers, called p^M.

- $p^M = p^*$ for $s \leq \overline{s}_{p^*}$ and
- $p^M = v_s - u_0$ for $s > \overline{s}_{p^*}$.

This result leads us to our proposition about the fully collusive outcomes.

Proposition 14 (Full Collusion Outcome)**.** *Suppose δ and L are high enough that full collusion is sustainable. Fully collusive price and corresponding consumer surplus weakly decrease in the search cost.*

An increase in the search cost can reduce the fully collusive price. It is because the monopoly has to ensure consumers are willing to search for their products. However, the price reduction is insufficient to increase the consumer surplus if the collusion sustains.

3.4.2 Sustainability of collusion

We then look at the sustainability of the monopoly when sellers use a punishment scheme to maintain the monopoly price. That means when a seller deviates, for example, by undercutting the price, other sellers will punish the deviating seller's future profit by lowering the prices and competing more severely. Deviation from the cartel can bring the seller some instant profit gain but will also jeopardize returns in future periods. Therefore, if the sellers are forward-looking, meaning the discount factor δ is high, the cartel becomes more sustainable.

Full collusion is sustainable with the largest set discount factor when punishment length $L = \infty$. In this case, condition (3.4) is always binding, and the right-hand side of condition (3.3) is the lowest. This strategy is also known as the Nash-reversion (because marginal cost pricing is a Bertrand-Nash equilibrium). We denote by $\hat{\delta}$ the critical discount factor under the Nash-reversion strategy (i.e., $p_c = p^M$, $p_p = c$, $L = \infty$). Condition (3) must bind at the critical discount factor.

To pin down the critical discount factor, we must study the one-shot deviation profit π_d when other sellers place the full collusion price p^M. Placing a price higher than p^M is not desirable because the deviating firm will end up with no demand. Slightly lowering the price makes the seller be visited first and get a discrete jump in demand. To guarantee the existence of the best one-shot deviation, we assume the deviating price $p_d \leq p^M$. By Corollary 1, the demand for the deviating firm can be written as

$$\tilde{D} = \int_{v_s+p_d-p^M}^{\infty} f(v)dv + \int_{u_0+p_d}^{v_s+p_d-p^M} F^{k-1}(v - p_d + p^M)f(v)dv \tag{3.6}$$

The first integration is the probability that a consumer stops searching after visiting the deviating store and buys the product from them. It happens if and only if the reward from the deviating seller $v - p_d$ exceeds the reservation value of other sellers $v_s - p^M$. The second integration represents the probability that a consumer buys from the deviating firm but not during the first visit. It only happens if and only if (1) the reward from the deviating seller does not exceed the reservation value of other sellers, (2) the reward from the deviating seller is higher than the outside option, and (3) the deviating seller has the highest reward among all k visited sellers.

The role of consumer search cost is illuminating in equation (3.6). Recall that p^M is constant in s for $s \leq \overline{s}_{p^*}$ while v_s is strictly decrease in s at all s. That means given any $p_d \leq p^M$, $\tilde{D}$ is strictly increasing in s for $s \leq \overline{s}_{p^*}$. The intuition is contained within equation (3.6). The consumers who are indifferent between continuing searching and buying from the first seller (who have $v = v_s + p_d - P^M$ from the deviating seller) will switch to buying from the seller when the search cost increases.

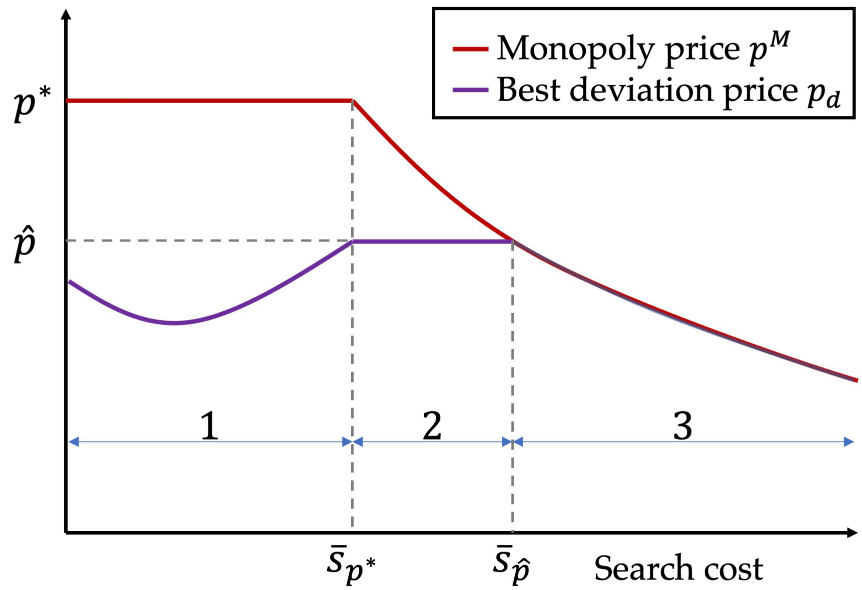

(a) Monopoly price and best deviation price

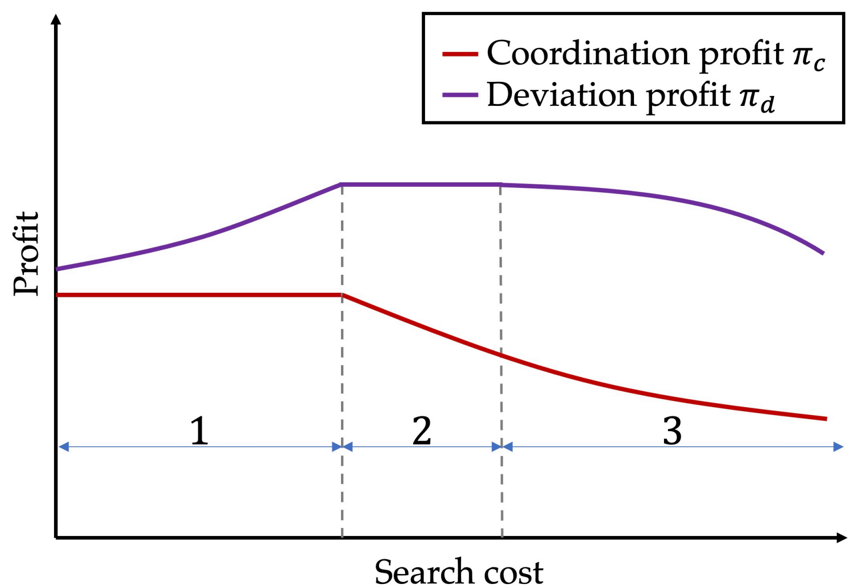

(b) Fully collusive profit and best deviation profit

Figure 3.1: Sustainability of full collusion

Because the demand for the deviating seller increases with a higher search cost, the direct result is that the maximum profit to the deviating firm is also strictly increasing in search cost. As a result, deviation becomes more tempting, and collusion is less sustainable. To formalize the result of the sustainability of the full collusion, we study how search cost affects the critical discount factor. At the critical discount factor $\hat{\delta}$, condition (3.3) is binding, which means

$$\frac{\pi_c(p^M)}{1-\hat{\delta}} = \pi_d. \tag{3.7}$$

For $s \leq \overline{s}_{p^*}$, fully collusive profit $\pi_c(p^M)$ is constant in s while deviation profit π_d increases in s. Therefore the critical discount factor $\hat{\delta}$ increases in search cost when $s \leq \overline{s}_{p^*}$. In fact, the critical discount factor increases with the search cost for the entire range. The technical proof is in the appendix.

Proposition 15 (Sustainability of Full Collusion)**.** *Full collusion is easier to sustain with lower search cost in the sense that the critical discount factor $\hat{\delta}$ increases in search cost s.*

As opposed to a market with a high price-transparency makes collusion easier to maintain because it is easier for sellers to punish deviation; a market with low consumer search cost makes collusion easier to maintain because it is less tempting for the sellers to deviate.

3.5 Symmetric collusion

In this section, we study a broader range of collusion. We allow sellers to collude at any price p_c above the competitive level (i.e., marginal cost).

3.5.1 Sustainability

Proposition 16 (Sustainability of Symmetric Collusion 1)**.** *Symmetric collusion is easier to maintain with lower search costs in the sense that the sustainable price set becomes smaller with higher search costs.*

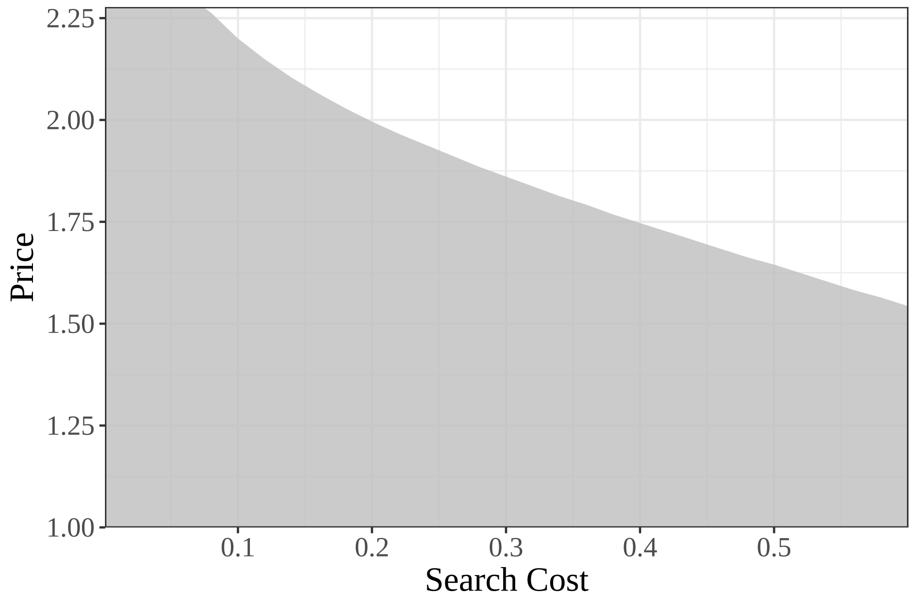

Figure 3.2: Sustainable collusion price set

Note. Simulations with marginal cost $c = 1$, discount factor $\delta = 0.95$, punishment length $L = \infty$, match value $v_i \sim \text{Gumbel}(2, 0.25)$, $n = 3$ and $k = 2$.

In general, the coordinated collusive price p_c can be considered as a random variable from a distribution $\mathcal{F}$ (can be degenerate). For example, one mechanism for sellers to coordinate on p_c is using a public randomization device. Full collusion is a special case when $\mathcal{F}$ degenerates at the monopoly price p^M. When $\mathcal{F}$ is not degenerate, p_c is a random variable, and the corresponding critical $\hat{\delta}$ is also a random variable. The collusion is fully characterized by the distribution $\mathcal{F}$. The following Proposition shows that $\mathcal{F}$ collusion is also harder to maintain with a higher search cost if the distribution $\mathcal{F}$ does not change with search cost.

Proposition 17 (Sustainability of Symmetric Collusion 2)**.** *Suppose $\mathcal{F}$ is the same for $s \in \{s_l, s_h\}$ and $s_l < s_h$. $\mathcal{F}$ collusion is easier to sustain with a lower search cost in the sense that the critical discount factor $\hat{\delta}$ increases in the first order stochastic dominance with higher search cost:*

$$F_{\hat{\delta}}(x|s = s_l) \geq F_{\hat{\delta}}(x|s = s_h)$$

for all for all x, with strict inequality at some x, and $s_l < s_h$.

To sum up, we have shown that the sustainability of collusion increases with lower search costs in various criteria. The critical factor to this result is that seller has more motivation to undercut the cartel price with higher search cost. Consumers will search less and tend to stay with the first seller they visit, leading to a higher demand for the deviating firm. In the following section, we will introduce our experiment design with AI pricing algorithms to study the effect of search cost on market outcomes and compare them with our theoretical results.

3.6 Experiment design

3.6.1 Q-learning

We simulate our model where each seller is equipped with independent reinforcement learning to help them price at each period. The reinforcement learning we use is Q-learning. This type of algorithm was first introduced by Watkins (1989) and has been used widely in the computer science literature.[7]

Q-learning is a reinforcement learning algorithm used widely due to its simplicity and the clear economic interpretations of its parameters. The algorithm involves iteratively updating a Q-value table, which maps state-action pairs to their expected rewards. This algorithm allows an agent to learn how to behave optimally in a given environment by selecting actions that maximize its expected reward.

Despite its simplicity and popularity, Q-learning is known to be slow to converge. This issue can be problematic in complex environments where the number of possible states and actions

[7]See Sutton and Barto (2018) for a comprehensive overview on Q-learning.

is large. However, convergence time is not typically an issue in our experiments with simpler setups. Therefore, Q-learning is often a good choice for these types of scenarios.

It is important to note that the choice of algorithm can have a significant impact on the outcome of an experiment. More complex algorithms may be able to adopt more sophisticated strategies. However, the choice of Q-learning can be appropriate to emphasize the results that even simple algorithms can exhibit collusive-like behavior, as will be shown in the next section.

To understand how Q-learning works, consider a dynamic programming problem in which, at every period, an agent observes a state $s \in S$ and can choose an action $a \in A$ where S and A are finite sets. For any state s and action a, the agent gets a payoff π and moves to a new state s'. The transition from (s, a) to (π, s') is a stationary Markov decision process, meaning the transition probability $F(\pi, s'|s, a)$ is time-invariant. The agent's problem is to find an optimal policy $a^*(s)$ that maximizes the expected discounted payoffs.

