

ESSAYS ON ECONOMIC THEORY AND BUSINESS APPLICATIONS

by

Yue Li

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Economics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2022

Date of final oral examination: 07/29/2022

The dissertation is approved by the following members of the Final Oral Committee:

Marek Weretka, Professor, Economics

Kenneth Hendricks, Professor, Economics

Xiaoyang Long, Assistant Professor, Operations and Information Management

Raghunath Rao, Professor, Marketing

© Copyright by Yue Li 2022
All Rights Reserved

*To my parents,
for their unconditional love and support.*

Acknowledgments

I appreciate the help and guidance of many people. They taught me to be independent, grateful, content, persistent, and optimistic during difficult times.

My dissertation committee chair and advisor, Marek Weretka, always supported my exploration of topics that interests me, regardless of whether they are considered traditional "economics" topics. I enjoyed all the meetings with him, in-person and virtual, where he provided countless insightful comments and suggestions. William Sandholm and Kenneth Hendricks provided detailed and insightful comments and suggestions on my field paper and dissertation chapters. As a junior researcher, I benefited a lot from their advice. Raghunath Rao, Paola Mallucci, and Xiaoyang Long guided me in applying economic theory to different business environments. I could not have completed this dissertation without them. Xiaoxia Shi, Kimberly Grocholski, and Becca George have also provided invaluable support in my PhD studies. I also wish to express my gratitude to them.

I thoroughly enjoyed spending time with Qi An, Mingjian Li, Jianrong Tian, Ziwei Wang, Yixi Yang, Jian Zhang, Kun Zhang, Ruixiang Zhang, Hanjing Zhu, and Wanjia Zhu. It is truly a blessing to have encountered all these wonderful friends in Madison. I also enjoyed meeting and working with Walaa Awad, Eric Hammond, Aminah Hannef, Hossein Panahi, and many others in the University Apartments. We came together as community leaders and made the University Apartments neighborhood a better place.

Last but not least, I wish to thank Yi Jiang for her company throughout this long journey. I also thank my parents Chunmei Yu and Changying Li for their continuous love and encouragement. No words can express my appreciation and gratitude to my family members. I dedicate this dissertation to them.

Contents

Contents iii

List of Figures vi

Abstract vii

1 A Theory of Limited Edition Contemporary Collectibles

| | | |
|-------|-------------------------------------------------------------------|----|
| 1.1 | <i>Introduction</i> | 1 |
| 1.1.1 | Related Literature | 4 |
| 1.2 | <i>Model Setup</i> | 5 |
| 1.3 | <i>Model Analysis</i> | 7 |
| 1.3.1 | The optimal product development time | 8 |
| 1.3.2 | Why do consumers value firm created scarcity? | 10 |
| 1.3.3 | When should the firm market limited edition products? | 10 |
| 1.3.4 | Discussion | 13 |
| 1.4 | <i>Extensions</i> | 13 |
| 1.4.1 | Premium pricing | 13 |
| 1.4.2 | A rationale to low retail prices and high secondary market prices | 15 |
| 1.4.3 | Partial commitment power on future product release dates . . . | 17 |
| 1.5 | <i>Conclusion</i> | 17 |
| 1.6 | <i>Appendix</i> | 18 |
| 1.6.1 | Solution concept | 18 |
| 1.6.2 | Proofs for the main model | 19 |
| 1.6.3 | Proofs for the Extensions | 22 |
| 1.6.4 | Premium pricing | 22 |
| 1.6.5 | A rationale to low retail prices and high secondary market prices | 25 |
| 1.6.6 | The firm with partial commitment power | 25 |

| | | |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 1.7 | <i>Bibliography</i> | 28 |
| 2 | Sequential Search with a Budget Constraint: An Approximation Algorithm | |
| 2.1 | <i>Introduction</i> | 30 |
| 2.1.1 | Related Literature | 31 |
| 2.2 | <i>The Search Problem</i> | 32 |
| 2.3 | <i>The Greedy Search Algorithm and its Properties</i> | 33 |
| 2.3.1 | The Greedy Search Algorithm (GSA) | 33 |
| 2.3.2 | Main result: asymptotic optimality of the GSA | 35 |
| 2.3.3 | When is the GSA optimal? | 37 |
| 2.4 | <i>Applications</i> | 38 |
| 2.5 | <i>Conclusion</i> | 40 |
| 2.6 | <i>Appendix</i> | 40 |
| 2.7 | <i>Bibliography</i> | 48 |
| 3 | Product Returns and Assortment Decisions: A Strategic Analysis of Online and Offline Competition (with Raghunath Rao and Paola Mallucci) | |
| 3.1 | <i>Introduction</i> | 50 |
| 3.2 | <i>Related Literature</i> | 53 |
| 3.3 | <i>Model Setup</i> | 55 |
| 3.3.1 | Products | 55 |
| 3.3.2 | Retailers | 57 |
| 3.3.3 | Consumers | 57 |
| 3.3.4 | Timing | 58 |
| 3.4 | <i>Model analysis</i> | 59 |
| 3.4.1 | An illustrative example | 59 |
| 3.4.2 | Benchmark: Only BMR present | 61 |
| 3.4.3 | The main model | 62 |
| 3.5 | <i>Application: Buy Online and Return in Store Agreements</i> | 70 |
| 3.5.1 | Model analysis | 71 |
| 3.5.2 | The impact and likelihood of a BORS agreement | 74 |
| 3.6 | <i>Extensions</i> | 76 |
| 3.6.1 | Consumers without access to the OR | 76 |

| | | |
|-------|------------------------------------------------------------------------------------|-----|
| 3.6.2 | Price competition | 78 |
| 3.7 | <i>Conclusion</i> | 80 |
| 3.8 | <i>Appendix</i> | 82 |
| 3.8.1 | Proofs for the main model | 82 |
| 3.8.2 | Proofs for the Application: Buy Online and Return in Store Agreements | 90 |
| 3.8.3 | Proofs for the Extension | 96 |
| 3.8.4 | Proofs for the extension 2: Price competition | 103 |
| 3.8.5 | An example with closed-form solutions | 106 |
| 3.8.6 | Numerical analysis | 107 |
| 3.9 | <i>Bibliography</i> | 108 |

List of Figures

| | | |
|-----|-------------------------------------------------------------------------------------------------------------------|----|
| 1.1 | The timeline. | 7 |
| 1.2 | Consumers' willingness to pay. | 11 |
| 1.3 | Firm profit as a function of product innovation time t (which is a decreasing function of supply S). | 12 |
| 1.4 | The effect of premium pricing. | 15 |
| 1.5 | The resale market structure. | 16 |
| 2.1 | Convergence speed of the GSA. | 37 |
| 3.1 | Optimal product assortment without online competition. | 62 |
| 3.2 | Optimal product assortment in the main model. | 69 |
| 3.3 | Optimal product assortment in the presence of a BORS agreement. | 73 |
| 3.4 | Optimal product assortment in the presence of loyal consumers. | 77 |
| 3.5 | Optimal product assortment in the presence of loyal consumers and a BORS agreement. | 78 |

Abstract

This dissertation contains three chapters, which use game-theoretic models to study firm strategies and consumer behavior in different business environments.

Chapter 1: A Theory of Limited Edition Contemporary Collectibles This chapter studies a contemporary collectibles producer that periodically introduces new product designs. When consumers cannot observe the firm's future product development plans, the equilibrium rate of product introductions is too fast. This inefficiency can be alleviated by a commitment to selling limited editions, i.e., supplying less than the market demand. Consumers value firm-created scarcity and their willingness to pay increases as supply decreases. Selling limited editions is optimal when products have high collection value.

Chapter 2: Sequential Search with a Budget Constraint: An Approximation Algorithm In this chapter, I study a sequential search problem where an agent faces a set of closed boxes with unknown values. She knows the purchase price and value distribution of each box, as well as the search cost to open each box and reveal its value. Box purchases are subject to a budget constraint, and the agent can only purchase open boxes. Her objective is to maximize the total value of all purchased boxes, plus any leftover purchase budget, minus search costs. The problem is to decide which boxes to open, in what order, when to stop opening boxes, and which open boxes to purchase. I present a computationally simple Greedy Search Algorithm and show that it well approximates the optimal solution when boxes are cheap relative to the total budget (regardless of whether the budget constraint binds).

Chapter 3: Product Returns and Assortment Decisions: A Strategic Analysis of Online and Offline Competition (with Raghunath Rao and Paola Mallucci) This chapter studies a model of product assortment decisions in the context of online-offline retail competition. A brick-and-mortar store facing online competition may trade off higher sales for more foot traffic by stocking its limited store shelves with harder-to-fit products. This provides a rationale for the steady growth of specialty

stores within the shrinking footprint of traditional retailing. The model is applied to study a nascent return policy that allows consumers to return online purchases directly to a competing brick-and-mortar store. This return policy could hurt marginalized communities with poor access to digital and financial resources, and thus, exacerbate the "digital divide" among marginalized communities.

CHAPTER 1

A THEORY OF LIMITED EDITION CONTEMPORARY COLLECTIBLES

1.1 Introduction

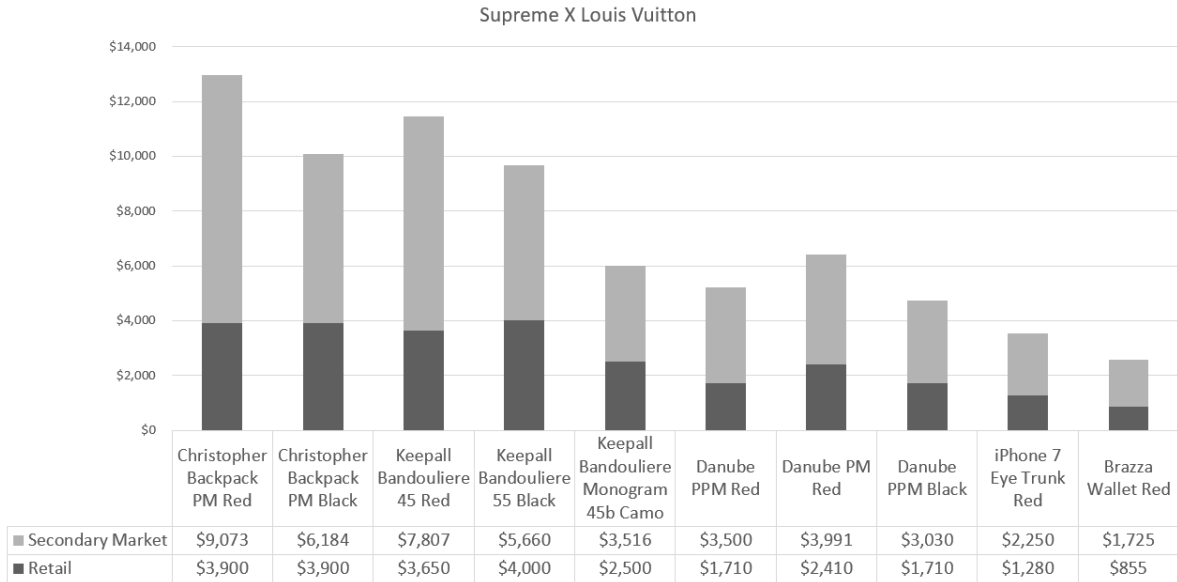
Contemporary collectibles include many different products such as trading cards, toys, comic books, luxury accessories, etc (Thomann, 2017). They can be manufactured at scale and distinctively display two characteristics. First, consumers derive a lump-sum collection value from obtaining each version of the product. Second, consumers enjoy a flow consumption value that is decreasing over time. For example, the Pokemon Company sold over 4 billion Pokemon trading cards in 2020. The trading cards are collected by consumers of all ages, and their producers hold numerous tournaments where collectors can trade and battle with their local community to earn prizes and cash rewards. However, the flow consumption value of each card declines over time because only new cards can be used in official tournaments.