Instead of working directly with Bellman's value function $V(s)$, Q-learning works with Q-function $Q(s, a)$, which is the discounted payoff of taking action a in state s. The connection is $V(s) \equiv \max_{a \in A} Q(s, a)$. Therefore,

$$Q(s, a) = E(\pi|s, a) + \delta E\left[\max_{a' \in A} Q(s', a')|s, a\right]. \tag{3.8}$$

The optimal policy is then $a^*(s) = \arg\max_{a \in A} Q(s, a)$.

Q-learning is a method to estimate the Q-function through iteration. For each particular pair of s and a, Q-learning starts with an initial estimation value $Q_e(s, a)$ and then updates it iteratively over time (subscript e to represent that it is an estimation). It is convenient to think of Q-function as a matrix with one state dimension and one action dimension. There are $|S|.|A|$ cells in the Q-matrix.

Suppose the system is in state s and chooses action a. It then gets a payoff π and moves to state s' (s' can be the same as s). In this case, the value in the cell (s, a) of the Q-matrix will get updated from $Q_e(s, a)$ to $Q'_e(s, a)$ as following:

$$Q'_e(s, a) = (1 - \alpha)Q_e(s, a) + \alpha \left[\pi + \delta \max_{a' \in A} Q_e(s', a') \right]. \tag{3.9}$$

For this period, only the value in cell (s, a) gets updated. Other cells remain the same. The learning rate is $\alpha \in (0, 1)$. After an action is made in a state, one cell in the Q-matrix gets updated.

For the system to estimate the true Q-function, it needs to update all cells, meaning some degree of exploration is necessary. We specify the experimentation probability ϵ_t as a time-declining function

$$\epsilon_t = e^{-\beta t} \tag{3.10}$$

with $\beta > 0$. The algorithms choose a random action with probability ϵ_t and choose the action with the highest estimation value with probability $1 - \epsilon_t$.

3.6.2 Model simulation

We constructed n independent Q-learning agents and let them repeatedly interact in period $t = 0, 1, 2, \ldots$. Each agent chooses an action (price) from an action set in each period. We discretize the interval $[\underline{p}, \overline{p}]$ to get 15 equally spaced price points, including the boundaries, where $\underline{p}$ is lower than marginal cost c, and $\overline{p}$ is higher than the monopoly price when search cost is zero $p^M_{s=0}$. To be more specific, we set $\underline{p} = c - 0.05(p^M_{s=0} - c)$ and $\overline{p} = p^M_{s=0} + 0.1(p^M_{s=0} - c)$.

After agents choose their prices, the payoff (profit) to each agent in that period is given by $\pi_i = (p_i - c)D_i$, where the demand D_i is determined by Corollary 6.

We use the set of all price permutations from n sellers $\{p_1, p_2, \ldots, p_n\}$ as the state set. The state of a period is determined by the chosen price permutation of the last period. Hence, the state set contains 15^n elements. The Q-matrix is then a 15×15^n matrix.

For the initial Q-matrix at $t = 0$, we set the value of each cell $Q_e(s, a) = \pi_0(s)/(1 - \delta)$ where $\pi_0(s)$ is the expected payoff to the agent in one period when all other sellers uniformly randomize their actions. To start the simulation, we randomly choose the initial state at $t = 0$.

When the simulation runs, the Q-matrix is updated, cell by cell. The optimal strategy from estimation $a_e(s) = \arg\max_{a \in A} Q_e(s, a)$ will also change (not necessarily after each period) during the learning process. We say the simulation converges when the optimal strategy from estimation for each agent does not change for 100,000 consecutive periods. When a simulation converges, we stop the simulation and study adopted actions and the relevant outcomes upon convergence. For each experiment (a fixed set of parameters), we run 1000 independent simulations.

In our baseline specification, the number of agents $n = 3$, the maximum number of visits $k = 2$, and the outside option value $u_0 = 0$. For other parameterization, we take a similar setup to Calvano et al. (2020) with quality index $\mu = 2$, product differentiation $\sigma = 0.25$, discount factor $\delta = 0.95$, and learning specification $\alpha = 0.15$ and $\beta = 10^{-5}$.

3.7 Experiment results

In this section, we study the roles of consumer search cost by doing experiments with a sequence of s. We first present a set of results upon convergence regarding prices, welfare, and other metrics. We then compare the experimental results with our theoretical results. This comparison includes testing our comparative statics results regarding market price and

	AI simulations by cycle length				
	1	2	3	≥ 4	All
$s = 0.01$					
Frequency	0.049	0.241	0.172	0.538	1
$s = 0.3$					
Frequency	0.049	0.323	0.204	0.424	1
Average price change (%)	-2.46	-5.76	-4.72	-4.34	-4.36
Average search time change (%)	-40.01	-42.46	-42.10	-42.14	-42.12
Average consumer surplus change (%)	-51.07	-44.70	-43.23	-43.31	-44.59
Average seller surplus change (%)	-7.14	-12.90	-9.97	-7.48	-8.68
Average total surplus change (%)	-33.70	-31.46	-30.98	-30.49	-31.10

Table 3.1: Descriptive Statistics

consumer surplus and matching AI-simulated results with a theoretical equilibrium.

Upon convergence, agents rarely end up choosing the same price repeatedly. Only around 5% of all simulations for $s \in [0, 0.3]$ that agents adopt the same price in every period. For the majority of simulations, agents loop over a set of prices. When the simulation converges, the algorithms update a cycle of cells in the Q-matrix. We refer to the number of cells in the loop as the price cycle of the simulation. More than 50% of the simulations have a price cycle length of 2 or 3 (for s is at the middle of our experiment range).

Figure 3.4 shows the AI average price compared to the fully collusive and competitive prices. The thin box plots show the distribution of AI prices. There is a clear downward-slopping trend of AI prices with consumer search costs. Compared to no search cost, average AI prices decrease by 4.36% when $s = 0.3$ and decrease by 7.41% when $s = 0.5$ We also find that AI agents manage to choose non-competitive prices in a frictional market. They adopt a reward-punishment scheme to keep their price above the competitive level. This scheme is similar to the price-trigger strategy in Green and Porter (1984), as punishment immediately follows a deviation. All algorithms reduce the price right after one algorithm undercuts the

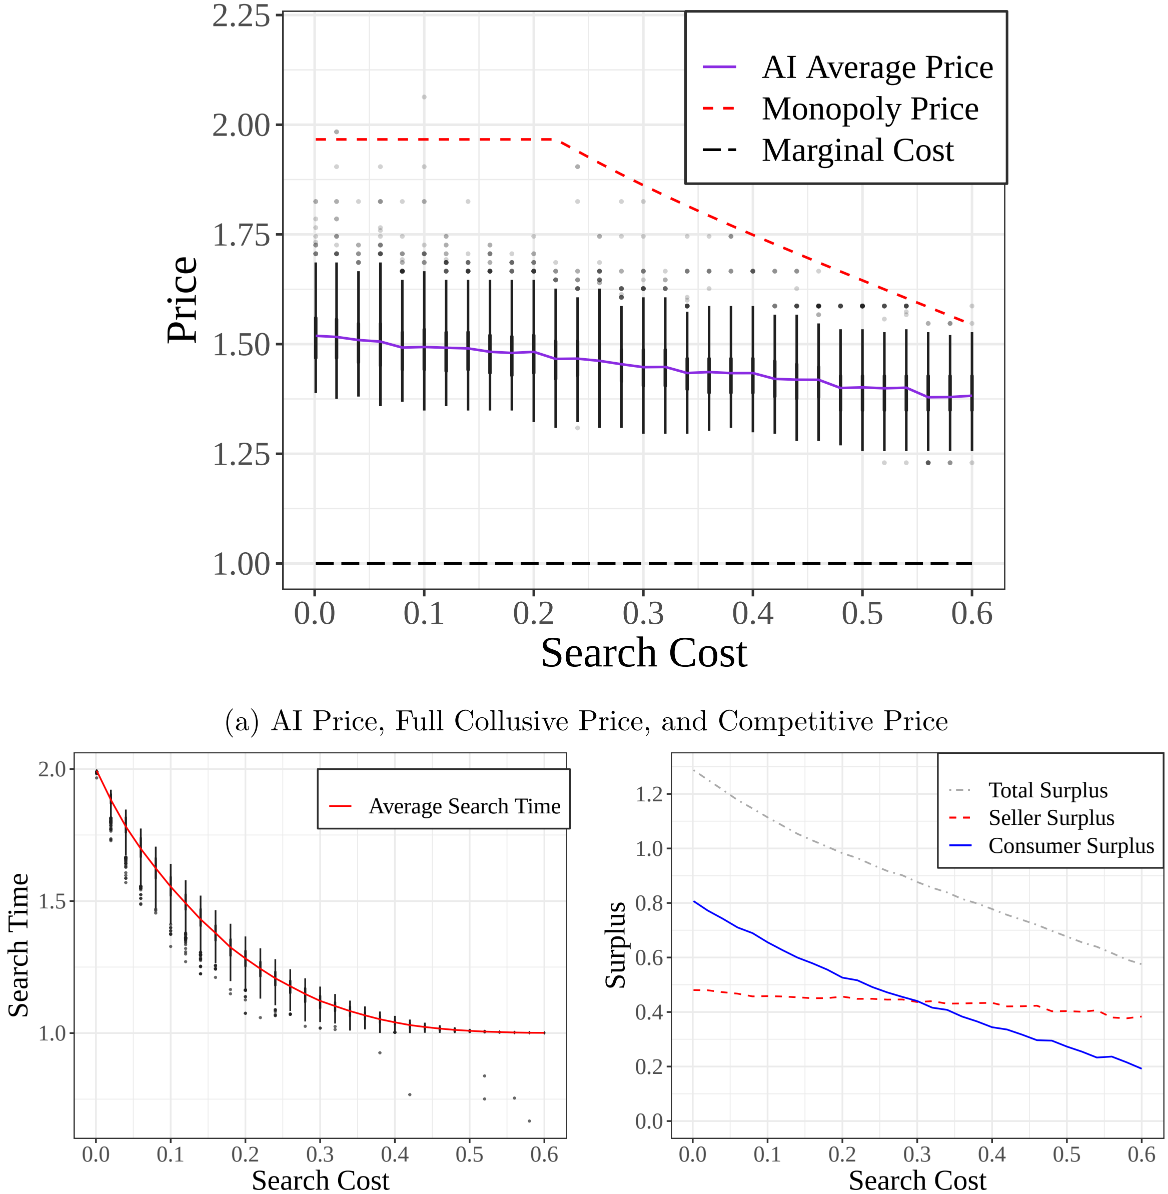

(a) AI Price, Full Collusive Price, and Competitive Price

(b) Consumer search time

(c) Welfare

Figure 3.3: AI prices, consumer search time and welfare

long-run price and then gradually increases the price. However, we find that AI algorithms do not return to the previous long-run price in the frictional market. They end up with a lower price than before the deviation. This result contrasts with Calvano et al. (2020), where algorithms in a market with no friction manage to return to the original price after a deviation.

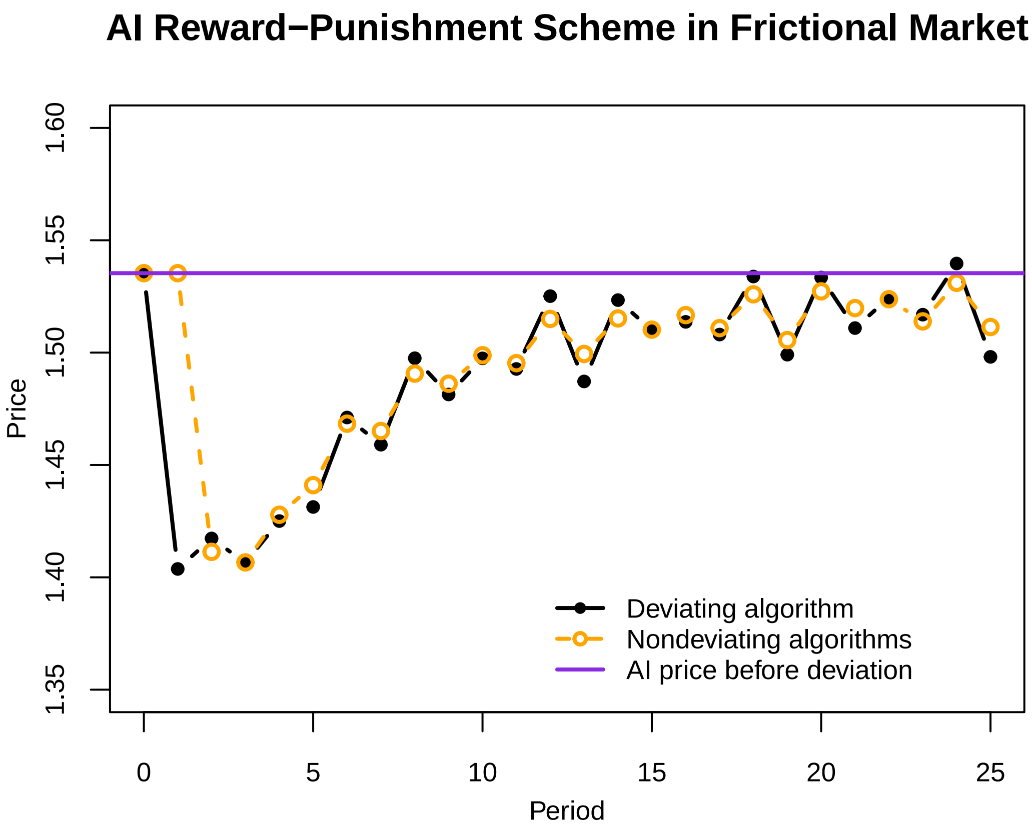

Figure 3.4: AI impulse response to deviation

Note. Market is frictional with $s = 0.2$. After AI strategies stabilize, we force one algorithm to deviate to the best static response. Then we let all algorithms respond to the deviation by letting them run with their optimal strategies in the following periods.