In this chapter, we focus on the phenomenon of firms selling limited edition collectibles. Trading cards such as Pokemon TCG and Yu-Gi-Oh TCG have limited print runs, and the resulting shortage forces major retailers such as Target to impose maximum purchase quantities (Kelly, 2021). This phenomenon is not unique to trading cards. For example, many fashion brands such as Supreme, Yeezy, and BAPE are also known for selling entire collections in limited quantities. On product release dates, long lines are formed outside their physical stores and products typically sell out within hours.

As shown in Table 1.1, the secondary market prices for Supreme and Louis Vuitton

collaboration products are much higher than their retail prices. Similar patterns also arise in other markets for collectibles, suggesting that the demand for limited editions consistently exceeds supply at retail prices. There are two puzzling questions. First, why do firms intentionally leave money on the table, and refrain from closing the demand-supply gap by either raising prices or increasing supply? Second, why do consumers exhibit a higher willingness to pay for limited edition products?

Table 1.1: This table shows the retail prices and secondary market prices for 10 Supreme X Louis Vuitton collaboration products. Data source: highsnobiety.com



To answer the first question, we analyze a firm's strategic choice of product development plans. When consumers cannot observe the firm's future product release dates, the firm releases new collectibles too frequently. Suppose that the firm has commitment power in selling limited editions. Then, selling limited editions can alleviate the inefficient rush to release new products. By restricting future supply, the firm credibly reduces the profitability of future product releases. This incentivizes the firm to wait longer for the next product release and slows down its inefficiently fast product clockspeed. Selling limited editions, however, comes at a direct cost of lower sales volumes. It is optimal for the firm to market limited editions if the collection value is large, the flow value of consumption is large, or the fixed cost of introducing new products is low. From a welfare perspective, my model contrasts the classical view that an under-provision of products is associated with higher deadweight loss and lower social welfare (e.g., Tirole, 1988).

To answer the second question, we abstract away from explicit preferences for scarce goods or social status, and assume that consumers only care about collecting and consuming a contemporary collectible (e.g., sneakers, handbags, luxury cars, trading cards, etc.). Due to consumers' preference for using the newest generation product (e.g., Pesendorfer, 1995), their willingness to pay depends on product clockspeed. In equilibrium, willingness to pay is higher for limited editions, because of a longer *anticipated* product clockspeed and better quality. This provides a new micro foundation for consumers' preference for limited editions.

While the main model focuses on the simple case of homogeneous consumers, relaxing this assumption generates further insights. In the first two extensions, we analyze a heterogeneous consumers model. First, we find that creating a luxurious brand image through premium pricing (i.e., always setting prices above the monopoly price) is strategically similar to marketing limited editions. Both strategies reduce the profitability of future product releases, which lowers the marginal cost of increasing product development time, and alleviates the firm's inefficient rush problem. This result draws a novel link between premium pricing and selling limited editions - both of which are commonly practiced in the contemporary collectibles industry. The second extension rationalizes the puzzling combination of low retail prices and high secondary market prices for limited editions. Refraining from raising retail prices to match secondary market prices is also a way to appropriate only part of the future profit, which restores efficiency by slowing down product clockspeed. With a frictionless secondary market, the equilibrium involves a buying frenzy where all consumers attempt to buy at retail price, and low valuation consumers profit by reselling on the secondary market.

The last extension examines the assumption of the firm's lack of commitment in future product development plans. Sometimes, it is more natural to assume that a fraction of (loyal) consumers know about the brand convention for new product releases. My results remain robust and still survive qualitatively, as long as some consumers are agnostic to the product clockspeed. Notably, if the firm commits to reducing supply, *all* consumers - both loyal and non-loyal - anticipate longer product clockspeed, and their willingness to pay rises.

1.1.1 Related Literature

Firm-created shortages have been studied in the literature of economics, marketing, and psychology. Early works in psychology provide evidence for consumer preferences for scarcity (Veblen and Galbraith, 1973; Worchel et al., 1975). Notably, consumers desire scarcity that is intentionally created but does not value the scarcity caused by accidents in the supply chain (Verhallen, 1982; Verhallen and Robben, 1994). We contribute to this literature by providing a novel micro foundation for why consumers value artificially created scarcity.

From the firm's perspective, very few papers rationalize the strategy of consistently limiting supply to be below market demand. Becker (1991) shows that, in the short run, a capacity-constrained restaurant owner who faces non-monotonic demand may want to increase prices to create excess demand. Becker further argues that, in the long run, equilibrium multiplicity and uncertainty may prevent the owner from investing to increase capacity. Alternatively, if sales take place before demand uncertainty is resolved, then the firm can sell limited editions to create a buying frenzy (DeGraba, 1995), signal product quality (Stock and Balachander, 2005; Balachander et al., 2009), or attract scalpers (Su, 2010). Other papers exogenously assume that consumers attach value to owning a product not possessed by others, and analyze whether two competing firms should add limited editions to their product line (Amaldoss and Jain, 2008; Balachander and Stock, 2009). In contrast, this chapter studies a model where increasing capacity is costless, there are no demand shocks, and consumers do not explicitly value scarcity. My model always has a unique equilibrium, in which consumers' willingness to pay rises as supply decreases, and the firm can optimally choose to sell limited editions.

The concept of limited editions is closely related to conspicuous consumption. Consumers often convey personal attributes through their brand choice (Kuksov, 2007). Previous lab experiments identify the existence of elite consumers, whose utility from a product decreases as more people consume the same product (e.g. Amaldoss and Jain, 2005a). Amaldoss and Jain (2010) show that offering limited editions can increase sales to elite consumers. Other papers analyze how conspicuous consumption affects competition (Amaldoss and Jain, 2005b), dynamic pricing (Rao and Schaefer, 2013), and branding (Amaldoss and Jain, 2015). However, works in this area do not explain why a firm may supply less than demand in equilibrium.

This chapter is also related to the works on product cycles of durable goods (e.g., Pesendorfer, 1995) and a firm's optimal research and development plans (e.g., Fishman

and Rob, 2000). Pesendorfer (1995) proposes a model of fashion, where products are used in a dating game to signal fashion. A product has the highest value at release, and its value falls when a new design is created. My model exogenously assumes that the flow value of a product declines after a new generation is released, and we analyze a firm's optimal product development time. As in Fishman and Rob (2000), we find that a durable good producer's choice of product development time is socially inefficient. My work further demonstrates how marketing limited edition products alleviates this inefficiency.

The rest of the chapter is organized as follows. We set up the main model in Section 1.2, and analyze it in Section 1.3. Section 1.4 studies the robustness of my model, and draws a link between creating scarcity and adopting luxury pricing. Section 1.5 concludes the chapter. Appendix 1.6.1 formally states the solution concept used in this chapter. All proofs can be found in Appendix.

1.2 Model Setup

We consider a continuous time model, where a firm and consumers share a common exponential present discount factor $r > 0$. The positive time discount guarantees a well-defined profit maximization and a finite equilibrium profit.

The firm. The firm faces an infinite horizon planning problem, where it introduces a series of new designs of a collectible at dates $\{t_n\}_{n=1}^{\infty}$. The quality of each design is determined by the product development time, i.e., the n th generation product has quality $v_n = v(t_n - t_{n-1})$. We assume that $v(0) = 0$, $v(t)$ is increasing, and $v''(t) < v'(t)(1 + \frac{r}{1 - e^{-rt}})$ for all $t > 0$. The final condition is satisfied if $v(t)$ is concave, linear, or no more convex than e^t .

The marginal cost of production is zero. But there is a fixed cost, $F > 0$, associated with each product release, which captures the firm's advertisement and marketing expenditures. This fixed cost also serves as friction that prevents the firm from introducing new products too frequently.

Consumers. There is a unit mass of homogeneous consumers. Each consumer obtains two sources of utility from obtaining a collectible - collection value and consumption value. The collection value is a non-negative constant, $\theta_c \geq 0$. The consumption value

is a flow payoff that depends on product quality. Standing at time $\tau = t$, a consumer who uses the n th generation product obtains a present discounted flow utility of

$$\begin{cases} e^{-r(\tau-t)}(1 + \theta_f)v_n, & \tau < t_{n+1} \\ e^{-r(\tau-t)}v_n, & \tau \geq t_{n+1} \end{cases}$$

at each future instant τ . Here, the parameter $\theta_f > 0$ represents the boosted flow consumption value, which is obtained only when using the newest generation. Therefore, at time $\tau = t_n$, consumers are willing to pay

$$\theta_c + \int_{t_n}^{\infty} e^{-r(\tau-t_n)}(1 + \theta_f \times 1_{\tau < t_{n+1}})v_n d\tau - \int_{t_n}^{\infty} e^{-r(\tau-t_n)}v_m d\tau \quad (1.1)$$

for replacing a generation $m < n$ product with a new product n .

This setup captures the two main characteristics of many contemporary collectibles such as trading cards, sneakers, handbags, etc. First, collection value is an important feature of contemporary collectibles (e.g., U-Dox, 2014). My results are qualitatively robust to alternative assumptions on the collection value, such as allowing collection value to increase in product quality or heterogeneity among consumers (see Section 1.4.1 and 1.4.2). Second, the flow consumption value decreases after a newer generation is released. This assumption captures the fact that only new trading cards can be used in tournaments, and accessories become less fashionable when newer designs are introduced (Pesendorfer, 1995).

Timing. At $\tau = 0$, the firm commits to a supply of S for each product generation. If $S < 1$, then we say that the firm sells *limited edition* products.

At $\tau = t_n > 0$, the n th generation product is released. First, consumers observe the quality of the new product v_n , form beliefs on the next product release date t_{n+1} , and decide whether to buy the new product. If the demand exceeds supply, then products are allocated randomly among all potential buyers. Then, the firm chooses the next product release date t_{n+1} . At $\tau = t_{n+1} > t_n$, the two steps above are repeated for the $(n + 1)$ th product generation.