With a higher search cost, AI agents are less likely to maintain the high prices. This experimental result is consistent with our theoretical results in two dimensions: (1) collusion is harder to maintain with higher search costs, and (2) collusive prices decrease with higher search costs.

As search costs increase, consumers are less willing to visit new stores. This change in consumer behavior has two effects: it reduces the demand for the sellers at the end of the order and increases the demand for front sellers, as consumers are more likely to stop searching and choosing them. That means undercutting the price to improve position in the consumer search order becomes more tempting. A slight reduction in the price could earn the seller a discrete jump in profit if they move the front of the search order. Therefore, a higher search cost makes the market as a whole become more competitive and brings down the price. Figure 3b shows the consumer average search time, which reflects that price competition intensifies as search cost increases.

Even though AI prices decrease with a higher search cost, it is not enough to offset the cost it incurs to consumers (figure 3c). This result is consistent with our theoretical results, which predict that consumer surplus reduces with higher search costs if collusion sustains. The experiment results match our theoretical comparative static results regarding the market price and social welfare. In the following analysis, we will look at how well the market outcome with AI agents fits into a certain equilibrium.

3.8 Half collusion comparison

This section compares our experiment results with the half-collusive equilibrium. In the half-collusive equilibrium, sellers uniformly randomize between fully collusive price and marginal cost. Figure 6 shows how the AI price fits into our theoretical equilibrium set, the shaded area, and the average half-collusive price.

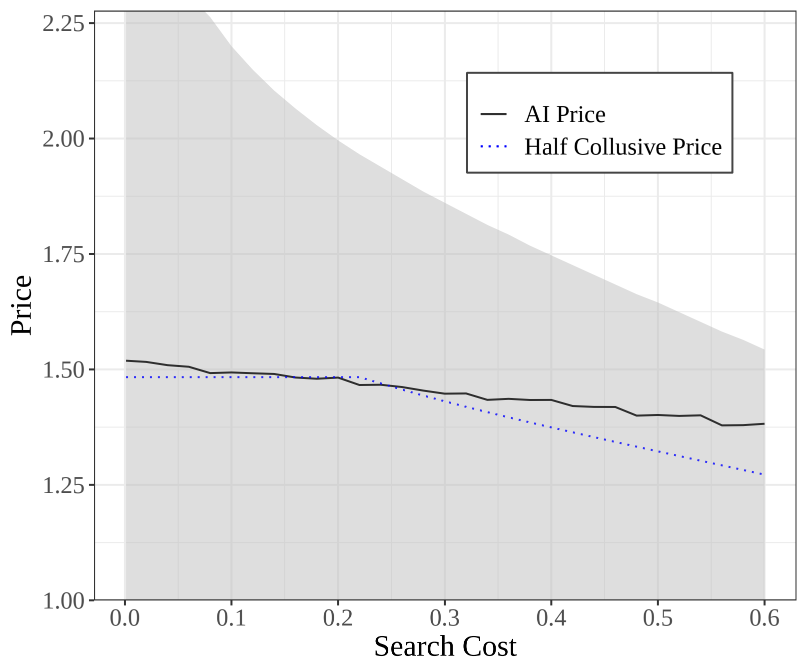

Price comparison

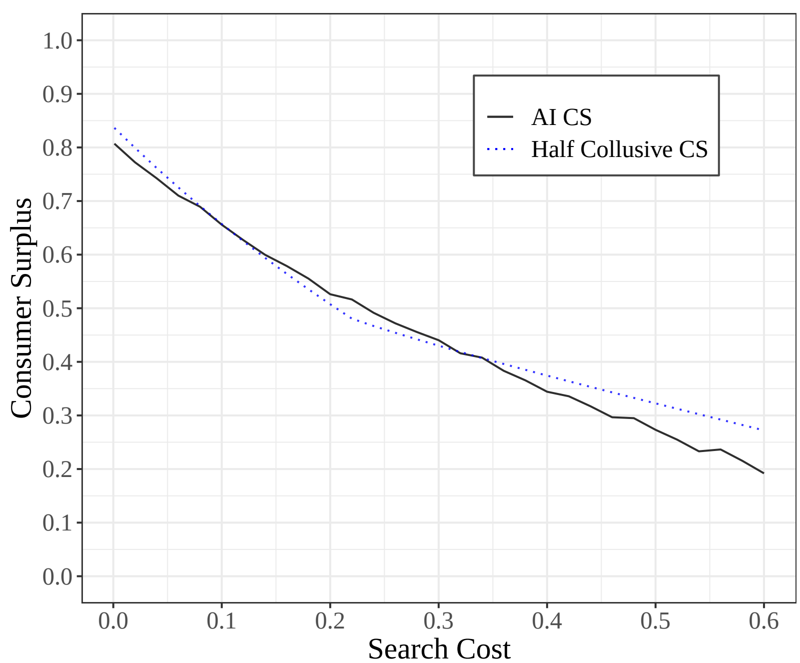

Consumer surplus comparison

Figure 3.5: Price and consumer surplus comparison with different search costs

Note. Simulations with marginal cost $c = 1$, discount factor $\delta = 0.95$, punishment length $L = \infty$, match value $v_i \sim \text{Gumbel}(2, 0.25)$, $n = 3$ and $k = 2$.

The AI price fits well in our theoretical equilibrium set and is close to the half-collusive price. Regarding welfare outcome, consumer surplus under half collusion matches AI simulations better than the randomized collusion, especially for $s \in [0, 0.3]$. We continue to compare our AI simulation results with half-collusion outcomes in different dimensions, including product quality, product differentiation, and production cost.

The simulation results confirm several theoretical predictions. First, higher product quality increases both price and consumer surplus. The extra utility from the better quality product is shared between sellers and consumers. It allows sellers to increase the price while still improving consumer surplus. Higher product differentiation increases both price and consumer surplus. When the products are more heterogeneous, the price competition will subside, leading to a higher price. On the other hand, with higher differentiated products, the search value for the consumer increases leading to higher consumer expected utility. Finally, higher production cost leads to higher price and lower consumer surplus. This result is intuitive.

In all cases, both price and consumer surplus under AI pricing simulations resemble the half-collusive equilibrium.

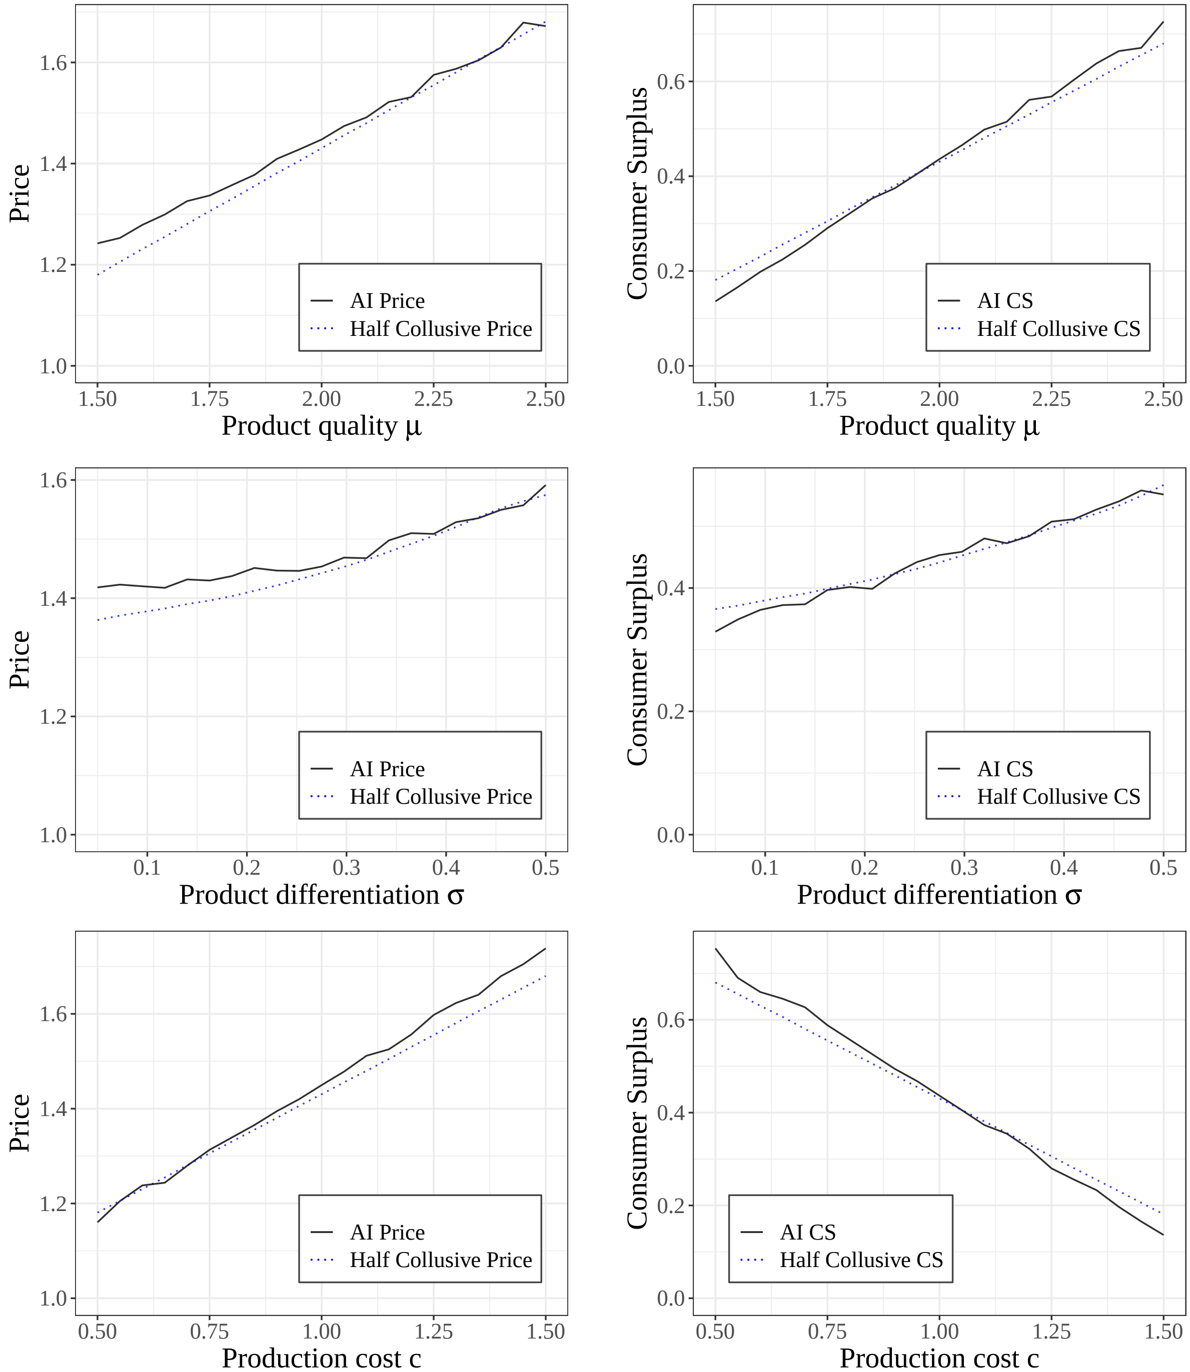

Figure 3.6: Price and consumer surplus comparison with different product quality, differentiation, and costs

3.9 Robustness

In this section, we will conduct robustness checks on the following hypotheses:

- Price decreases with higher search costs.
- Consumer surplus decreases with higher search costs.
- Price and consumer surplus with AI pricing algorithms resemble half-collusive equilibrium.

3.9.1 Product quality

We run our experiment with different product quality indexes μ for the robustness exercise. Taking the baseline specification as the middle point, we set $\mu = 1.5$ for the low product quality and $\mu = 2.5$ for the high product quality.