Solution concept. As in Fishman and Rob (2000) and Plambeck and Wang (2009), we consider a pure strategy, stationary, rational expectations equilibrium. We drop the

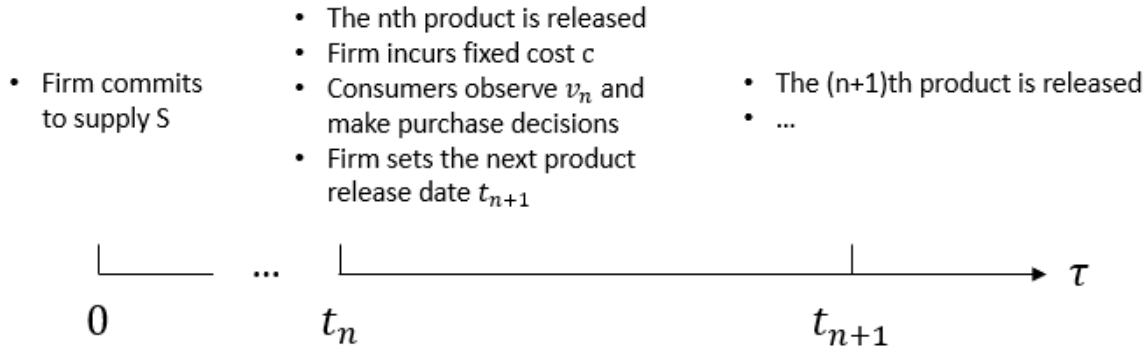


Figure 1.1: The timeline.

subscript, n , to denote variables in the stationary equilibrium. The properties of the equilibrium are described below.

First, the firm releases new products at a constant pace, $t_{n+1} - t_n \equiv t$. Consequently, the quality of new products is constant over time, i.e., $v_n = v(t_n - t_{n-1}) \equiv v(t)$. Second, consumers have correct beliefs about future product release dates, i.e., $\tilde{t}_{n+1} - t_n \equiv \tilde{t} = t$. Third, consumers make optimal purchase decisions. Lastly, the firm's supply S and product clockspeed t maximize its present discounted profit. A more detailed description of the solution concept is presented in Appendix A.

1.3 Model Analysis

We first fix the supply S and derive the firm's profit maximizing product development time in Subsection 1.3.1. We show that the inability to commit to future product release dates leads to inefficiently fast product clockspeed. In Subsection 1.3.2, we demonstrate that selling limited edition products (i.e., choose $S < 1$) alleviates this inefficiency at the cost of decreasing sales. In equilibrium, consumers exhibit a higher willingness to pay for limited edition products, even though they do not have an explicit preference for scarcity. In Subsection 1.3.3, we analyze when should a firm market limited edition products. Finally, we discuss the welfare implications and contrast my findings to the antitrust literature in Subsection 1.3.4. Proofs are collected in the Appendix.

Table 1.2: A summary of notations.

| Notation | Definition |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------|
| τ | Time. |
| r | The continuous time present discount factor. |
| t_n | The release date for the n th generation product. |
| \tilde{t}_n | Consumers' anticipated release date for the n th generation product. (In equilibrium, $\tilde{t}_n \equiv t_n, \forall n.$) |
| t | The product innovation time in the stationary equilibrium. (In equilibrium, $t \equiv t_n - t_{n-1}, \forall n.$) |
| \tilde{t} | Consumers' anticipated product innovation time in the stationary equilibrium. (In equilibrium, $\tilde{t} \equiv t.$) |
| $\beta(t)$ | $\beta(t) = \frac{1-e^{-rt}}{r}.$ |
| v_n | The quality of the n th generation product. (In equilibrium, $v_n \equiv v(t), \forall n.$) |
| θ_c | The non-negative collection value of a product. |
| θ_f | The boosted flow value of using a new product. |
| F | The fixed cost of launching a new product. |
| S | The total supply for each product generation. |

1.3.1 The optimal product development time

In equilibrium, product quality is constant across generations. So, we can simplify Equation (1), and a consumer's willingness to pay for a new product is always

$$\theta_c + \int_t^{t+\tilde{t}} e^{-r(\tau-t)} \theta_f v(t) d\tau = \theta_c + \theta_f v(t) \beta(\tilde{t}), \quad (1.2)$$

where $\beta(\tilde{t}) = \frac{1-e^{-r\tilde{t}}}{r}$.¹ Notice that consumers can observe the quality of the current product, but cannot observe the firm's future plans. So, $t + \tilde{t}$ is a consumer's *anticipated* release date of the next product.

Since consumers are homogeneous, the firm optimally sets the price equal to con-

¹We assume that, at $\tau = 0$, consumers are initially endowed with a product of quality $v_0 = v(t)$. This simplifies the model so that the initial product release is no different from all future product releases.

sumers' willingness to pay, $p = \theta_c + \theta_f v(t)\beta(\tilde{t})$. Given the supply $S \leq 1$, the profit maximization can be recursively expressed as

$$\begin{aligned} \pi^m = \max_t e^{-rt} [pS - F + \pi^m] \\ \text{s.t. } p = \theta_c + \theta_f v(t)\beta(\tilde{t}) \end{aligned} \quad (1.3)$$

We impose $F > \theta_c$, so that the firm cannot make an infinite profit by setting price equal to the collection value and releasing new products very frequently.

The first order condition with respect to t shows the trade-offs from increasing product development time,

$$\underbrace{e^{-rt}\theta_f v'(t)\beta(\tilde{t})S}_{MB} - \underbrace{\frac{re^{-rt}}{\beta(t)}[\theta_c S + \theta_f v(t)\beta(\tilde{t})S - F]}_{MC} = 0. \quad (1.4)$$

If the firm increases t the consumers will observe a higher quality $v(t)$ upon the next product release. This leads to a marginal benefit of higher future prices and profits. However, as t increases, the profit from the next generation gets delayed further into the future. To maximize profit, the firm equals the marginal benefit with the marginal cost.

Lemma 1.1. *Given supply $S > 0$, an equilibrium exists and is unique. The firm's equilibrium product development time is implicitly given by the unique solution to the transcendental equation²*

$$v(t)\beta(t) - v'(t)[\beta(t)]^2 = \frac{F - \theta_c S}{\theta_f S}. \quad (1.5)$$

Now, briefly consider the case where consumers can observe future product release dates. Then, consumers' willingness to pay in Equation (2) changes to $\theta_c + \theta_f v(t)\beta(t)$. In this case, the firm should still set a price equal to consumers' willingness to pay. However, the marginal benefit of raising t will increase from $e^{-rt}\theta_f v'(t)\beta(\tilde{t})S$ to $e^{-rt}\theta_f v'(t)\beta(t)S + e^{-rt}\theta_f v(t)\beta'(t)S$. That is, if the firm's future plans are publicly observable, consumers' current willingness to pay increases because they know that the current product will have a flow value for a longer period of time. This marginal benefit of raising t is absent when consumers cannot observe future product release dates, leading to an inefficiently low equilibrium product development time.

²Transcendental equations often do not have closed-form solutions.

Proposition 1.1. *When consumers cannot observe future product release dates, the firm's equilibrium product development time and product clockspeed are inefficiently short.*

1.3.2 Why do consumers value firm created scarcity?

Why do consumers value firm-created scarcity? This subsection provides one possible answer.


Suppose that the firm chooses to market limited edition products (i.e., $S < 1$). Notice that the right hand side of Equation (5) is decreasing in S , and the left hand side is increasing in t . Thus, the equilibrium product development time t is decreasing in S . Since consumer's willingness to pay, $\theta_c + \theta_f v(t)\beta(t)$, is an increasing function of t , we have that consumer's willingness to pay also increases as supply decreases.

Proposition 1.2. *In equilibrium, consumers' willingness to pay is decreasing in S .*

This result provides a micro foundation for consumers' value of firm-created scarcity. Even if consumers do not have an explicit preference for limited edition products, in equilibrium, they are still willing to pay more for firm-created scarcity. Intuitively, consumers are reluctant to pay for the current generation, because they worry that the firm will release the next generation soon and the current design will quickly become obsolete. By committing to selling limited editions, the firm reduces its future profitability and alleviates the inefficient rush to release the next product. The longer product clockspeed increases (future) product quality $v(t)$, and each product (including the current generation) will have a flow value for an extended period of time. Both factors increase consumers' equilibrium willingness to pay in equilibrium. Figure 1.2 illustrates consumers' willingness to pay as a function of the firm's product development time t (which, in equilibrium, is determined by the firm's supply S).

1.3.3 When should the firm market limited edition products?

While committing to limited edition products alleviates the inefficient rush to release the next product, it comes at an obvious cost of lower sales for all generations. In this subsection, we analyze when should a firm market limited edition products.



Note. This figure plots consumers' willingness to pay with $v(t) = t$, $r = \theta_c = 0.1$, and $F = \theta_f = 1$. Notice that a larger equilibrium t is equivalent to a smaller S .

Figure 1.2: Consumers' willingness to pay.

From Equation (3), we can rewrite the equilibrium profit as

$$\pi^m = \frac{e^{-rt}}{1 - e^{-rt}} \{[\theta_c + \theta_f v(t)\beta(t)]S - F\}.$$

Using $\beta(t) = \frac{1 - e^{-rt}}{r}$ and Equation (5), we can transform the profit into a function of the equilibrium product development time, t ,

$$\pi^m(t(S)) = \frac{\theta_f F}{r} \times \frac{[1 - r\beta(t)]\beta(t)v'(t)}{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 + \theta_c}.$$


Here, t is implicitly determined by supply S (see Lemma 1.1). At time $\tau = 0$, the firm chooses S and determines whether to sell limited edition products,

$$\max_S \pi^m(t(S)).$$

The next proposition specifies the conditions for marketing limited editions and commit to a supply that is consistently below demand.

Proposition 1.3. *The firm optimally markets limited edition products (i.e., choose $S < 1$) if $\theta_c > 0$ and $\frac{F - \theta_c S}{\theta_f S}$ is sufficiently small.*

Proposition 1.3 shows that a necessary condition for marketing limited editions is that the product must have some collection value, at least for a fraction of the consumers (see Sections 1.4.1 and 1.4.2). This explains why limited editions are more prevalent



Note. The left (right) shows a case where creating limited edition products is (not) profitable. Here, we use parameters $v(t) = t$, $r = \theta_c = 0.1$, and $F = \theta_f = 1$.

Figure 1.3: Firm profit as a function of product innovation time t (which is a decreasing function of supply S).

among products such as sneakers, handbags, luxury cars, trading cards, etc. It is optimal to sell limited editions if (1) the boosted flow value of a new product is large (θ_f is large); or (2) the collection value, θ_c , is large; or (3) the fixed cost of launching a product, F , is small; or (4) the fixed cost associated with a product release, F , is small. Conditions (1) - (3) increase the profitability of the next generation, and condition (4) reduces the cost of releasing new designs. All four conditions lower the equilibrium product development time, and exaggerate the firm's inefficient rush to release future generations. Under extreme parameter values, the equilibrium product development time and profit can both approach zero if the entire demand is met.