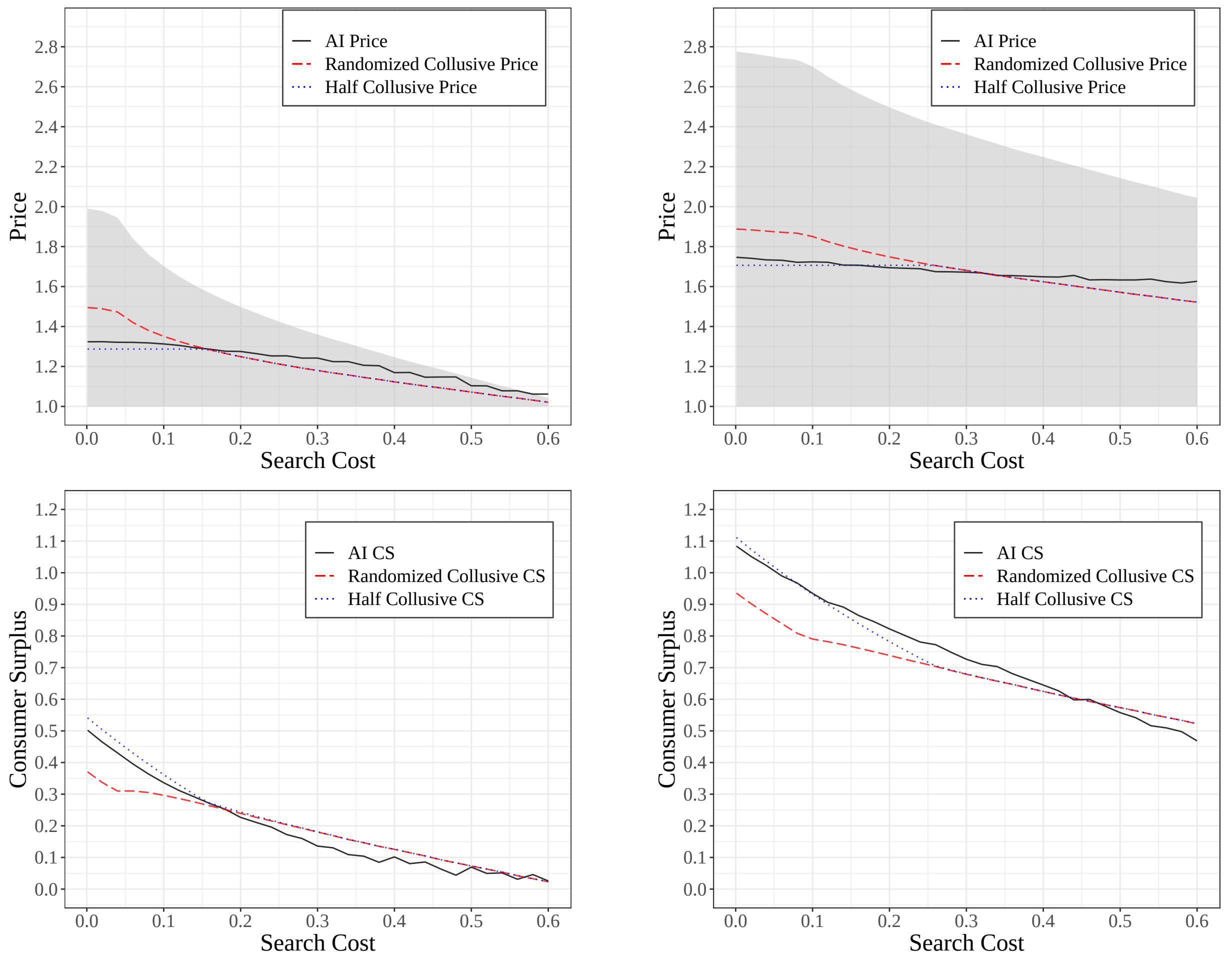

Low Product Quality $\mu = 1.5$ High Product Quality $\mu = 2.5$

Figure 3.7: Product quality robustness

With a high product quality, our experiment shows both AI price and consumer surplus increases. These results are intuitive and consistent with classical theory results, which predict that price and surpluses increase with higher product quality. The AI price still fits well in our theoretical equilibrium set under different product qualities. When the theoretical equilibrium set is narrow (under low product quality and high search cost), the AI price remains mainly within the theoretical equilibrium set. At $s = 6$, the AI price slightly exceeds the equilibrium upper boundary. However, this does not come as a surprise because of the discretization of the price action set. Under the low product quality specification, only one price (1.018) is higher than c and lower than $\overline{p}$. All other prices are below marginal cost or are too high that the consumer will never visit. If the agents fully colluded, they would all choose 1.018 as their price, and the price cycle has a length of one. However, AI agents generally cannot obtain such a price in practice. If yes, it is by chance. AI agents generally end up with a price cycle with the price of 1.018 as a part of it, meaning they can earn positive profit only sometimes in the cycle.

Regarding the comparative statics result, our experiment shows that, in general, AI prices and consumer surplus decrease with higher search costs. This result is consistent with our main theoretical results. However, we observe a zigzag shape on both price and consumer surplus under low product quality and high search cost, around 0.4-0.6. The uneven shape is pronounced with consumer surplus when the AI line abruptly increases several times. We believe this is due to the discretization of the price grid. The number of prices higher than c and lower than $\overline{p}_s$ change rapidly from 5 to 1 from $s = 0.4$ to $s = 0.6$. When a price becomes higher than $\overline{p}_s$ as the search cost increases, the profit from choosing that price becomes 0 every time because consumers will not visit the store at such a price. Each sudden drop in AI price and each jump in consumer surplus corresponds with each time $\overline{p}_s$ goes lower than one more price when search cost increases.

3.9.2 Product differentiation

In this robustness experiment, we run the simulations with different product differentiation $\sigma = 0.05$ and $\sigma = 0.5$ (in baseline specification, $\sigma = 0.25$). Figure 8 shows that AI price is lower with lower product differentiation. This result is intuitive because price competition is more intense when products are close substitutes. On the other hand, consumers are better of with higher product differentiation because of the higher search value. Consumers generally search more with high product differentiation (Figure 9).

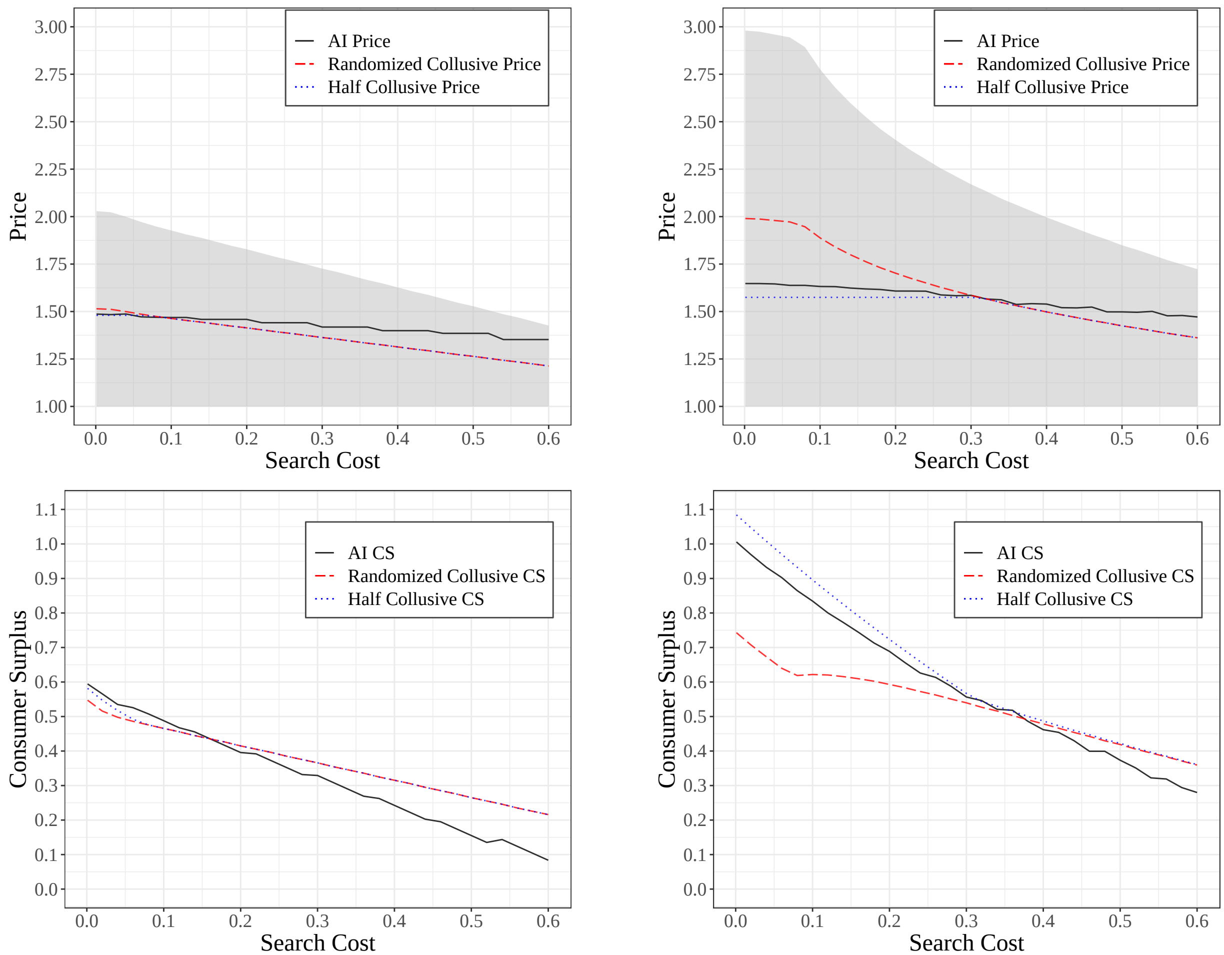

Low Product Differentiation $\sigma = 0.05$ High Product Differentiation $\sigma = 0.5$

Figure 3.8: Product differentiation robustness

Regarding the comparative statics, our experimental result is consistent with our theoretical results, which show that both price and consumer surplus decrease with a higher search cost. The AI price fits well in our theoretical equilibrium set. The price and consumer surplus from experiments are close to half-collusive equilibrium.

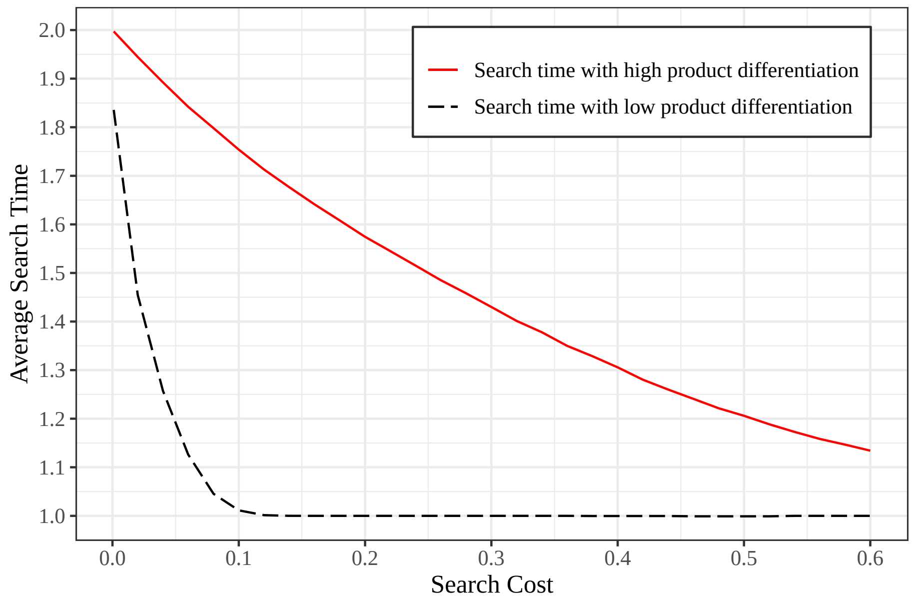

Figure 3.9: Consumer average search time

3.9.3 Production cost

In this robustness exercise, we conduct our experiment with different production costs. The low production cost takes the value of 0.5, and the high production cost takes the value of 1.5 (in baseline specification, $c = 1$).

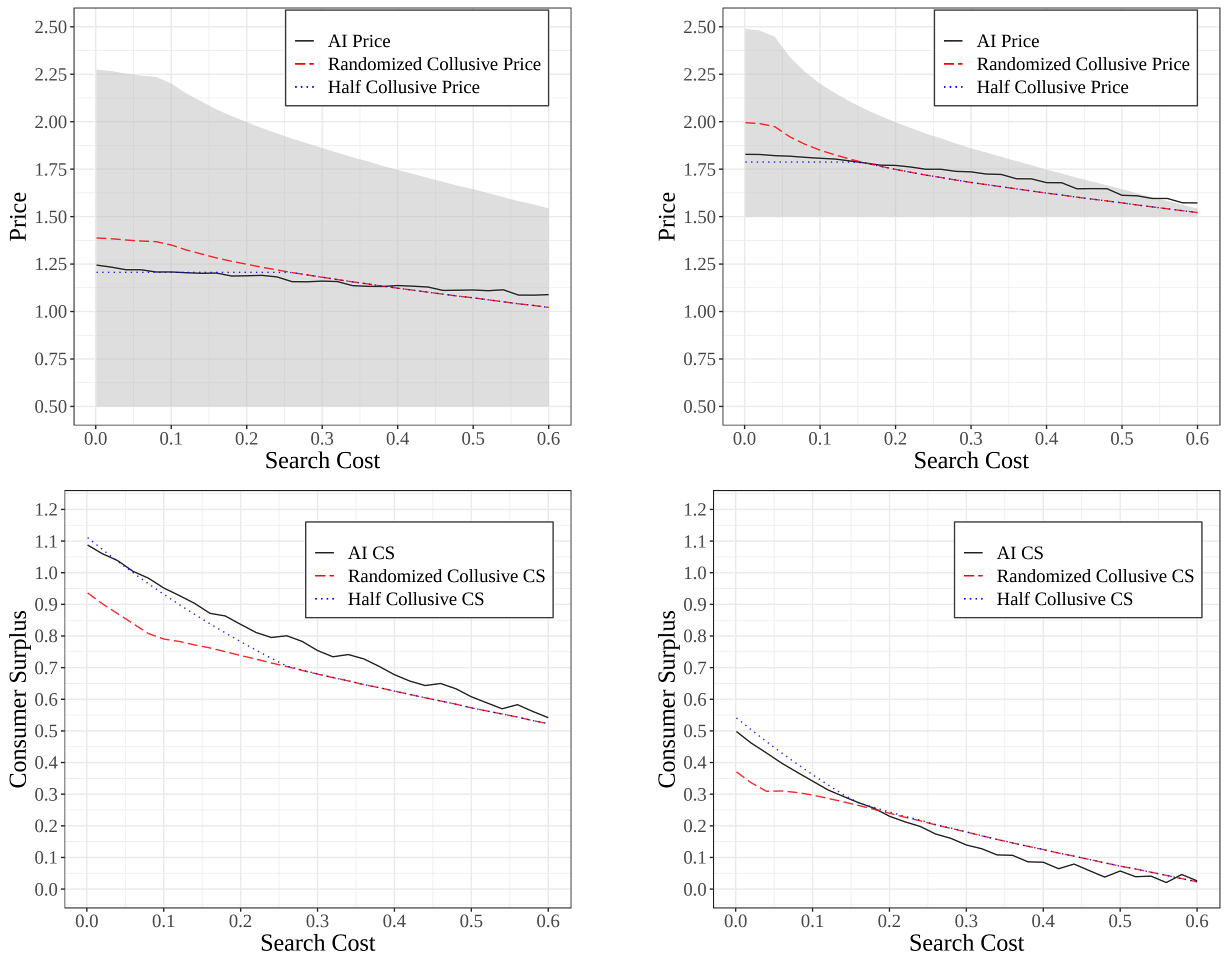

Low Production Cost $c = 0.5$ High Production Cost $c = 1.5$

Figure 3.10: Production cost robustness

Regarding the comparative statics, our experiments show that AI price and consumer surplus generally decrease with a higher search cost, which is consistent with our theoretical results. However, consumer surplus with AI agents has several slight jumps when search cost increases. We believe that this is due to the discretized price grid. Each dip of AI price and each jump of consumer surplus corresponds to each time a price in a price grid becomes lower than $\overline{p}_s$.

Our experiment shows that AI price increases and consumer surplus decreases with a higher production cost. The AI price generally fits well in our theoretical equilibrium set, except for high production cost and search cost ($s = \in [0.55, 0.6]$), the AI price is slightly above the equilibrium range. This outcome is similar to the low product quality case. The reason is the same in both cases and is discussed in subsection 9.2. The half-collusive equilibrium price and consumer surplus match the simulation outcome relatively well.

3.10 Conclusion

This article studies the roles of consumer search friction on collusion; and implications on market prices and consumer welfare. We tackle these questions theoretically by revisiting collusion model of Green and Porter (1984) with search friction and experimentally by simulating sellers' interactions with AI-powered pricing algorithms.

We find that lower search costs make it easier for firms to collude. This finding is distinctive because it differs from existing consumer search literature results and should probably ring an alarm to antitrust policies, especially for e-commerce marketplaces. The mechanism is that when prices are observable, lower search costs make undercutting the cartel price less profitable. Therefore, the sellers have less incentive to deviate, and the cartel becomes more sustainable. On the other hand, we also show that higher search costs can reduce market prices. However, this price reduction cannot offset the search burden on consumers if the collusion sustains.

With decreasing search costs due to new technological advancement, our findings raise an important question on how we should regulate consumer search friction. Lowering search costs can increase consumer surplus per se; however, the sellers can collude more easily. This change in collusion incentive puts the consumers at the stake of being worse off eventually.

Our paper also puts a step forward to characterize the collusion equilibrium that can realize in practice. This task is challenging if done only theoretically because of a large equilibrium set.

However, our experiment results show that markets with AI pricing algorithms have a similar outcome to half-collusive equilibrium, where sellers randomize between fully collusive pricing and competitive pricing. This result can help determine the optimal policy for markets where sellers use AI-powered pricing software.

Last but not least, we find that Q-learning, a simple reinforcement learning algorithm, can adopt a collusive strategy to maintain prices above the competitive level in the frictional market. The punishment scheme is similar to a price-trigger strategy in which a deviation is followed by punishment with lower prices from all agents. Prices gradually increase after the first hard punishment period. However, they converge to a lower price than before the deviation. This finding joins the policy concern on how to prosecute collusion when sellers use algorithmic pricing. These algorithms may autonomously learn to collude without having any overt communication or agreement.

Bibliography

Abreu, D., D. Pearce, and E. Stacchetti (1986). Optimal cartel equilibria with imperfect monitoring. *Journal of Economic Theory 39*(1), 251–269.

Abreu, D., D. Pearce, and E. Stacchetti (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica 58*(5), 1041–1063.

Anderson, E. T., K. Hansen, and D. Simester (2009). The option value of returns: Theory and empirical evidence. *Marketing Science 28*(3), 405–423.

Armstrong, M. (2017, 06). Ordered Consumer Search. *Journal of the European Economic Association 15*(5), 989–1024.

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

Arya, A. and B. Mittendorf (2004). Using return policies to elicit retailer information. *The RAND Journal of Economics 35*(3), 617–630.

Assad, S., R. Clark, D. Ershov, and L. Xu (2020). Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market. Technical report.

Brown, Z. and A. MacKay (2022). Competition in pricing algorithms. *American Economic Journal: Microeconomics (Forthcoming).*

Burdett, K. and K. L. Judd (1983). Equilibrium price dispersion. *Econometrica 51*(4), 955–969.

Calvano, E., G. Calzolari, V. Denicolò, and S. Pastorello (2020, October). 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.

Campbell, C., G. Ray, and W. A. Muhanna (2005). Search and collusion in electronic markets. *Management Science 51*(3), 497–507.

Che, Y.-K. (1996). Customer return policies for experience goods. *The Journal of Industrial Economics 44*(1), 17–24.

Chen, L., A. Mislove, and C. Wilson (2016). An empirical analysis of algorithmic pricing on amazon marketplace. In *Proceedings of the 25th International Conference on World Wide Web*, WWW '16, Republic and Canton of Geneva, CHE, pp. 1339–1349. International World Wide Web Conferences Steering Committee.

Choi, M., A. Y. Dai, and K. Kim (2018). Consumer search and price competition. *Econometrica 86*(4), 1257–1281.

Compte, O. (2002). On failing to cooperate when monitoring is private. *Journal of Economic Theory 102*(1), 151–188.

Courty, P. and H. Li (2000, 10). Sequential Screening. *The Review of Economic Studies 67*(4), 697–717.

Davis, S., E. Gerstner, and M. Hagerty (1995). Money back guarantees in retailing: matching products to consumer tastes. *Journal of Retailing 71*(1), 7 – 22.

Davis, S., M. Hagerty, and E. Gerstner (1998). Return policies and the optimal level of "hassle". *Journal of Economics and Business 50*(5), 445–460.

De Loecker, J., J. Eeckhout, and G. Unger (2020, 01). The Rise of Market Power and the Macroeconomic Implications*. *The Quarterly Journal of Economics 135*(2), 561–644.

Delrahim, M. (2018). "life in the fast lane": Antitrust in a changing telecommunications landscape. *U.S. Deparment of Justice*.

Escobari, D. and P. Jindapon (2014). Price discrimination through refund contracts in airlines. *International Journal of Industrial Organization 34*, 1 – 8.

Ezrachi, A. and M. E. Stucke (2017). Artificial intelligence and collusion: When computers inhibit competition. *University of Illinois Law Review*, 1775–1810.

Fudenberg, D., D. Levine, and E. Maskin (1994). The folk theorem with imperfect public information. *Econometrica 62*(5), 997–1039.

Green, E. J. and R. H. Porter (1984). Noncooperative collusion under imperfect price information. *Econometrica 52*(1), 87–100.

Grossman, S. (1981). The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics 24*(3), 461–83.

Haan, M., J. Moraga-Gonzalez, and V. Petrikaite (2018). A model of directed consumer search. *International Journal of Industrial Organization 61*(C), 223–255.

Harrington, J. E. (2018). Developing Competition Law For Collusion By Autonomous Artificial Agents. *Journal of Competition Law and Economics 14*(3), 331–363.

Hess, J. D., W. Chu, and E. Gerstner (1996). Controlling product returns in direct marketing. *Marketing Letters 7*(4), 307–317.

Hinnosaar, T. and K. Kawai (2020). Robust pricing with refunds. *The RAND Journal of Economics 51*(4), 1014–1036.

Hotelling, H. (1929). Stability in competition. *The Economic Journal 39*(153), 41–57.

Inderst, R. and M. Ottaviani (2013). Sales talk, cancellation terms and the role of consumer protection. *The Review of Economic Studies 80*(3 (284)), 1002–1026.

Inderst, R. and G. Tirosh (2015). Refunds and returns in a vertically differentiated industry. *International Journal of Industrial Organization 38*, 44 – 51.

Johnson, J., A. Rhodes, and M. Wildenbeest (2022, October). Platform design when sellers use pricing algorithms. *SSRN Electronic Journal*.

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

Krähmer, D. and R. Strausz (2015, 01). Optimal Sales Contracts with Withdrawal Rights. *The Review of Economic Studies 82*(2), 762–790.

Lutz, N. A. and V. Padmanabhan (1998). Warranties, extended warranties, and product quality. *International Journal of Industrial Organization 16*(4), 463–493.

Maskin, E. and J. Tirole (1988). A theory of dynamic oligopoly, ii: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica 56*(3), 571–599.

Mason, R. and J. Välimäki (2011). Learning about the arrival of sales. *Journal of Economic Theory 146*(4), 1699–1711.

Matsushima, H. and M. Kandori (1998, 05). Private observation, communication and collusion. *Econometrica 66*, 627–652.

Mehra, S. (2016). Antitrust and the robo-seller: Competition in the time of algorithms. *Minnesota Law Review 100*, 1323–1375.

Moorthy, S. and K. Srinivasan (1995). Signaling quality with a money-back guarantee: The role of transaction costs. *Marketing Science 14*(4), 442–466.

Nilsson, A. (1999, February). Transparency and Competition. SSE/EFI Working Paper Series in Economics and Finance 298, Stockholm School of Economics.

Obradovits, M. and P. Plaickner (2022). Price-Directed Search, Product Differentiation and Competition. Technical report.

OECD (2017). Algorithms and collusion: Competition policy in the digital age. *Organisation for Economic Co-operation and Development*.

Palfrey, T. and T. Romer (1983). Warranties, performance, and the resolution of buyer-seller disputes. *The Bell Journal of Economics 14*(1), 97–117.

Petrikaitė, V. (2016). Collusion with costly consumer search. *International Journal of Industrial Organization 44*, 1–10.

Petrikaitė, V. (2018). A search model of costly product returns. *International Journal of Industrial Organization 58*, 236 – 251.

Posselt, T., E. Gerstner, and D. Radic (2008). Rating e-tailers' money-back guarantees. *Journal of Service Research 10*(3), 207–219.

Schultz, C. (2005). Transparency on the consumer side and tacit collusion. *European Economic Review 49*(2), 279–297.

Shulman, J. D., A. T. Coughlan, and R. C. Savaskan (2011). Managing consumer returns in a competitive environment. *Management Science 57*(2), 347–362.

Stahl, D. (1989). Oligopolistic pricing with sequential consumer search. *American Economic Review 79*(4), 700–712.

Sutton, R. S. and A. G. Barto (2018). *Reinforcement Learning: An Introduction* (Second ed.). The MIT Press.

Waltman, L. and U. Kaymak (2008). Q-learning agents in a cournot oligopoly model. *Journal of Economic Dynamics and Control 32*(10), 3275–3293.

Watkins, C. (1989, 01). Learning from delayed rewards. *PhD thesis. King's College*.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica 47*(3), 641–654.

Wolinsky, A. (1986). True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics 101*(3), 493–511.

Appendix A

Supplementary material for Chapter 1

Proof of Proposition 1. The explicit function of l^* comes from the fact that l^* maximizes A_{l^*}. Let $\hat{p}^*$ be the equilibrium expected expenditure. The following argument is to show $\hat{p}^* < p' < p^*$ if $s\lambda(1-\alpha) > m$. Recall that

$$\hat{p}^* = \frac{1 - F\left(\frac{\hat{p}^*}{Q_0}\right)}{f\left(\frac{\hat{p}^*}{Q_0}\right)} Q_0 + c - A_{l^*}, \tag{A.1}$$

$$p' = \frac{1 - F\left(\frac{p'}{Q_0}\right)}{f\left(\frac{p'}{Q_0}\right)} Q_0 + c. \tag{A.2}$$

Given $s\lambda(1-\alpha) > m$ and $\frac{1-F}{f}$ is monotone non-increasing, then $A_{l^*} > 0$, $\hat{p}^* < p'$, and

$$\frac{1 - F\left(\frac{\hat{p}^*}{Q_0}\right)}{f\left(\frac{\hat{p}^*}{Q_0}\right)} \geq \frac{1 - F\left(\frac{p'}{Q_0}\right)}{f\left(\frac{p'}{Q_0}\right)}. \tag{A.3}$$

By (A.1), the definitions of $\hat{p}$ and A_l, we have

$$p^* - (1-\alpha)\int_0^{l^*}(p^* - s\mathrm{e}^{-\sigma t})\lambda\mathrm{e}^{-\lambda t}\mathrm{d}t - ml^* = \frac{1-F\left(\frac{\hat{p}^*}{Q_0}\right)}{f\left(\frac{\hat{p}^*}{Q_0}\right)}Q_0 + c. \tag{A.4}$$

By contradiction, suppose $p^* \leq p'$. By (A.2) and (A.4), we have

$$\frac{1-F\left(\frac{\hat{p}^*}{Q_0}\right)}{f\left(\frac{\hat{p}^*}{Q_0}\right)}Q_0 + c < p^* \leq p' = \frac{1-F\left(\frac{p'}{Q_0}\right)}{f\left(\frac{p'}{Q_0}\right)}Q_0 + c,$$

which violates (A.3). We conclude by contradiction that $p^* > p'$. □

Proof of Corollary 1. From the explicit function of l^* (Proposition 1), the partial derivative $\partial l^*/\partial\lambda$ is given by

$$\frac{\lambda+\sigma-\lambda\ln\left[\frac{(1-\alpha)s\lambda}{m}\right]}{\lambda(\lambda+\sigma)^2}.$$

By contradiction, suppose there exist $\lambda_1 < \lambda_2$ such that $\partial l^*/\partial\lambda\big|_{\lambda=\lambda_i} = 0$ for all $i \in \{1,2\}$. Therefore,

$$\lambda_i + \sigma = \lambda_i \ln\left[\frac{s\lambda_i(1-\alpha)}{m}\right] \ \forall i \in \{1,2\},$$

which implies $\ln[s\lambda_i(1-\alpha)/m] > 1$. Because $\ln[s\lambda_i(1-\alpha)/m]$ is increasing in λ_i and $\lambda_1 + \sigma = \lambda_1\ln[s\lambda_1(1-\alpha)/m]$, we must have $\lambda_2+\sigma < \lambda_2\ln[s\lambda_2(1-\alpha)/m]$, which is contradictory.