To increase product development time and restore efficiency, the firm needs to credibly reduce the profitability of future product releases. This can be achieved by a commitment to selling limited editions, which forgoes future profit by reducing supply. It is optimal to market limited editions if the marginal benefit (i.e., alleviating the inefficient rush to release future generations) exceeds the marginal cost (i.e., lower sales volume). Figure 1.3 illustrates the firm's profit as a function of its product development time t (which, in equilibrium, is determined by its supply S).

1.3.4 Discussion

1.3.4.1 Commitment to product release dates

So far, we have assumed that the firm can commit to selling limited editions, but has no commitment power in its future product release dates. Firm-created scarcity can be easily conveyed to consumers through unique product serial numbers, empty store shelves, consumers forming long lines outside the stores at release dates, etc. In contrast, it can be harder for the firm to communicate its future plans. Even when the firm has established a reputation and can credibly convey its future plans to a fraction of the (loyal) consumers, it may still sell limited editions (see Section 1.4.3). Notably, in this case, a supply reduction raises the willingness to pay for *all* consumers.

1.3.4.2 Welfare

In the traditional antitrust literature, a decrease in equilibrium quantity is typically associated with larger monopoly power, higher deadweight loss, and lower social welfare. This chapter shows an interesting counter-example. In my simple framework with homogeneous consumers, the firm optimally sets prices equal to all consumers' willingness to pay and the consumer surplus is zero. Thus, social welfare is equal to the firm's profit. Proposition 1.3 shows that, under certain conditions, a lower supply can actually increase social welfare. The reason is that marketing limited editions serves as a remedy for the inefficiently short product clockspeed.

This result is not an artifact of my homogeneous consumer assumption. As shown in Sections 1.4.1 and 1.4.2, when consumers have private types and the demand function is downwards sloping, selling limited editions can increase both firm profit and consumer surplus.

1.4 Extensions

1.4.1 Premium pricing

My framework provides a new rationale for premium pricing.

To analyze pricing decisions, we need to relax the simplifying assumption of homogeneous consumers and obtain a downward sloping demand curve. Suppose consumers' collection value, θ_c , is independent and identically distributed according to the cumula-

tive distribution function $G(\cdot)$. Given the current product's development time t and consumers' anticipation of the next product's development time \tilde{t} , the profit of the current generation is given by $p \times Pr[\theta_c + \theta_f v(t)\beta(\tilde{t}) \geq p] = p\{1 - G[p - \theta_f v(t)\beta(\tilde{t})]\}$. Suppose the distribution of θ_c has increasing hazard rate.³ Then, the profit of each period has a unique maximizer which solves the equation $p^m = \frac{1 - G[p^m - \theta_f v(t)\beta(\tilde{t})]}{g[p^m - \theta_f v(t)\beta(\tilde{t})]}$, where g is the density of θ_c .


Without a commitment on pricing, the firm always sets price $p = p^m$ in a stationary equilibrium, and sells to consumers with $\theta_c \geq p^m - \theta_f v(t)\beta(t)$. When the firm commits to premium pricing and sets price $p > p^m$, sales decline and the firm obtains less profit from subsequent product releases. Analogous to the main model, this commitment also reduces the inefficient rush of releasing new products. Thus, the product clockspeed slows down, each consumer's willingness to pay increases, and the demand curve shifts out. Similar to Proposition 1.3, it is optimal to commit to premium pricing if the boosted flow value of a new product is large, the collection value is large, or the fixed cost of releasing new products is low. Figure 1.4 illustrates a situation where premium pricing is profitable.

Proposition 1.4. *A commitment to premium pricing (set price $p > p^m$)*

- *increases the equilibrium product development time t ;*
- *shifts the demand curve shifts out;*
- *and increases profits if θ_f is large, or θ_c is large, or F is small.*

Proposition 1.4 shows that premium pricing and selling limited editions are strategically similar, both of which slow down the product clockspeed and increase consumers' willingness to pay. This equivalence result sheds new light on the frequent usage of premium pricing and limited editions in the contemporary collectibles industry. The model details and proofs for this extension can be found in the Appendix.

³Many distributions such as uniform, normal, and exponential distributions all have monotone hazard rate.



Note. Without premium pricing, the firm sets monopoly price p_1^m and profits $\pi_1 = p_1^m S_1$ each period. Premium pricing ($p_2 > p_2^m$) shifts the demand curve outwards, and raises profit to $\pi_2 = p_2^m S_2$ each period.


Figure 1.4: The effect of premium pricing.

1.4.2 A rationale to low retail prices and high secondary market prices

Consider the same heterogeneous consumer framework as in Section 1.4.1. As established in Propositions 1.3 and 1.4, a commitment to partially forfeit future profit is needed to alleviate the inefficient rush of releasing new products. For this purpose, if a combination of low supply and low retail price is deployed, then we observe excess demand at the given retail price. For a concrete example, recall from Section 4.1 that, to maximize the profit of the current generation, the firm sets retail price $p = p^m$ and sells to consumers with $\theta_c \geq p^m - \theta_f v(t) \beta(t)$. One way to reduce the profitability of future releases is to keep supply at the monopoly level and lower price to $p = \alpha p^m$, $\alpha < 1$. This way, the firm creates an excess demand of $D - S = G[p^m - \theta_f v(t) \beta(t)] - G[\alpha p^m - \theta_f v(t) \beta(t)] > 0$.

Now, suppose there is a secondary market, where new products can be resold frictionlessly. Assume that each consumer can purchase at most one product.⁴ Due to the low retail price and low supply, all consumers are incentivized to buy at the retail price. Low valuation scalpers (those with $\theta_c < p^m - \theta_f v(t) \beta(t)$) who successfully obtain a product sells in the secondary market to high valuation consumers (those with

⁴For example, Supreme holds free lotteries to ration demand at release dates. Each phone number can only enter in the drawing once.



Note. When the firm supplies the entire demand, the per period profit is $\pi_1 = p_1^m S_1$. When the product is limited edition ($S_2 < D_2$), the demand curve of each period shifts outwards, and the firm obtains a higher profit $\pi_2 = p_2^m S_2$ for each product generation. Shortage at the retail price p_2^m leads to a high secondary market price $p_2^{secondary}$.

Figure 1.5: The resale market structure.

$\theta_c \geq p^m - \theta_f v(t)\beta(t)$) who did not successfully obtain a product in the primary market. The secondary market clears at an equilibrium price of p^m , which is strictly greater than the retail price αp^m . (See the Appendix for details.) The firm refrains from raising the retail price to match secondary market prices, to appropriate only part of the future profit, and resolve the inefficient short product clockspeed. The logic is analogous to Propositions 1.3 and 1.4.

Proposition 1.5. *Suppose there is a frictionless secondary market for unopened products. When θ_f is large, or θ_c is large, or F is small, the firm optimally sells limited editions. In equilibrium, the secondary market clearing price is higher than the retail price, and scalpers benefit from the price gap.*

As illustrated in Figure 1.5, the firm can also create scarcity by fixing the monopoly supply and lowering the supply. This is strategically equivalent to reducing price and fixing supply at the monopoly level. Both strategies forfeit some future profit to slow down product clockspeed, and activate a secondary market.

1.4.3 Partial commitment power on future product release dates

In practice, some producers have a convention of releasing new products at a constant pace. For example, Hermes releases two new scarves every year since 1986. It is natural to assume that a fraction of loyal consumers know this tradition, and believes that Hermes has credibility in future product release plans.

To capture this in my model, we assume that only a fraction $\gamma \in (0, 1]$ of the consumers is unaware of the firm's commitment to future product releases. We find that, as long as $\gamma > 0$, the imperfect commitment on future product release dates still leads to a rush to release new products. Marketing limited editions alleviates this inefficiency, so that *all* consumers anticipate longer product cycles and have a higher willingness to pay in equilibrium.

Denote \hat{t} as the unique positive solution to the equation $[1 - (1 - \gamma)e^{-r\hat{t}}]v(\hat{t}) - v'(\hat{t})\beta(\hat{t}) = 0$. We have the following result.

Proposition 1.6. *Suppose that a fraction $\gamma \in (0, 1]$ of the consumers is unaware of the firm's commitment to future product releases.*

- *In equilibrium, all consumers' willingness to pay decreases in S .*
- *If \hat{t} is small enough, then the firm sells limited editions when θ_f is large, or θ_c is large, or F is small.*

Here, \hat{t} can be interpreted as the shortest possible equilibrium product development time. In the case of my main model where $\gamma = 1$, we have $\hat{t} = 0$. So, with $\gamma = 1$, the firm's inefficient rush can lead to near zero profit under extreme parameter values. Proposition 1.6 states that, when \hat{t} is small enough, then we recover the large inefficiency of the main model, so that the firm opts to sell limited editions. The assumption of small \hat{t} is satisfied if $v(t)$ is sufficiently concave or if γ is large. The model details and proofs can be found in the Appendix.

1.5 Conclusion

This chapter proposes a rationale for the prevalence of limited-edition contemporary collectibles. We also provide a micro foundation for consumers' preference for firm-

created exclusivity. The extensions also produce insights in scalping, reselling, and premium pricing.

However, this dissertation chapter only takes a small step in understanding the market of contemporary collectibles. Contemporary collectibles exhibit many other interesting and distinct features. First, contemporary collectibles have heterogeneous rarities (which are designed by the firm) and are often sold in blind boxes. On resale markets, where the uncertainty of blind boxes is resolved, rare collectibles sell for much higher than retail price, while common collectibles (also known as bulks) are traded at very low prices. Second, collectibles perish over time, and many collectors invest significant time and effort to preserve their collections. Finally, collectible producers such as trading card companies attempt to boost the collectible's flow consumption value by investing in tournaments and competitions. In the future, follow-up works will be combined with this chapter to study the optimal rarity design of blind boxes, reselling, price trends for perishable collectibles, and the optimal design of events for collectibles (e.g., trading card tournaments).

1.6 Appendix

1.6.1 Solution concept

Here, we provide a detailed description of the equilibrium concept. At each product release date t_n , the consumers decide whether to purchase the new product. Let $\chi_n = \{0, 1\}$ denote the consumers' purchase decisions for the n th design, where $\chi_n = 0$ denotes not wanting to purchase and $\chi_n = 1$ denotes wanting to purchase (and entering the lottery when $S < 1$). We consider a pure strategy, stationary, Markov perfect equilibrium with rational expectations. The equilibrium selection criteria are listed below.

In a stationary Markov perfect equilibrium, state variables are constant across periods and only include variables that are directly payoff relevant (Maskin and Tirole, 1988). Therefore, the state variables are the design quality of the current product $v_n \equiv v(t)$ and the development time for the next generation $\tilde{t}_{n+1} - t_n \equiv \tilde{t}$. Note that the state variables should also include the vector of previous generation design qualities, as this also influences consumers' purchase decisions. But in the steady state, all entries of this vector must be equal to $v(t)$.