□

Proof: *$\hat{p}$ increases in α and decreases in λ.* Because $f/(1-F)$ is monotone non-decreasing and Q_0 increases in α and decreases in λ, by (1.7), we only need to show that A_{l^*} decreases in α and increases in λ. It is clear that A_{l^*} decreases in α. The following is to show that A_{l^*} increases in λ.

- Suppose $m/[(1-\alpha)s] < \lambda < \dot{\lambda}$. By Corollary 1, $\partial l^*/\partial\lambda > 0$. By (2.12), we have

$\partial A_{l^*}/\partial l^* = s\lambda e^{-(\lambda+\sigma)l^*} > 0$. Thus $\partial A_{l^*}/\partial \lambda > 0$.

- Suppose $\lambda > \dot{\lambda}$. By (2.12), we can represent A_{l^*} as following

$$A_{l^*} = s(1-\alpha)\frac{\lambda}{\lambda+\sigma}\left(1 - \int_{l^*}^{\infty}(\lambda+\sigma)s e^{-(\lambda+\sigma)t}\mathrm{d}t\right) - ml^*. \tag{A.5}$$

Let l_1^* and l_2^* be the two equilibrium return periods corresponding to λ_1 and λ_2 such that $\dot{\lambda} \leq \lambda_1 < \lambda_2$. Thus $l_2^* < l_1^*$ by Corollary 1. By the explicit function of l^* in Proposition 1, we have $s(1-\alpha)\lambda e^{-(\lambda+\sigma)l^*} = m$. Thus

$$s(1-\alpha)\int_{l_2^*}^{l_1^*}\lambda_2 e^{-(\lambda_2+\sigma)t}\mathrm{d}t < m\left(l_1^* - l_2^*\right) \tag{A.6}$$

and $\lambda_2 e^{-(\lambda_2+\sigma)l_1^*} < \lambda_1 e^{-(\lambda_1+\sigma)l_1^*}$. Note that $e^{-(\lambda_2+\sigma)t}/e^{-(\lambda_1+\sigma)t}$ decreases in t. Therefore,

$$s(1-\alpha)\int_{l_1^*}^{\infty}\lambda_2 e^{-(\lambda_2+\sigma)t}\mathrm{d}t < s(1-\alpha)\int_{l_1^*}^{\infty}\lambda_1 e^{-(\lambda_1+\sigma)t}\mathrm{d}t. \tag{A.7}$$

Thus $A(l_1^*) < A(l_2^*)$ is a direct result from (A.5), (A.6), and (A.7).

□

Proof of Proposition 2. Let p_1^* and p_2^* be the two corresponding equilibrium prices for α_1 and α_2 such that $0 < \alpha_1 < \alpha_2 < \tilde{\alpha}$. Consider the case where s/m is high enough such that $s\lambda(1-\tilde{\alpha}) > m$ and

$$\frac{\lambda}{\gamma} < \left[(1-\tilde{\alpha})\lambda\frac{s}{m}\right]^{\frac{\lambda}{\sigma+\lambda}} - 1. \tag{A.8}$$

We will prove that $p_1^* > p_2^*$.

From the explicit function of l^* in Proposition 1, the inequality (A.8) implies

$$\frac{\lambda}{\gamma} < \mathrm{e}^{\lambda l^*(\alpha)} - 1 = \frac{1 - \mathrm{e}^{-\lambda l^*(\alpha)}}{\mathrm{e}^{-\lambda l^*(\alpha)}} \quad \forall \alpha \in [\alpha_1, \alpha_2].$$

Because $l^*(\alpha)$ decreases in α, we have

$$\frac{\partial}{\partial \alpha}\left[\frac{\alpha\left(1 - \mathrm{e}^{-\lambda l^*(\alpha)}\right) + \mathrm{e}^{-\lambda l^*(\alpha)}}{\alpha\lambda + \gamma}\right] > 0 \quad \forall \alpha \in [\alpha_1, \alpha_2]. \tag{A.9}$$

Note that $\Omega_l = 1 - (1-\alpha)\int_0^{l^*(\alpha)} \lambda \mathrm{e}^{-\lambda t}\mathrm{d}t = \alpha\left(1 - \mathrm{e}^{-\lambda l^*(\alpha)}\right) + \mathrm{e}^{-\lambda l^*(\alpha)}$, the inequality (A.9) implies

$$\frac{1 - (1-\alpha_1)\int_0^{l_1^*} \lambda \mathrm{e}^{-\lambda t}\mathrm{d}t}{Q_1} < \frac{1 - (1-\alpha_2)\int_0^{l_2^*} \lambda \mathrm{e}^{-\lambda t}\mathrm{d}t}{Q_2}, \tag{A.10}$$

where $Q_1 = Q_0\big|_{\alpha=\alpha_1}$, $Q_2 = Q_0\big|_{\alpha=\alpha_2}$, $l_1^* = l^*(\alpha)\big|_{\alpha=\alpha_1}$, and $l_2^* = l^*(\alpha)\big|_{\alpha=\alpha_2}$.

By contradiction, suppose $p_1^* \leq p_2^*$. Because $\hat{p} = \Omega_l p$, we have the following inequalities

$$\frac{\hat{p}_1^*}{Q_1} < \frac{\hat{p}_2^*}{Q_2}, \tag{A.11}$$

$$\frac{(p_1^* - s)\left[1 - (1-\alpha_1)\int_0^{l_1^*} \lambda \mathrm{e}^{-\lambda t}\mathrm{d}t\right]}{Q_1} < \frac{(p_2^* - s)\left[1 - (1-\alpha_2)\int_0^{l_2^*} \lambda \mathrm{e}^{-\lambda t}\mathrm{d}t\right]}{Q_2}. \tag{A.12}$$

Note that $l_1^* > l_2^*$ and $Q_1 < Q_2$, we also have the following inequalities

$$\frac{s - c - ml_1^*}{Q_1} < \frac{s - c - ml_2^*}{Q_2} \tag{A.13}$$

and

$$\frac{-s(1-\alpha_1)\int_0^{l_1^*}(1-\mathrm{e}^{-\sigma t})\,\lambda\mathrm{e}^{-\lambda t}\mathrm{d}t}{Q_1} < \frac{-s(1-\alpha_2)\int_0^{l_2^*}(1-\mathrm{e}^{-\sigma t})\,\lambda\mathrm{e}^{-\lambda t}\mathrm{d}t}{Q_2}. \tag{A.14}$$

By adding (A.12), (A.13), and (A.14) all together, we get

$$\frac{\hat{p}_1^* + A(l_1^*) - c}{Q_1} < \frac{\hat{p}_2^* + A(l_2^*) - c}{Q_2}. \tag{A.15}$$

Given $(1-F)/f$ is monotone non-increasing, then (A.11) and (A.15) lead to a violation to (1.7). We conclude from the contradiction that $p_1^* > p_2^*$. □

Proof of Proposition 3. By the explicit function of l^* in Proposition 1, we note that $\partial\left[(\lambda+\sigma)\mathrm{e}^{-\lambda l^*}\right]/\partial\lambda < 0$ for all $\lambda > m/[s(1-\alpha)]$. Because $\gamma \geq \sigma$, we have

$$\frac{\partial\left[(\lambda+\gamma)\mathrm{e}^{-\lambda l^*}\right]}{\partial\lambda} < 0 \quad \text{for all } \lambda > \frac{m}{s(1-\alpha)}. \tag{A.16}$$

Because $l^*\big|_{\lambda=m/[s(1-\alpha)]} = 0$, there exists a $\ddot{\lambda} > m/[s(1-\alpha)]$ such that

$$(\lambda+\sigma)\mathrm{e}^{-\lambda l^*} > \sigma \quad \text{for all } \lambda \in \left[\frac{m}{s(1-\alpha)}, \ddot{\lambda}\right]. \tag{A.17}$$

By Corollary 1, there exists a $\dot{\lambda}$ such that l^* increases in λ if $\lambda \in [m/[s(1-\alpha)], \dot{\lambda}]$. We will prove the Proposition 3 by choosing $\tilde{\lambda}$ equal to $\min\{\dot{\lambda}, \ddot{\lambda}\}$. Let p_1^* and p_2^* be the two corresponding equilibrium prices for λ_1 and λ_2 such that $m/[s(1-\alpha)] < \lambda_1 < \lambda_2 < \tilde{\lambda}$. We will show that $p_1^* < p_2^*$.

By (A.16) and (A.17), we have

$$(\lambda_1+\sigma)\mathrm{e}^{-\lambda_1 l_1^*} > (\lambda_2+\sigma)\mathrm{e}^{-\lambda_2 l_2^*} > \sigma \tag{A.18}$$

where $l_1^* = l^*\big|_{\lambda=\lambda_1}$ and $l_2^* = l^*\big|_{\lambda=\lambda_2}$. The inequality chain (A.18) yields

$$\frac{(1-\alpha)\mathrm{e}^{-\lambda_1 l_1^*}+\alpha}{\frac{1-\alpha}{\lambda_1+\gamma}+\frac{\alpha}{\gamma}} > \frac{(1-\alpha)\mathrm{e}^{-\lambda_2 l_2^*}+\alpha}{\frac{1-\alpha}{\lambda_2+\gamma}+\frac{\alpha}{\gamma}}. \tag{A.19}$$

Note that $\Omega_l = 1-(1-\alpha)\int_0^{l^*(\lambda)} \lambda \mathrm{e}^{-\lambda t}\mathrm{d}t = (1-\alpha)\mathrm{e}^{-\lambda l^*(\lambda)}+\alpha$, the inequality (A.19) is equivalent to

$$\frac{1-(1-\alpha)\int_0^{l_1^*} \lambda_1 \mathrm{e}^{-\lambda_1 t}\mathrm{d}t}{Q_1} > \frac{1-(1-\alpha)\int_0^{l_2^*} \lambda_2 \mathrm{e}^{-\lambda_2 t}\mathrm{d}t}{Q_2}, \tag{A.20}$$

where $Q_1 = Q_0\big|_{\lambda=\lambda_1}$ and $Q_2 = Q_0\big|_{\lambda=\lambda_2}$.

By contradiction, suppose $p_1^* \geq p_2^*$. Because $\hat{p} = \Omega_l p$, we have the following inequalities

$$\frac{\hat{p}_1^*}{Q_1} > \frac{\hat{p}_2^*}{Q_2}, \tag{A.21}$$

$$\frac{(p_1^*-s)\left[1-(1-\alpha)\int_0^{l_1^*} \lambda_1 \mathrm{e}^{-\lambda_1 t}\mathrm{d}t\right]}{Q_1} > \frac{(p_2^*-s)\left[1-(1-\alpha)\int_0^{l_2^*} \lambda_2 \mathrm{e}^{-\lambda_2 t}\mathrm{d}t\right]}{Q_2}. \tag{A.22}$$

Because $l_1^* < l_2^*$ and $Q_1 > Q_2$, we also have the following inequality

$$\frac{s-c-ml_1^*}{Q_1} > \frac{s-c-ml_2^*}{Q_2}. \tag{A.23}$$

Finally, we will show that

$$\frac{-s(1-\alpha)\int_0^{l_1^*}\left(1-\mathrm{e}^{-rt}\right)\lambda_1 \mathrm{e}^{-\lambda_1 t}\mathrm{d}t}{Q_1} > \frac{-s(1-\alpha)\int_0^{l_2^*}\left(1-\mathrm{e}^{-rt}\right)\lambda_2 \mathrm{e}^{-\lambda_2 t}\mathrm{d}t}{Q_2}. \tag{A.24}$$

Note that $\partial l^*/\partial\lambda\big|_{\lambda=\dot{\lambda}} = 0$ implies that $\dot{\lambda} = 1/l^*\big|_{\lambda=\dot{\lambda}}$. Because l^* increases in λ for $\lambda \in [m/[s(1-\alpha)], \dot{\lambda}]$, we have $\lambda < 1/l^*$ for $\lambda \in [m/[s(1-\alpha)], \dot{\lambda}]$. Therefore, $\lambda_1 < \lambda_2 < 1/l_1^* < 1/l_2^*$.

Thus $\lambda_1 < \lambda_2 < 1/t$ for all $t \in (0, l_1^*)$. Hence $\lambda_1 \mathrm{e}^{-\lambda_1 t} < \lambda_2 \mathrm{e}^{-\lambda_2 t}$ for all $t \in (0, l_1^*)$. The inequality (A.24) is then a direct result.

By adding (A.22), (A.23), and (A.24) all together, we get

$$\frac{\hat{p}_1^* + A(l_1^*) - c}{Q_1} > \frac{\hat{p}_2^* + A(l_2^*) - c}{Q_2}. \tag{A.25}$$

Given $(1-F)/f$ is monotone non-increasing, then (A.21) and (A.25) lead to a violation to (1.7). We conclude from the contradiction that $p_1^* < p_2^*$. □

Proof of Lemma 1. By (1.8), $\mathrm{d}t/\mathrm{d}\pi_t = 1/[\pi_t(1-\pi_t)\lambda]$. Let $\pi_t\big|_{t=T} = \beta$. Because $\pi_t\big|_{t=0} = \alpha$, we have $\int_\alpha^\beta 1/[\pi_t(1-\pi_t)\lambda]\mathrm{d}\pi_t = T$. Thus

$$\begin{aligned} T\lambda &= \int_\alpha^\beta \frac{1}{\pi_t}\mathrm{d}\pi_t + \int_\alpha^\beta \frac{1}{1-\pi_t}\mathrm{d}\pi_t \\ &= \ln\left[\frac{\beta(1-\alpha)}{(1-\beta)\alpha}\right]. \end{aligned}$$

Therefore, we have $\pi_t\big|_{t=T} = \beta = \alpha/[(1-\alpha)\mathrm{e}^{-\lambda T} + \alpha]$. □

Proof of Proposition 4. From the definitions of Ω_l, equation (1.3), and Q_l, equation (1.9), and the explicit function of π_l, equation (**??**), we have

$$\Omega_l Q_l = \frac{\gamma + \frac{\lambda\alpha}{(1-\alpha)\mathrm{e}^{-\lambda l}+\alpha}}{(\gamma+\lambda)\gamma}\mathrm{e}^{-\gamma l}\left[\alpha + (1-\alpha)\mathrm{e}^{-\lambda l}\right].$$

Thus $\partial(\Omega_l Q_l)/\partial l = -\left[\alpha\left(\mathrm{e}^{\lambda l} - 1\right) + 1\right]\mathrm{e}^{-(\gamma+\lambda)l}$, which is strictly negative, so by equation (1.10), the critical expected expenditure decreases in l. □

Proof of Proposition 5. Let $\hat{p}_l^*$ be the expected expenditure that maximizes the seller's profit for a given l. Thus

$$\hat{p}_l^* = \frac{1 - F\left(\frac{\hat{p}_l^*}{\Omega_l Q_l}\right)}{f\left(\frac{\hat{p}_l^*}{\Omega_l Q_l}\right)} \Omega_l Q_l - A_l + S_l^n. \tag{A.26}$$

Note that $A_{l=0} = 0$, $S_{l=0}^n = s$, and $\Omega_{l=0} = 1$. Given any $l \geq 0$, we have $\lim_{\mu \to \infty} \hat{p}_l^* = \infty$ and $\lim_{\mu \to \infty} \Pi(\hat{p}_l^*) = \hat{p}_l^* + A_l - c$. For any $\tau > 0$:

- Suppose $A_{l=\tau} \leq 0$. It is clear that $\Pi(\hat{p}|l = \tau) < \Pi(\hat{p}|l = 0)$ for all $\hat{p}$.
- Suppose $A_{l=\tau} > 0$. Because $\Omega_{l=\tau} Q_{l=\tau} < Q_0$ and $S_{l=\tau}^n < s$, there exists a finite $\tilde{\mu}$ such that if $\mu > \tilde{\mu}$, then $\hat{p}_{l=\tau}^* < \hat{p}_{l=0}^* - A_{l=\tau}$. This implies, there exists a finite $\tilde{\mu}$ such that if $\mu > \tilde{\mu}$, then $\Pi(\hat{p}_{l=\tau}^*) < \Pi(\hat{p}_{l=0}^*)$.

Therefore, there exist a finite $\tilde{\mu}$, such that if $\mu > \tilde{\mu}$ then $\Pi(\hat{p}, l)$ is maximized only if $l = 0$. Let $\hat{p}_B^*$ be the equilibrium expected expenditure in the baseline model. By equation (1.7), we have

$$\hat{p}_B^* = \frac{1 - F\left(\frac{\hat{p}_B^*}{Q_0}\right)}{f\left(\frac{\hat{p}_B^*}{Q_0}\right)} Q_0 - \max_{l \geq 0} A_l + c. \tag{A.27}$$

Because $\max_{l \geq 0} A_l > 0$, there exists a $\tilde{s} < c$ such that $c - \max_{l \geq 0} A_l < s$ for $s > \tilde{s}$. Note that $(1 - F)/f$ is monotone non-increasing. Thus $\hat{p}_B^* < \hat{p}_{l=0}^*$ if $s > \tilde{s}$ and $\mu > \tilde{\mu}$. $\square$

Appendix B

Supplementary material for Chapter 2

Proof of Proposition 6.

Part 1: It is clear that v_k is strictly decreasing in h_e and v_b is non-decreasing in h_e. I claim that if $h_e = p - h$, then $v_b > 0$. Indeed, suppose by contradiction $v_b \leq 0$ then $\max\left(U(v_b\mathrm{e}^{-\delta l}, \pi_l), p - h - h_e\right) = 0$. Because $B(v_b) = p$, then

$$\begin{aligned} p - (p-h)(1-\alpha)\int_0^l \lambda\mathrm{e}^{-\lambda t}dt =& (1-\alpha)\int_0^l\int_0^y v_b\mathrm{e}^{-\delta x}\mathrm{d}x\lambda\mathrm{e}^{-\lambda y}\mathrm{d}y \\ &+ \left[(1-\alpha)\mathrm{e}^{-\lambda l} + \alpha\right]\left[\int_0^l v_b\mathrm{e}^{-\delta x}\mathrm{d}x\right]. \end{aligned} \tag{B.1}$$

This is contradictory as the left hand side of (B.1) is positive whereas the right hand side is negative. Hence, $v_b|_{h_e=p-h} > 0$. Note that $v_k|_{h_e=p-h} = 0$. Therefore, if $v_b|_{h_e=a} < v_k|_{h_e=a}$, then there exists a unique $b > a$ such that $v_b|_{h_e=b} = v_k|_{h_e=b}$.

Part 2: If $h_e \leq b$, then $v_b \leq v_k$. The seller's profit is

$$\Pi = \left[1 - F(v_k)\right]\Pi_1 + \left[F(v_k) - F(v_b)\right]\Pi_2$$

where

$$\Pi_1 \equiv p - (1-\alpha)\int_0^l (p - s\mathrm{e}^{-\delta t})\lambda \mathrm{e}^{-\lambda t}\mathrm{d}t - ml - c$$

$$\Pi_2 \equiv p - (1-\alpha)\int_0^l (p - s\mathrm{e}^{-\delta t})\lambda \mathrm{e}^{-\lambda t}\mathrm{d}t - \left[\alpha + (1-\alpha)\mathrm{e}^{-\lambda l}\right]\left(p - s\mathrm{e}^{-\delta l}\right) - ml - c.$$

Because $p > c$ (otherwise, $\Pi < 0$) and $c > s$ (by assumption), then $p > s$. Hence, $\Pi_1 > \Pi_2$ and $\Pi_2 < 0$. Because v_k is decreasing in h_e and v_b is increasing in h_e, then $F(v_k)$ is decreasing in h_e and $F(v_b)$ is increasing in h_e. Therefore, Π increases in h_e.

Part 3: If $h_e \geq b$, then $v_b \geq v_k$. The seller's profit is $\Pi = [1 - F(v_k)]\,\Pi_1$. Because $v_b \geq v_k$ and $B(v_b) = p$, then

$$p = U(v_b, \alpha) + (p - h)(1-\alpha)\int_0^l \lambda \mathrm{e}^{-\lambda t}\mathrm{d}t,$$

which shows that v_b is constant in h_e. Therefore, Π is constant in h_e.

By results of the three parts, we conclude that $\Pi|_{h_e \geq b} = \max_{h_e \geq 0} \Pi(h_e)$. □

Proof of Corollary 2. By the definition of $A(l)$, $\frac{\partial A(l)}{\partial l} = (1-\alpha)(s\mathrm{e}^{-\delta l} - h)\lambda \mathrm{e}^{-\lambda l} - m$. Because l^* maximizes $A(l)$, then $l^* > 0$ if and only if $\left.\frac{\partial A(l)}{\partial l}\right|_{l=0} > 0$. Therefore, $l^* > 0$ if and only if $(1-\alpha)(s-h)\lambda - m > 0$. □

Appendix C

Supplementary material for Chapter 3

Proof. (**Proposition 13**) For $k = n$, the proof can be found in Weitzman (1979). For $k < n$, suppose the policy is optimal if each consumer can only visit m more sellers ($m \leq k$). To use the induction in the Weitzman proof, we only need to show that the policy is optimal when $m = 1$.

If $m = 1$, the optimal option is to choose the seller with the highest reservation value (the lowest price) or the outside option. From the definition of the reservation value, it is optimal to choose the outside option if and only if u_0 exceeds the highest reservation value. Therefore, the policy is optimal when $m = 1$. $\square$

Proof. (**Lemma 3**) Convert to standard Type I Extreme Value:
Let $\epsilon_i \equiv \dfrac{v_i - \mu}{\sigma}$. Therefore, the c.d.f. and p.d.f. of ϵ_i have the standard form: $G(\epsilon_i) = \exp(-\exp(-\epsilon_i))$, and $g(\epsilon_i) = \exp(-\epsilon_i)\exp(-\exp(-\epsilon_i))$. Consider seller i against the outside option. The consumer purchases i if $v_i - p_i > u_0$, which is equivalent to

$$\epsilon_i > \frac{u_0 - \mu}{\sigma} + \frac{p_i}{\sigma}.$$

Consider seller i against seller j. The consumer purchases i if $v_i - p_i > v_j - p_j$, which is equivalent to

$$\epsilon_j < \frac{p_j}{\sigma} - \frac{p_i}{\sigma} + \epsilon_i.$$

For simplicity, let $\rho \equiv \frac{p}{\sigma}$ and $b \equiv \frac{u_0 - \mu}{\sigma}$. The demand for seller i is

$$D_i = \int_{\rho_i + b}^{\infty} \prod_{j \neq i}^{k} G(\rho_j - \rho_i + \epsilon_i) g(\epsilon_i) d\epsilon_i. \tag{C.1}$$

Substitute $G(\epsilon_j)$ and $g(\epsilon_i)$ we can simplify the inner product to

$$\begin{aligned} D_i &= \int_{\rho_i + b}^{\infty} \prod_{j=1}^{k} \exp\left(-\exp\left(-\left(\rho_j - \rho_i + \epsilon_i\right)\right)\right) \exp(-\epsilon_i) d\epsilon_i \\ &= \int_{\rho_i + b}^{\infty} \exp\left(-\sum_{j=1}^{k} \exp(-(\rho_j - \rho_i + \epsilon_i))\right) \exp(-\epsilon_i) d\epsilon_i \\ &= \int_{\rho_i + b}^{\infty} \exp\left(-\exp(-\epsilon_i) \sum_{j=1}^{k} \exp(-(\rho_j - \rho_i))\right) \exp(-\epsilon_i) d\epsilon_i. \end{aligned} \tag{C.2}$$

Set $t = -\exp(-\epsilon_i)$. This means $dt = \exp(-\epsilon_i) d\epsilon_i$. The demand for seller i can be rewritten to

$$\begin{aligned} D_i &= \int_{-\exp(-\rho_i - b)}^{0} \exp\left(t \sum_{j=1}^{k} \exp(-(\rho_j - \rho_i))\right) dt \\ &= \left. \frac{\exp\left(t \sum_{j=1}^{k} \exp(\rho_i - \rho_j)\right)}{\sum_{j=1}^{k} \exp(\rho_i - \rho_j)} \right|_{-\exp(-\rho_i - b)}^{0}. \end{aligned} \tag{C.3}$$

Because $\frac{1}{\sum_{j=1}^{k}\exp(\rho_i-\rho_j)} = \frac{\exp(-\rho_i)}{\sum_{j=1}^{k}\exp(-\rho_j)}$, then

$$
\begin{aligned}
D_i &= \frac{\exp(-\rho_i)}{\sum_{j=1}^{k}\exp(-\rho_j)} - \frac{\exp\left(-\exp(-\rho_i-b)\frac{\sum_{j=1}^{k}\exp(-\rho_j)}{\exp(-\rho_i)}\right)\exp(-\rho_i)}{\sum_{j=1}^{k}\exp(-\rho_j)} \\
&= \frac{\exp(-\rho_i)\left[1-\exp\left(-\exp(-b)\sum_{j=1}^{k}\exp(-\rho_j)\right)\right]}{\sum_{j=1}^{k}\exp(-\rho_j)} \\
&= \frac{\exp(-\rho_i)\left(1-\exp\left(-\eta\Sigma\right)\right)}{\Sigma}
\end{aligned}
\tag{C.4}
$$

where

$$
\begin{aligned}
&\eta = \exp(-b), \\
&\Sigma = \sum_{j=1}^{k}\exp(-\rho_j).
\end{aligned}
\tag{C.5}
$$

By the quotient rule,

$$
\begin{aligned}
\frac{\partial D_i}{\partial \rho_i} &= \frac{\begin{array}{l}\left[-\exp(-\rho_i)(1-\exp(-\eta\Sigma)) - \exp(-\eta\Sigma)\eta\exp(-2\rho_i)\right]\Sigma \\ +\exp(-2\rho_i)(1-\exp(-\eta\Sigma))\end{array}}{\Sigma^2} \\
&= D_i\left(-1-\frac{\eta\exp(-\eta\Sigma)}{1-\exp(-\eta\Sigma)}\exp(-\rho_i)+\frac{\exp(-\rho_i)}{\Sigma}\right)
\end{aligned}
\tag{C.6}
$$