Markov strategies can only condition on the current state, so strategies are also stationary. Optimality of the firm requires that, given the state and anticipating consumers' purchase decisions χ , the firm chooses $p : \{v, \tilde{t}\} \rightarrow \mathbb{R}_+$ and $t : \{v, \tilde{t}\} \rightarrow \mathbb{R}_+$ to maximize profit. Optimality of the consumers requires that, given the state and firm's choices p and t , consumers choose $\chi : \{v, \tilde{t}, p\} \rightarrow \{0, 1\}$ to maximize utility

$$U(v, \tilde{t}) = \begin{cases} \theta_c + \theta_f \int_{t_n}^{\tilde{t}_{n+1}} e^{-r(\tau-t_n)} v_n d\tau - p, & \chi = 1 \\ 0, & \chi = 0 \end{cases}$$

$$= \begin{cases} \theta_c + \theta_f v(t) \beta(\tilde{t}) - p, & \chi = 1 \\ 0, & \chi = 0 \end{cases}$$

We assume that consumers break indifference in favor of the firm (i.e., $\chi = 1$), because the firm can charge price $p - \epsilon$ and induce all consumers to purchase (or enter the lottery). Both consumers and the firm use pure strategies.

Rational expectations is the fixed point of the perceived and actual law of motion. In my context, rational expectations together with stationarity require $\tilde{t} \equiv t$ both on and off the equilibrium path. This assumption is also imposed in other papers such as Fishman and Rob (2000).

We seek the unique equilibrium under these selection criteria.

1.6.2 Proofs for the main model

Proof of Lemma 1.1. Given S , the firm's optimization can be re-written as

$$\pi^m = \max_t e^{-rt} \{[\theta_c + \theta_f v(t) \beta(\tilde{t})] S - F + \pi^m\}.$$

The first order condition with respect to t is

$$e^{-rt} \theta_f v'(t) \beta(\tilde{t}) S - r e^{-rt} \{[\theta_c + \theta_f v(t) \beta(\tilde{t})] S - F + \pi^m\} = 0.$$

In equilibrium, consumers hold correct beliefs, i.e., $\tilde{t} = t$. Since $\pi^m = \frac{e^{-rt}}{1-e^{-rt}} \{[\theta_c + \theta_f v(t)\beta(t)]S - F\}$ and $\beta(t) = \frac{1-e^{-rt}}{r}$, we have

$$\begin{aligned} e^{-rt}\theta_f v'(t)\beta(t)S - \frac{re^{-rt}}{1-e^{-rt}} \{[\theta_c + \theta_f v(t)\beta(t)]S - F\} &= 0 \\ \theta_f v'(t)\beta(t)S - \frac{1}{\beta(t)} \{[\theta_c + \theta_f v(t)\beta(t)]S - F\} &= 0 \\ \theta_f v'(t)[\beta(t)]^2 S - [\theta_c + \theta_f v(t)\beta(t)]S + F &= 0. \end{aligned}$$

Rearranging, we have

$$v(t)\beta(t) - v'(t)[\beta(t)]^2 = \frac{F - \theta_c S}{\theta_f S}.$$

By assumption, the left hand side is strictly increasing from zero to infinity. Since the right hand side is a positive constant, there exists a unique t that solves the equation. Therefore, an equilibrium exists and is unique. \square

Proof of Proposition 1.1. When the firm can commit to future product release dates, its optimal product development time is given by the optimization

$$\begin{aligned} \pi^m &= \max_t e^{-rt} [pS - F + \pi^m] \\ s.t. \quad p &= \theta_c + \theta_f v(t)\beta(t) \end{aligned}$$

Compared with Equation 3, the price changes from $\theta_c + \theta_f v(t)\beta(\tilde{t})$ to $\theta_c + \theta_f v(t)\beta(t)$, because consumers can observe the entire product cycle for the current product.

The first order condition with respect to t is

$$e^{-rt}\theta_f v'(t)\beta(t)S + e^{-rt}\theta_f v(t)\beta'(t)S - re^{-rt} \{[\theta_c + \theta_f v(t)\beta(t)]S - F + \pi^m\} = 0.$$

Comparing this with Equation 4, the ability to commit to future product release dates increases the marginal benefit of raising t by $e^{-rt}\theta_f v(t)\beta'(t)S > 0$. Since the marginal cost doesn't change, the firm's optimal product development time increases when it can commit to future product release dates. \square

Proof of Proposition 1.2. Consumer's willingness to pay is $\theta_c + \theta_f v(t)\beta(t)$, where the equilibrium product development time t is implicitly determined by Equation 5.

By the implicit function theorem, we have

$$\begin{aligned} \frac{dt}{dS} &= -\frac{-\frac{\partial}{\partial S} \frac{F-\theta_c S}{\theta_f S}}{\frac{\partial}{\partial t} \{v(t)\beta(t) - v'(t)[\beta(t)]^2\}} \\ &= \begin{cases} -\frac{\frac{F}{\theta_f S^2}}{\frac{\partial}{\partial t} \{v(t)\beta(t) - v'(t)[\beta(t)]^2\}}, & S \leq 1 \\ 0, & S > 1 \end{cases} \end{aligned}$$

By assumption, $v(t)\beta(t) - v'(t)[\beta(t)]^2$ is increasing, so $\frac{\partial}{\partial t} \{v(t)\beta(t) - v'(t)[\beta(t)]^2\} > 0$. Since $\frac{F}{\theta_f S^2} > 0$, we have

$$\frac{dt}{dS} \begin{cases} < 0, & S \leq 1 \\ = 0, & S > 1 \end{cases}$$

Therefore,

$$\begin{aligned} \frac{d}{dS} [\theta_c + \theta_f v(t)\beta(t)] &= \frac{d}{dt} [\theta_c + \theta_f v(t)\beta(t)] \frac{dt}{dS} \\ &\begin{cases} < 0, & S \leq 1 \\ = 0, & S > 1 \end{cases} \end{aligned}$$

So consumers' willingness to pay is decreasing in S . □

Proof of Proposition 1.3. Given S , the equilibrium profit can be rewritten as

$$\begin{aligned} \pi^m(S) &= \frac{e^{-rt}}{1 - e^{-rt}} \{[\theta_c + \theta_f v(t)\beta(t)]S - F\} \\ &= \frac{1 - r\beta(t)}{r} \theta_f v'(t)\beta(t)S \\ &= \frac{1 - r\beta(t)}{r} \theta_f v'(t)\beta(t) \frac{F}{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 + \theta_c} \\ &= \frac{\theta_f F}{r} \frac{v'(t)\beta(t)[1 - r\beta(t)]}{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 + \theta_c} \end{aligned}$$

where the second and third steps use the first order condition with respect to t . Thus, we express profit as a function of the equilibrium $t(S)$. Optimizing over S is equivalent to optimizing over t instead.

Notice that $\beta'(t) = e^{-rt}$. Taking derivative with respect to t , we have

$$\begin{aligned} \frac{\partial \pi^m(t)}{\partial t} &= \frac{\theta_f F}{r} \left\{ \frac{\{v''(t)\beta(t)[1 - r\beta(t)] + v'(t)(1 - 2r\beta(t))e^{-rt}\} \{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 + \theta_c\}}{\{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 + \theta_c\}^2} \right. \\ &\quad \left. - \frac{v'(t)\beta(t)[1 - r\beta(t)]\theta_f \{v'(t)\beta(t) + v(t)e^{-rt} - v''(t)[\beta(t)]^2 + 2v'(t)\beta(t)e^{-rt}\}}{\{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 + \theta_c\}^2} \right\} \\ &\propto \{v''(t)\beta(t)[1 - r\beta(t)] + v'(t)(1 - 2r\beta(t))e^{-rt}\} \{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 + \theta_c\} \\ &\quad - v'(t)\beta(t)[1 - r\beta(t)]\theta_f \{v'(t)\beta(t) + v(t)e^{-rt} - v''(t)[\beta(t)]^2 + 2v'(t)\beta(t)e^{-rt}\} \end{aligned}$$

As $t \rightarrow 0$,

$$\begin{aligned} \frac{\partial \pi^m(t)}{\partial t} &\propto \{v''(0)\beta(0)[1 - r\beta(0)] + v'(0)(1 - 2r\beta(0))\}\theta_c - 0 \\ &= v'(0)\theta_c > 0 \end{aligned}$$

Thus, $\frac{\partial \pi^m(t)}{\partial t} > 0$ when the equilibrium t is small, and the firm will have an incentive to increase the equilibrium t . Since $\frac{dt(S)}{dS} \leq 0$, it is optimal to decrease S .

Finally, notice that $t \rightarrow 0$ if $\frac{F - \theta_c S}{\theta_f S} = \frac{F}{\theta_f S} - \frac{\theta_c}{\theta_f}$ is small. \square

1.6.3 Proofs for the Extensions

1.6.4 Premium pricing

Suppose consumers' collection value, θ_c , is independent and identically distributed according to the cumulative distribution function $G(\cdot)$. Given the current product's development time t and consumer's anticipation of the next product's development time \tilde{t} , the current product has quality $v(t)$ and will remain as the newest generation for \tilde{t} . Thus, profit of the current generation is given by

$$H(p; t, \tilde{t}) = p \times Pr[\theta_c + \theta_f v(t)\beta(\tilde{t}) \geq p] = p\{1 - G[p - \theta_f v(t)\beta(\tilde{t})]\}.$$

To obtain a unique equilibrium product development time, we impose the assumption that $\{1 - \theta_f g[p - \theta_f v(t)\beta(t)]v'(t)[\beta(t)]^2\}H(p; t, t)$ is increasing in t . For example, if θ_c is uniformly distributed, this assumption holds trivially. Taking the first order condition with respect to p , we have

$$\{1 - G[p - \theta_f v(t)\beta(\tilde{t})]\}S + p \times g[p - \theta_f v(t)\beta(\tilde{t})]S = 0.$$

Suppose the distribution of θ_c has increasing hazard rate, i.e., $\frac{g(\theta)}{1-G(\theta)}$ is increasing in θ . Then, the derivative of the profit function is first increasing and then decreasing. So, the profit of each period has a unique maximizer which solves the equation $p^m = \frac{1-G[p^m - \theta_f v(t)\beta(\tilde{t})]}{g[p^m - \theta_f v(t)\beta(\tilde{t})]}$. This is the equilibrium price without commitment.

Proof of Proposition 1.4. Recall that $p^m = \arg \max_p H(p)$. Raising the price to $p > p^m$ decreases the profit obtained in each generation. So, we can equivalently model a commitment to $p > p^m$ as a commitment to profiting $\alpha H(p)$ from each product release, where $\alpha < 1$.