For $m \in \{1,2,...,k\}$ and $m \neq i$, by the quotient rule,

$$
\begin{aligned}
\frac{\partial D_m}{\partial \rho_i} &= \frac{\left[-\exp(-\eta\Sigma)\eta\exp(-\rho_i-\rho_m)\right]\Sigma + \exp(-\rho_i-\rho_m)(1-\exp(-\eta\Sigma))}{\Sigma^2} \\
&= D_i\left(-\frac{\eta\exp(-\eta\Sigma)}{1-\exp(-\eta\Sigma)}\exp(-\rho_m)+\frac{\exp(-\rho_m)}{\Sigma}\right)
\end{aligned}
\tag{C.7}
$$

Since $\pi = \sum_{i=1}^{k}(p_i - c)D_i$, therefore,

$$\begin{aligned}\frac{\partial \pi}{\partial p_i} &= D_i + \frac{\partial D_i}{\partial \rho_i} \cdot \frac{p_i - c}{\sigma} + \sum_{m \neq i} \frac{\partial D_m}{\partial \rho_i} \cdot \frac{p_m - c}{\sigma} \\ &= D_i \left[1 - \frac{p_i - c}{\sigma} + \left(\frac{\eta \exp(-\eta \Sigma)}{\exp(-\eta \Sigma) - 1} + \frac{1}{\Sigma} \right) \sum_{j=1}^{k} \exp(-p_j) \frac{p_j - c}{\sigma} \right]\end{aligned} \tag{C.8}$$

The third part in the large bracket is a constant. If $p_i < p_j$ where $i, j \in \{1, 2, ..., k\}$ then, $D_i > D_j$. Therefore, $\dfrac{\partial \pi}{\partial p_i} > \dfrac{\partial \pi}{\partial p_j}$. □

Proof. (**Uniqueness of** p^*) Let ω be the random variable with the cumulative distribution function $G(\omega) = F^k(\omega)$. The probability density function of ω is then $g(\omega) = kf(\omega)F^{k-1}(\omega)$. Note that the product of log-concave functions is also log-concave. Since f and F are log-concave, g is also log-concave. The joint profit can be written to

$$\pi_r(p) = (1 - G(u_0 + p))(p - c). \tag{C.9}$$

At p^*, the first order condition is

$$p^* - c = \frac{1 - G(u_0 + p^*)}{g(u_0 + p^*)} \tag{C.10}$$

Because g is log-concave, $\frac{g}{1-G}$ is monotone non-decreasing. The left-hand side of (C.10) strictly increases in p^* while the right-hand side is monotone non-increasing. Therefore, the maximizer of $\pi_r(p)$ is unique.

By Prekopa's Theorem, $1 - G$ is also log-concave. Therefore, π_r is also log-concave, which means π_r is quasi-concave. Note that $\pi_r'(p) = 0$ has one unique solution. Therefore, if $\overline{u} \leq p^*$, then $\overline{u}$ is the unique maximizer of $\pi_r(p \mid p \leq \overline{u})$. □

Proof. (**Lemma 4**) This proof has two parts. The first part is to show that $\pi(\vec{p})$ is monotone non-increasing in s for $s \geq 0$. The second part is to show that $\pi(\vec{p})$ strictly decreases in s for $s \in [0, s']$ where $p_k = \overline{p}(s')$.

Part 1: For any fixed $\vec{p}$ with $p_1 < p_k$, consider 2 scenarios: A $= \{\vec{p}, s = s_a\}$ and B $= \{\vec{p}, s = s_b\}$ with $0 \leq s_a < s_b$.

Step 1. The probability that a consumer does not buy any product under A and buys some product under B is zero.

Suppose a consumer ends up not buying any product in scenario A. For this consumer, the rewards from sellers he visited are all lower than u_0. We can safely ignore the case where the reward is equal to u_0 because the probability of that event is infinitesimally small. Let S_A be the set of sellers that this consumer visited under scenario A and S_B be the set of sellers that he will visit under scenario B. By Corollary 1, $S_B \subseteq S_A$. Therefore the rewards from S_B are also lower than u_0, and the consumer will also not buy any product under scenario B.

Step 2. Probability a consumer buys a product with p_l under A and buys a product with $p_h > p_l$ under B is zero.

Let l and h be any two prices among $\vec{p}$ with $p_l < p_h$. Therefore, $z_l > z_h$. Consider a consumer who buys product l under scenario A. Let $w_i \equiv \min\{z_i, x_i\}$ for each i. By Proposition 2, the consumer must have $w_l(s_a) > w_h(s_a)$ ①. We can safely ignore the equality because $z_l > z_h$, the probability that $w_l(s_a) = w_h(s_a)$ is arbitrarily small. We will show that the consumer will not purchase from seller h under scenario B by using proof by contradiction. Suppose this consumer purchases h under scenario B. By Corollary 1, $w_l(s_b) \leq w_h(s_b)$ ②. Since $w(s)$ is non-increasing in s, then $w_h(s_b) \leq w_h(s_a)$ ③. By ①, ② and ③, $w_l(s_a) < w_l(s_b)$. Since the reward x_i is constant in 2 scenarios, we must have $w_l(s_b) = z_l(s_b)$ Then by ② we have $z_l(s_2) < \min\{x_h, z_h(s_2)\}$. However, $z_l(s_2) > z_h(s_2)$ because $p_l < p_h$. By contradiction, the proof is complete.

Hence, $\pi(\vec{p})$ is monotone non-increasing in s for $s \geq 0$.

Part 2: For any fixed $\vec{p}$ with $p_1 < p_k$, consider 2 scenarios: A = $\{\vec{p}, s = s_a\}$ and B = $\{\vec{p}, s = s_b\}$ with $0 \leq s_a < s_b \leq s'$ where $p_k = \overline{p}(s')$. Consider consumers with the following properties:

1. $x_1 \in (z_k(s_b), z_k(s_a)) \iff v^*(s_b) + p_1 - p_k < v_1 < v^*(s_a) + p_1 - p_k$
2. $x_2, x_3, ..., x_{k-1} < z_k(s_b) \iff v_i < v^*(s_b) + p_i - p_k \ \forall i \in \{2, 3, ..., k-1\}$
3. $x_k > z_k(s_a) \iff v_k > v^*(s_a)$

The probability that a random consumer has those three properties is

$$[F(v^*(s_a) + p_1 - p_k) - F(v^*(s_a) + p_1 - p_k)]\,[1 - F(v^*(s_a))] \prod_{i=2}^{k-1} F(v^*(s_b) + p_i - p_k),$$

which is strictly positive when $s_a < s_b$. Let $w_i \equiv \min\{z_i, x_i\}$. Then $w_k(s_a) = z_k(s_a) > w_j(s_a)\ \forall j \neq k$. Note that $s_a < s_b \leq s'$ where $p_k = \overline{p}(s')$. Therefore, $z_k(s_a) > z_k(s_b) \geq u_0$. By Corollary 1, these consumers eventually buy from seller k under scenario A.

Because $p_1 < p_k$, then $z_1(s_b) > z_k(s_b)$. From property 1, $x_1 > z_k(s_b)$. We must have $w_1 > z_k(s_b)$. Therefore, $w_1(s_b) > w_j(s_b)\ \forall j \neq 1$ and $w_1(s_b) > u_0$. By Corollary 1, these consumers eventually buy from seller 1 under scenario B. There is a strictly positive probability that a consumer buys from seller k with s_a but switch to seller 1 with s_b. Note that we have already shown in part 1 that if s increases, the probability that a consumer ends up buying from a higher price is 0. Therefore, for any fixed $\vec{p}$ with $p_1 < p_k$, $\pi(\vec{p})$ strictly decreases in s for $s \in [0, s']$ where $p_k = \overline{p}(s')$. $\square$

Lemma 5. *Given any $\vec{p}$ with $p_k > \overline{p}(s)$ where $s < \overline{s}$, then $\pi(\vec{p}) < \pi(\boldsymbol{p'})$ where $p'_i = \min\{p_i, \overline{p}\}\ \forall i \in \{1, 2, ..., k\}$.*

Lemma 3 technically shows that for any price profile, if the highest price, p_k exceeds the upper limit $\overline{p}$ in equilibrium, the cartel can be better off by reducing the price of seller k to $\overline{p}$ given all other prices stay the same. It implies that in the equilibrium, the cartel will not set any of the k prices higher than $\overline{p}$. All k firms will get some consumers visited.

Proof. (**Lemma 5**) For a fixed $\vec{p}$ with $p_k > \overline{p}(s)$ where $s < \overline{s}$

Step 1: Probability that a consumer buys from some seller under $\vec{p}$ and buys from a lower price seller under $\boldsymbol{p'}$ is zero.

Proof: Consider consumers who buy from seller j with the price profile $\vec{p}$. Let p_{j-} (if exists) be a lower price than p_j in $\vec{p}$. Note that $p_j = p'_j$ and $p_{j-} = p'_{j-}$. Let $w_i \equiv \min\{z_i, x_i\}$ and $w'_i \equiv \min\{z'_i, x_i\}$ for each i. By Corollary 1, these consumers must have $w_j \geq w_{j-}$ if the cartel prices are $\vec{p}$. Since $z_j < z_{-j}$, the probability that $w_j \geq w_{j-}$ is infinitesimally small. Given $w_j > w_{j-}$, then $w'_j > w'_{j-}$, by Proposition 2, these consumers will also not buy from seller $j-$ under profile $\boldsymbol{p'}$.

Step 2: Probability that a consumer buys from some seller under $\vec{p}$ and takes the outside option under $\boldsymbol{p'}$ is zero.

Proof: Consider consumers who buys from seller j under price profile $\vec{p}$. These consumers must have $x_j > u_0$ and $w_j \geq u_0$. We can safely ignore the case $x_j = u_0$ because of its infinitesimally small probability. Under $\boldsymbol{p'}$, a consumer opts to outside option only if $x_j \leq 0$ given $w'_j \geq u_0$. Because $w_j = w'_j$, the probability that a consumer buys from seller j under $\vec{p}$ and takes the outside option under $\boldsymbol{p'}$ is zero.

Step 3: Probability that a consumer takes the outside option under $\vec{p}$ but buys the product with the price $\overline{p}$ under $\boldsymbol{p'}$ is strictly positive.

Proof: Probability that a consumer does not buy any product under $\vec{p}$

Let $m \geq 0$ be the number of prices in $\vec{p}$ that are not higher than $\overline{p}$. That means $p_m \leq \overline{p}$ and $p_{m+1} > \overline{p}$. Consider consumers with the following properties

1. $x_i < u_0 \; \forall i \leq m$
2. $\exists i > m \mid x_i > u_0$.

The probability that a consumer has these properties is strictly positive for all $m \geq 0$. Consumers with these properties take an outside option under $\vec{p}$ but buy the product with the price $\overline{p}$ under $\boldsymbol{p'}$. □

Proof. (**Proposition 15**) For $s > \overline{s}_{p^*}$, $p^M = v_s - u_0$. For any deviation price $p_d \leq p^M$, the one-shot deviation profit can be written by

$$(1 - F(u_0 + p_d))\,(p_d - c). \tag{C.11}$$

Because $F(p_d)$ is log-concave, (C.11) is strictly quasi concave. Denote by $\hat{p}$ the unique maximizer of (C.11). It can be shown that $\hat{p} < p^*$. For $s \in (\overline{s}_{p^*}, \overline{s}_{\hat{p}})$, the maximum one-shot deviation profit π_d is constant in s while the full collusion profit decreases in s. Therefore, critical discount factor $\hat{\delta}$ increases in search cost when $s \in (\overline{s}_{p^*}, \overline{s}_{\hat{p}})$

For $s \geq \overline{s}_{\overline{s}_{\hat{p}}}$, the one-shot deviation profit is maximized when the deviating price is closest to the monopoly price: $p^D = p^M = v_s - u_0$. In this case,

$$\begin{aligned}
\pi_c &= \frac{1}{n}\left(1 - F^k(v_s)\right)(v_s - u_0 - c) \\
\pi_d &= (1 - F(v_s))\,(v_s - u_0 - c) \\
\frac{\pi_c}{\pi_d} &= \frac{1}{n}\left(1 + F(v_s) + \cdots + F^{k-1}(v_s)\right).
\end{aligned}$$

Because $\frac{\pi_c}{\pi_d}$ decreases in s, by equation (3.7), $\hat{\delta}$ increases in search cost. □

Proof. ($\hat{p} < p^*$) By definitions, p^* is the maximizer of $(1 - F^k(u_0 + p))(p - c)$ and $\hat{p}$ is the maximizer of $(1 - F(u_0 + p))(p - c)$. Therefore,

$$p^* - c = \frac{1 - F^k(u_0 + p^*)}{kF^{k-1}(u_0 + p^*)f(u_0 + p^*)} \tag{C.12}$$

$$\hat{p} - c = \frac{1 - F(u_0 + \hat{p})}{f(u_0 + \hat{p})}. \tag{C.13}$$

Note that for $k \geq 2$

$$\frac{1 - F^k}{kF^{k-1}f} = \frac{1 - F}{f}\left(\frac{1}{k} + \frac{1}{kF} + \frac{1}{kF^2} + ... + \frac{1}{kF^{k-1}}\right) > \frac{1 - F}{f} \tag{C.14}$$

By contradiction, suppose $\hat{p} \geq p^*$. By (C.12) and (C.14), $p^* - c > \frac{1-F(u_0+p^*)}{f(u_0+p^*)}$. Because f is log-concave, by Prekopa's Theorem, $\frac{1-F}{f}$ is monotone non-increasing, which implies $\frac{1-F(u_0+p^*)}{f(u_0+p^*)} \geq \frac{1-F(u_0+\hat{p})}{f(u_0+\hat{p})}$. Therefore, $\hat{p} - c > \frac{1-F(u_0+\hat{p})}{f(u_0+\hat{p})}$, which contradicts to equation (C.13). □