Given any α , the firm solves

$$\pi^m(\alpha) = \max_t e^{-rt} \{\alpha H(p; t, \tilde{t}) - F + \pi^m(\alpha)\}.$$

Taking the first order condition with respect to t , we have

$$e^{-rt} \alpha H(p; t, \tilde{t}) \frac{\partial H}{\partial t} - r e^{-rt} \{\alpha H(p; t, \tilde{t}) - F + \pi^m\} = 0.$$

In equilibrium, consumers hold correct beliefs, i.e., $\tilde{t} = t$. Since $\pi^m = \frac{e^{-rt}}{1-e^{-rt}} \{\alpha H(p; t, \tilde{t}) - F\}$ and $H(p; t, \tilde{t}) = p\{1 - G[p - \theta_f v(t)\beta(\tilde{t})]\}S$, we have

$$e^{-rt} \alpha H(p; t, t) \theta_f g[p - \theta_f v(t)\beta(t)] v'(t) \beta(t) S - \frac{r e^{-rt}}{1 - e^{-rt}} \{\alpha H(p; t, t) - F\} = 0.$$

Substituting in $\beta(t) = \frac{1-e^{-rt}}{r}$, we have

$$\begin{aligned} \alpha H(p; t, t) \theta_f g[p - \theta_f v(t)\beta(t)] v'(t) \beta(t) - \frac{1}{\beta(t)} \{\alpha H(p; t, t) - F\} &= 0 \\ \alpha \theta_f g[p - \theta_f v(t)\beta(t)] v'(t) [\beta(t)]^2 H(p; t, t) - \alpha H(p; t, t) + F &= 0 \\ H(p; t, t) - \theta_f g[p - \theta_f v(t)\beta(t)] v'(t) [\beta(t)]^2 H(p; t, t) &= \frac{F}{\alpha}. \end{aligned}$$

By assumption, there exists a unique equilibrium product development time $t(\alpha)$. Since the left hand side of the above is increasing in t , we have that $t(\alpha)$ is a decreasing function of α . That is, as the price rise higher above p^m , the product development time increases. This proves the first part of the proposition.

Consumers' willingness to pay is given by $\theta_c + \theta_f v(t)\beta(\tilde{t})$. Consumers hold correct beliefs in equilibrium, so $\tilde{t} = t$. Since $v'(t) > 0$ and $\beta'(t) = e^{-rt} > 0$, we have that

consumers' willingness to pay is decreasing in α . In other words, as α decreases, all consumers' willingness to pay increases, and the demand curve shifts out. This proves the second part of the proposition.

At the beginning of the game, the firm solves $\max_{\alpha} \pi^m(\alpha)$. Given α , the equilibrium profit can be rewritten as

$$\begin{aligned} \pi^m(\alpha) &= \frac{e^{-rt}}{1 - e^{-rt}} \{\alpha H(p; t, t) - F\} \\ &= \frac{e^{-rt}}{1 - e^{-rt}} \frac{1 - e^{-rt}}{re^{-rt}} \alpha H(p; t, t) \theta_f g[p - \theta_f v(t) \beta(t)] v'(t) \beta(t) \\ &= \frac{\theta_f}{r} H(p; t, t) g[p - \theta_f v(t) \beta(t)] v'(t) \beta(t) \frac{F}{H(p; t, t) - \theta_f g[p - \theta_f v(t) \beta(t)] v'(t) [\beta(t)]^2 H(p; t, t)} \\ &= \frac{\theta_f F}{r} \frac{g[p - \theta_f v(t) \beta(t)] v'(t) \beta(t)}{1 - \theta_f g[p - \theta_f v(t) \beta(t)] v'(t) [\beta(t)]^2} \end{aligned}$$

where the second and third steps use the first order condition with respect to t . Thus, we express profit as a function of the equilibrium $t(\alpha)$. Optimizing over α is equivalent to optimizing over t instead.

Notice that $\beta'(t) = e^{-rt}$. Taking derivative with respect to t , we have

$$\begin{aligned} \frac{\partial \pi^m(t)}{\partial t} &= \frac{\theta_f F}{r} \left\{ \frac{-g'[p - \theta_f v(t) \beta(t)] \theta_f [v'(t) \beta(t)]^2 + g[p - \theta_f v(t) \beta(t)] [v''(t) \beta(t) + v'(t) e^{-rt}]}{1 - \theta_f g[p - \theta_f v(t) \beta(t)] v'(t) [\beta(t)]^2} \right. \\ &\quad - \frac{g[p - \theta_f v(t) \beta(t)] v'(t) \beta(t)}{\{1 - \theta_f g[p - \theta_f v(t) \beta(t)] v'(t) [\beta(t)]^2\}^2} \\ &\quad \times \{\theta_f g'[p - \theta_f v(t) \beta(t)] [v'(t) \beta(t) + v(t) e^{-rt}] v'(t) [\beta(t)]^2 \\ &\quad \left. + \theta_f g'[p - \theta_f v(t) \beta(t)] \{v''(t) [\beta(t)]^2 + 2v'(t) \beta(t) e^{-rt}\}\} \right\} \\ &\propto \{-g'[p - \theta_f v(t) \beta(t)] \theta_f [v'(t) \beta(t)]^2 + g[p - \theta_f v(t) \beta(t)] [v''(t) \beta(t) + v'(t) e^{-rt}]\} \\ &\quad \times \{1 - \theta_f g[p - \theta_f v(t) \beta(t)] v'(t) [\beta(t)]^2\} \\ &\quad - g[p - \theta_f v(t) \beta(t)] v'(t) \beta(t) \\ &\quad \times \{\theta_f g'[p - \theta_f v(t) \beta(t)] [v'(t) \beta(t) + v(t) e^{-rt}] v'(t) [\beta(t)]^2 \\ &\quad \left. + \theta_f g'[p - \theta_f v(t) \beta(t)] \{v''(t) [\beta(t)]^2 + 2v'(t) \beta(t) e^{-rt}\}\} \end{aligned}$$

As $t \rightarrow 0$,

$$\begin{aligned} \frac{\partial \pi^m(t)}{\partial t} &\propto \{0 + g(p)[0 + v'(0)]\} \times [1 - 0] - 0 \\ &= g(p)v'(0) > 0 \end{aligned}$$

Thus, $\frac{\partial \pi^m(t)}{\partial t} > 0$ when the equilibrium t is small, and the firm will have an incentive to increase the equilibrium t . Since $\frac{dt(\alpha)}{d\alpha} \leq 0$, it is optimal to decrease α . Thus, when θ_f is large, or θ_c is large, or F is small, we have $t \rightarrow 0$, and it is optimal to decrease α by increasing p . This proves the third part of the proposition. \square

1.6.5 A rationale to low retail prices and high secondary market prices

Proof of Proposition 1.5. Following the same setup as in the previous subsection, notice that a lowering supply is equivalent to selecting a smaller α and obtaining $\alpha H(p; t, \tilde{t})$ profit at each product release date. By the proof of Proposition 1.4, the firm optimally lowers α when θ_f is large, or θ_c is large, or F is small. This proves the first part of the proposition.

At the price p^m , selling limited editions means that the firm supplies less than $\{1 - G[p^m - \theta_f v(t)\beta(t)]\}$. Suppose the firm supplies $\alpha\{1 - G[p^m - \theta_f v(t)\beta(t)]\}$, where $\alpha < 1$. Then, the secondary market clearing price p_2 solves $1 - G(p_2) = \alpha\{1 - G[p^m - \theta_f v(t)\beta(t)]\}$. Since $G(\cdot)$ is increasing, we have $p_2 > p^m$, and

$$p_2 = G^{-1}\{1 - \alpha + \alpha G[p^m - \theta_f v(t)\beta(t)]\}.$$

Since reselling is costless, consumers with $\theta_c < p_2 - \theta_f v(t)\beta(t)$ attempts to buy at retail price, and if successfully purchased, will sell on the secondary market to make profit $p_2 - p^m > 0$. Consumers with $\theta_c \geq p_2 - \theta_f v(t)\beta(t)$ also attempt to buy at retail price. If these consumers successfully purchased from the firm, they will consume the product. Otherwise, these consumers will buy from the secondary market at price p_2 and consume. Therefore, all consumers attempt to buy at retail price. \square

1.6.6 The firm with partial commitment power

Suppose a fraction $\gamma \in (0, 1]$ of the consumers do not believe that the firm can commit to future product release dates. For simplicity, the consumer types are assumed to be publicly known, so the firm can charge different prices to the two consumer groups. Alternatively, we can assume that γ is large enough, so that it is not profitable to price discriminate.

Proof of Proposition 1.6. Given S , the firm's optimization can be re-written as

$$\pi^m = \max_t e^{-rt} \{ \theta_c S + \theta_f v(t) [\gamma \beta(\tilde{t}) + (1 - \gamma) \beta(t)] S - F + \pi^m \}.$$

The first order condition with respect to t is

$$\begin{aligned} e^{-rt} \theta_f v'(t) [\gamma \beta(\tilde{t}) + (1 - \gamma) \beta(t)] S + e^{-rt} \theta_f v(t) (1 - \gamma) \beta'(t) S \\ - r e^{-rt} \{ \theta_c S + \theta_f v(t) [\gamma \beta(\tilde{t}) + (1 - \gamma) \beta(t)] S - F + \pi^m \} = 0 \end{aligned}$$

In equilibrium, consumers hold correct beliefs, i.e., $\tilde{t} = t$. Since $\pi^m = \frac{e^{-rt}}{1 - e^{-rt}} \{ [\theta_c + \theta_f v(t) \beta(t)] S - F \}$ and $\beta(t) = \frac{1 - e^{-rt}}{r}$, we have

$$\begin{aligned} e^{-rt} \theta_f [v'(t) \beta(t) + v(t) (1 - \gamma) \beta'(t)] S - \frac{r e^{-rt}}{1 - e^{-rt}} \{ [\theta_c + \theta_f v(t) \beta(t)] S - F \} = 0 \\ \theta_f [v'(t) \beta(t) + v(t) (1 - \gamma) \beta'(t)] S - \frac{1}{\beta(t)} \{ [\theta_c + \theta_f v(t) \beta(t)] S - F \} = 0 \\ \theta_f v'(t) [\beta(t)]^2 S + \theta_f v(t) (1 - \gamma) S \beta'(t) \beta(t) - [\theta_c + \theta_f v(t) \beta(t)] S + F = 0. \end{aligned}$$

Rearranging, we have

$$v(t) \beta(t) - v'(t) [\beta(t)]^2 - v(t) (1 - \gamma) \beta'(t) \beta(t) = \frac{F - \theta_c S}{\theta_f S}.$$

□

This first order condition is different from the main model. Assume the left hand side is strictly increasing in t . (A sufficient condition is that $v(t)$ is sufficiently concave.) Since the right hand side is constant in t , there is a unique solution to the first order condition.

By assumption $v(t) \beta(t) - v'(t) [\beta(t)]^2 - v(t) (1 - \gamma) \beta'(t) \beta(t)$ is strictly increasing in t . Thus, given any given supply $S > 0$, an equilibrium exists and is unique. The firm's equilibrium product development time is implicitly given by the unique solution to the transcendental equation

$$v(t) \beta(t) - v'(t) [\beta(t)]^2 - v(t) (1 - \gamma) \beta'(t) \beta(t) = \frac{F - \theta_c S}{\theta_f S}.$$

The ability to commit to future product release dates increases the marginal benefit of raising t by $e^{-rt} \theta_f v(t) \gamma \beta'(t) S > 0$. Since the marginal cost doesn't change, the lack

of commitment reduces t .

Given S , the equilibrium profit can be written as

$$\begin{aligned}
\pi^m &= \frac{e^{-rt}}{1 - e^{-rt}} \{[\theta_c + \theta_f v(t)\beta(t)]S - F\} \\
&= \frac{1 - r\beta(t)}{r} \theta_f [v'(t)\beta(t) + v(t)(1 - \gamma)\beta'(t)]S \\
&= \frac{1 - r\beta(t)}{r} \theta_f [v'(t)\beta(t) + v(t)(1 - \gamma)\beta'(t)]S \\
&\quad \times \frac{F}{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 - \theta_f v(t)(1 - \gamma)\beta'(t)\beta(t) + \theta_c} \\
&= \frac{\theta_f F S}{r} \frac{[v'(t)\beta(t) + v(t)(1 - \gamma)\beta'(t)][1 - r\beta(t)]}{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 - \theta_f v(t)(1 - \gamma)\beta'(t)\beta(t) + \theta_c}
\end{aligned}$$

where the second and third steps use the first order condition with respect to t . Thus, we express profit as a function of the equilibrium $t(S)$. Optimizing over S is equivalent to optimizing over the equilibrium product development time.

Notice that $\beta'(t) = e^{-rt}$. Taking derivative with respect to t , we have

$$\begin{aligned}
\frac{\partial \pi}{\partial t} &\propto \frac{[v''(t)\beta(t) + v'(t)e^{-rt} + v'(t)(1 - \gamma)e^{-rt} - rv(t)(1 - \gamma)e^{-rt}][1 - r\beta(t)]}{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 - \theta_f v(t)(1 - \gamma)\beta'(t)\beta(t) + \theta_c} \\
&\quad - \frac{[v'(t)\beta(t) + v(t)(1 - \gamma)\beta'(t)]re^{-rt}}{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 - \theta_f v(t)(1 - \gamma)\beta'(t)\beta(t) + \theta_c} \\
&\quad - \frac{[v'(t)\beta(t) + v(t)(1 - \gamma)\beta'(t)]\{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 - \theta_f v(t)(1 - \gamma)\beta'(t)\beta(t) + \theta_c\}'}{\{\theta_f v(t)\beta(t) - \theta_f v'(t)[\beta(t)]^2 - \theta_f v(t)(1 - \gamma)\beta'(t)\beta(t) + \theta_c\}^2}
\end{aligned}$$

As the equilibrium $t \rightarrow 0$,

$$\begin{aligned}
\frac{\partial \pi}{\partial t} &\propto \frac{v''(0)\beta(0) + (2 - \gamma)v'(0) - r(1 - \gamma)v(0)}{\theta_c} - \frac{r(1 - \gamma)v(0)}{\theta_c} \\
&\quad - \frac{v(0)(1 - \gamma)[\theta_f v(0) - \theta_f v(0)(1 - \gamma)]}{\theta_c^2} \\
&\propto (2 - \gamma)\theta_c v'(0) - 2r(1 - \gamma)\theta_c v(0) - \theta_f(1 - \gamma)[v(0)]^2 \{1 - (1 - \gamma)\} \\
&= (2 - \gamma)\theta_c v'(0) > 0
\end{aligned}$$

If $\frac{\partial \pi}{\partial t} > 0$, the firm will have an incentive to increase the equilibrium t when it is very small. To finish the proof, notice that $\frac{dt(S)}{dS} < 0$ and the equilibrium $t \rightarrow 0$ if $\frac{F - \theta_c S}{\theta_f S}$ is small and \hat{t} is small.

1.7 Bibliography

- AMALDOSS, W. AND S. JAIN (2005a): “Conspicuous consumption and sophisticated thinking,” *Management science*, 51, 1449–1466.
- (2005b): “Pricing of conspicuous goods: A competitive analysis of social effects,” *Journal of Marketing Research*, 42, 30–42.
- (2008): “Research note - Trading up: A strategic analysis of reference group effects,” *Marketing science*, 27, 932–942.
- (2010): “Reference groups and product line decisions: An experimental investigation of limited editions and product proliferation,” *Management Science*, 56, 621–644.
- (2015): “Branding conspicuous goods: An analysis of the effects of social influence and competition,” *Management Science*, 61, 2064–2079.
- BALACHANDER, S., Y. LIU, AND A. STOCK (2009): “An empirical analysis of scarcity strategies in the automobile industry,” *Management Science*, 55, 1623–1637.
- BALACHANDER, S. AND A. STOCK (2009): “Limited edition products: When and when not to offer them,” *Marketing Science*, 28, 336–355.
- BECKER, G. S. (1991): “A note on restaurant pricing and other examples of social influences on price,” *Journal of political economy*, 99, 1109–1116.
- DEGRABA, P. (1995): “Buying frenzies and seller-induced excess demand,” *The RAND Journal of Economics*, 331–342.
- FISHMAN, A. AND R. ROB (2000): “Product innovation by a durable-good monopoly,” *The RAND journal of Economics*, 237–252.
- KELLY, B. (2021): “Target resumes in-store sales of Pokemon trading cards but not MLB, NFL and NBA cards,” .
- KUKSOV, D. (2007): “Brand value in social interaction,” *Management Science*, 53, 1634–1644.

- MASKIN, E. AND J. TIROLE (1988): “A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles,” *Econometrica: Journal of the Econometric Society*, 571–599.
- PESENDORFER, W. (1995): “Design innovation and fashion cycles,” *The american economic review*, 771–792.
- PLAMBECK, E. AND Q. WANG (2009): “Effects of e-waste regulation on new product introduction,” *Management Science*, 55, 333–347.
- RAO, R. S. AND R. SCHAEFER (2013): “Conspicuous consumption and dynamic pricing,” *Marketing Science*, 32, 786–804.
- STOCK, A. AND S. BALACHANDER (2005): “The making of a “hot product”: A signaling explanation of marketers’ scarcity strategy,” *Management science*, 51, 1181–1192.
- SU, X. (2010): “Optimal pricing with speculators and strategic consumers,” *Management Science*, 56, 25–40.
- THOMANN, L. (2017): “The 10 Most Popular Collectible Items (And How to Store Them),” .
- TIROLE, J. (1988): *The theory of industrial organization*, MIT press.
- U-DOX (2014): *Sneakers: The Complete Collectors’ Guide*, Thames & Hudson.
- VEBLEN, T. AND J. K. GALBRAITH (1973): *The theory of the leisure class*, Houghton Mifflin Boston.
- VERHALLEN, T. M. (1982): “Scarcity and consumer choice behavior,” *Journal of Economic Psychology*, 2, 299–322.
- VERHALLEN, T. M. AND H. S. ROBBEN (1994): “Scarcity and preference: An experiment on unavailability and product evaluation,” *Journal of economic psychology*, 15, 315–331.
- WORCHEL, S., J. LEE, AND A. ADEWOLE (1975): “Effects of supply and demand on ratings of object value.” *Journal of personality and social psychology*, 32, 906.

CHAPTER 2

SEQUENTIAL SEARCH WITH A BUDGET CONSTRAINT: AN APPROXIMATION ALGORITHM

2.1 Introduction

Consider a consumer shopping in a store with multiple products. The consumer knows the price of each product and holds a prior belief on product match values. She can spend time and effort to inspect a product and reveal its match value. The consumer faces a monetary budget constraint, so her total expenditure on purchases cannot exceed her total budget. Her objective is to maximize the expected payoff from the shopping journey, which is the total match value of purchases, plus the money leftover, minus the total inspection hassle. Conditioning on the match value of inspected products, the consumer decides whether to continue product inspection. If the consumer continues product inspection, she must choose which product to inspect next; If the consumer stops product inspection, she chooses which inspected products to purchase.

This search problem nests the classic “Pandora’s problem” (Weitzman, 1979) as a special case, and also applies to other environments such as - a marketer optimizing its advertisement expenditure, a firm interviewing and hiring multiple workers, and a retailer choosing the product assortment of its store, etc. (See Section 2.4 for more details.) To be consistent with the sequential search literature, in the rest of the chapter, I refer to each unknown alternative as a closed box. Opening a closed box refers to the

costly search to uncover the hidden box value.

In general, finding the optimal search strategy is computationally difficult, and would likely need to exhaust all possible search orders and purchase combinations. The optimal search strategy must specify the first box to open; then condition on any possible value realization, specify which box to open next; and so on. Furthermore, when the search terminates, the optimal strategy should also specify which open boxes to purchase. So, there is no reason to expect to be able to even write down such a strategy concisely.

This chapter presents a Greedy Search Algorithm (GSA) and shows that it well approximates optimality when the ratio of the largest box price over the total budget is small. In the GSA, each box is assigned a ratio index, which is equal to its (option) value over its price for an open (closed) box. Open boxes are placed into a “knapsack” (which specifies how each dollar of the budget is tentatively allocated) in descending ratio index order. In each iteration, the algorithm considers the closed box j with the highest ratio index. The agent opens this box if she is, on expectation, willing to reallocate the last p_j dollars in the knapsack to purchase j . The knapsack is updated after each box opening. The search terminates when the agent is unwilling to open the next closed box or if all boxes are open. When the search terminates, an open box is purchased if it is entirely included in the knapsack. The GSA is computationally simple in the sense that, when there are n boxes in total, it computes at most $2n$ box indices and sorts boxes no more than n times.

2.1.1 Related Literature

This chapter belongs to the sequential search literature. Our search problem nests the classic “Pandora’s problem” (Weitzman, 1979) as a special case - Weitzman (1979) assumes that all boxes have prices equal to the budget constraint, so the agent can only purchase at most one box. Our GSA degenerates to the optimal algorithm in the Pandora’s problem. Our chapter is also related to Olszewski and Weber (2015a) and Olszewski and Weber (2015b). They consider an agent without a budget constraint, but allow for multi-unit demand and non-additive preferences. Olszewski and Weber (2015b) show that a “generalized Pandora’s rule” is optimal in a special case where the agent has additive preferences and buys a fixed number of boxes. Our search problem nests this special case of Olszewski and Weber (2015b), and our GSA degenerates to their “generalized Pandora’s rule” in this special case.

A few earlier works also analyze search problems with budget constraints. But they assume that all boxes have identical value distributions and search costs, and the search cost is monetary and enters the budget constraint (Manning and Morgan, 1982; Veendorp, 1984; Aharon and Veendorp, 1983; Manning and Manning, 1997). Manning and Morgan (1982) and Manning and Manning (1997) consider a fixed-sample-size search where the agent samples a number of alternatives at once. Aharon and Veendorp (1983) and Veendorp (1984) study sequential search and show that the optimal policy does not have a reservation price property. In contrast, this chapter studies the case where boxes are ex-ante heterogeneous and search costs are non-monetary hassles.

The remaining parts of this chapter are organized as follows. In Section 2.2, I set up the search problem. I formally state the GSA in Subsection 2.3.1 and analyze its asymptotic properties in Subsection 2.3.2. Subsection 2.3.3 provides conditions under which the GSA is optimal. Section 2.4 discusses some potential applications. Section 2.5 concludes the chapter.

2.2 The Search Problem

An agent faces a search problem $\mathcal{S} = \{\mathbb{B}, B\}$, which consists of a purchase budget, $B > 0$, and a set of finitely many unknown alternatives, \mathbb{B} . To be consistent with the literature, I refer to each unknown alternative as a closed box. Each closed box $i \in \mathbb{B}$ has a known price $p_i \in (0, \bar{p}]$ and an (independently distributed) unknown value denoted by the random variable X_i . The value distribution F_i is known by the agent ex-ante. At a search cost of $c_i \geq 0$, the agent can open box i to reveal the value realization x_i . Let O be the set of open boxes and P be the set of boxes that the agent purchases. The total purchase expenditure cannot exceed the agent's budget B , and the agent must open a box before purchasing, so $\sum_{i \in P} p_i \leq B$ and $P \subset O$. The agent cannot search or purchase a fraction of a box. Conditioning on the value realizations of open boxes, the agent chooses whether to continue or terminate searching. If the agent continues searching, she must decide which closed box to open next; If the agent terminates the search, she must decide which open box(es) to purchase. The objective is to maximize the expected payoff

$$\mathbb{E}[\sum_{i \in P} X_i + (B - \sum_{i \in P} p_i) - \sum_{i \in O} c_i]. \quad (2.1)$$

Here, the first term is the total consumption utility, the second term is the value of the leftover budget, and the third term is the total search cost.

Without loss of generality, I assume $\bar{p} \leq B$ to rule out trivially unaffordable boxes. For simplicity of stating the algorithm, I also assume that p_i and B can take any positive values. Nonetheless, a similar version of my results also holds if there exists a smallest currency (e.g. one cent), so that $p_i, B \in \mathbb{N}$. It is also worth emphasizing that only the money spent on box purchases is subject to the budget constraint, not the hassle cost of opening boxes. Applications of this search problem are discussed in Section 2.4.

2.3 The Greedy Search Algorithm and its Properties

2.3.1 The Greedy Search Algorithm (GSA)

In this subsection, I present the Greedy Search Algorithm (GSA). First, I define the *ratio index*, which is a real number that is assigned to each box. This index can be intuitively understood as the “bang-per-buck” for purchasing a box. The ratio index is simple to compute in the sense that the ratio index of each box is independent of the status and value of other boxes.

Definition 2.1. The *ratio index* of box $i \in \mathbb{B}$ is $r_i = \frac{w_i}{p_i}$, where w_i is the Weitzman index (Weitzman, 1979) defined as

$$w_i = \begin{cases} x_i, & i \in O \\ \inf\{y \mid \int_{-\infty}^{\infty} \max\{X_i, y\} dF_i(X_i) - c_i \leq y\}, & i \in O^c \end{cases} \quad (2.2)$$

Before stating the GSA, I need to define a *knapsack function* that, given the set of open boxes, tentatively specifies how to allocate the purchase budget.

Definition 2.2. Relabel the set of open boxes $i \in O$ in descending ratio index order. The *knapsack function* $K_O : [0, B] \rightarrow [1, \infty)$ is a (weakly) decreasing step-wise function such that:

- If $O = \emptyset$, then $K_O(b) = 1$ for all $b \in [0, B]$.

- If $O \neq \emptyset$, then $K_O(b) = \max\{r_1, 1\}$ for all $b \in [0, p_1]$. For any $b \in (p_1, B]$, $K_O(b) = \max\{r_j, 1\}$ if there exists $j \in O$ such that $b \in (\sum_{i \leq j-1} p_i, \sum_{i \leq j} p_i]$, and $K_O(b) = 1$ otherwise.

Loosely speaking, the knapsack function temporarily allocates the first p_1 dollars to box 1, the next p_2 dollars to box 2, and so on, until the next box i has $r_i < 1$ or the budget is depleted. The leftover purchase budget is allocated to the outside option, 1. An example of the knapsack function is provided below.

Example 2.1. Suppose $B = 10$ and $O = \{1, 2\}$. If $(r_1, p_1) = (3, 6)$ and $(r_2, p_2) = (0.5, 6)$, then

$$K_O(b) = \begin{cases} 3, & b \in [0, 6] \\ 1, & b \in (6, 10] \end{cases} \quad (2.3)$$

Instead, if $(r_1, p_1) = (2, 6)$ and $(r_2, p_2) = (2.5, 6)$, then

$$K_O(b) = \begin{cases} 2.5, & b \in [0, 6] \\ 2, & b \in (6, 10] \end{cases} \quad (2.4)$$

◇

The *Greedy Search Algorithm (GSA)* is as follows:

1. Set $O = \emptyset$.
2. Sort and relabel all open boxes $i \in O$ in descending ratio index order. Update the knapsack function K_O .
3. If there are no more closed boxes (i.e., $O^c = \emptyset$), go to Step 4. Otherwise, consider the closed box j with the largest ratio index (i.e., $j = \arg \max_{i \in O^c} r_i$).
 - If $r_j \geq K_O(B - p_j)$, open box j and go to Step 2.
 - If $r_j < K_O(B - p_j)$, terminate the search and go to Step 4.
4. Starting from the open box $i = 1$, purchase open boxes in order, until all boxes with $r_i > 1$ are purchased or the next box is unaffordable.

Since the total number of boxes is finite, the GSA terminates in finite steps. This algorithm is computationally efficient in the sense that it computes at most $2n$ box

indices and sorts boxes no more than n times. I give a simple example to illustrate the GSA below.

Example 2.2. Suppose the agent faces a search problem $\mathcal{S}_1 = \{\mathbb{B}_1, 10\}$, where the total purchase budget is $B = 10$ and \mathbb{B}_1 contains three boxes with identical search costs, $c_1 = c_2 = c_3 = 1$. Box 1 has price $p_1 = 6$ and takes values 18 or 12, each with a half probability. Boxes 2 and 3 have prices $p_2 = 6$ and $p_3 = 4$, and a deterministic value of 15 and 7, respectively.

The GSA starts with $K_\emptyset \equiv 1$. One can compute $r_1 = \frac{8}{3} > r_2 = \frac{7}{3} > r_3 = \frac{3}{2}$. Since $r_1 > 1 = K_\emptyset(10 - 6)$, the GSA opens box 1 first. (Case 1) If box 1 has value 18, then its index changes to $r_1 = \frac{18}{6} = 3$ and the knapsack becomes Equation 3. Since $r_2 = \frac{7}{3} < 3 = K_{\{1\}}(10 - 6)$, the search terminates and the algorithm purchases only box 1. Since the leftover purchase budget is 4, the agent obtains a payoff of $-1 + 18 + 4 = 21$. (Case 2) If box 1 has value 12, then its index changes to $r_1 = \frac{12}{6} = 2$, and the knapsack $K_{\{1\}}(b)$ takes value 2 for $b \in [0, 6]$ and takes value 1 for $b \in (6, 10]$. Since $r_2 = \frac{7}{3} > 2 = K_{\{1\}}(10 - 6)$, the GSA continues to open box 2. The open box 2 has index $r_2 = \frac{15}{6} = \frac{5}{2}$, so the knapsack becomes Equation 4. Then, the search terminates because $r_3 = \frac{3}{2} < \frac{5}{2} = K_{\{1,2\}}(10 - 4)$. Since $r_2 = \frac{5}{2} > 2 = r_1$, the algorithm purchases only box 2 and the leftover budget is $4 < p_1 = 6$. In this case, the agent obtains a payoff of $-2 + 15 + 4 = 17$. Each of the two cases happens with a half probability, so the expected utility of using the GSA is $U^{GSA} = \frac{21+17}{2} = 19$.

In this simple example, is not hard to see that the optimal strategy is to first open box 1. If $x_1 = 18$, then open box 3 and purchase boxes 1 and 3; Otherwise, open and purchase boxes 2 and 3. This leads to the maximum expected payoff of $U^* = -1 + \frac{1}{2}(18 - 1 + 7) + \frac{1}{2}(-1 + 15 - 1 + 7) = 21$. So the GSA is generally not optimal. However, the next subsection shows that the GSA is near optimal when all boxes are small relative to the total budget. \diamond

2.3.2 Main result: asymptotic optimality of the GSA

For a given search problem \mathcal{S} , denote $U^*(\mathcal{S})$ as the expected utility obtained from using the optimal strategy, and $U^{GSA}(\mathcal{S})$ as the expected utility from using the GSA. Define the *payoff loss* from using the GSA as $L(\mathcal{S}) = \frac{U^*(\mathcal{S}) - U^{GSA}(\mathcal{S})}{U^*(\mathcal{S})}$. The main result of this section is as follows.

Theorem 2.1. *For any search problem $\mathcal{S} = \{B, \mathbb{B}\}$, the payoff loss $L(\mathcal{S}) \leq 1 - (1 - \frac{\bar{p}}{B})^2$, where $\bar{p} = \max_{i \in \mathbb{B}} p_i$.*

I sketch the proof of Theorem 2.1 here and present the full proof in the Appendix. The loss bound is proven without finding the optimal search rule for the general problem. This is done in four steps. Fix an arbitrary search problem $\mathcal{S} = \{B, \mathbb{B}\}$. First, I construct a “modified optimal algorithm” - take the search rule of the optimal algorithm and additionally impose the GSA’s search termination rule (i.e., let $j = \arg \max_{i \in O^c} r_i$, terminate the search whenever $r_j < K_O(B - p_j)$) and purchase strategy (i.e., when search terminates, purchase open boxes in decreasing ratio index order). Denote the expected utility obtained from this modified optimal algorithm as $U_{mod}^*(\mathcal{S})$. Lemma 2.1 uses the monotonic decreasing property of the knapsack function to show that $\frac{U_{mod}^*(\mathcal{S})}{U^*(\mathcal{S})} \geq 1 - \frac{\bar{p}}{B}$. Second, I construct a “compensated GSA” - follow the GSA search rule, slightly compensate the knapsack to $\tilde{K}_O(b) \leq \max\{K_O(b), \max_{j \in O^c} r_j\}$ when the search terminates, and take the entire compensated knapsack as the payoff. Denote the expected utility obtained from this compensated GSA as $U_{comp}^{GSA}(\mathcal{S})$. Lemma 2.2 uses the GSA’s search termination rule and the monotonic decreasing property of the compensated knapsack function to show that $\frac{U_{comp}^{GSA}(\mathcal{S})}{U_{mod}^*(\mathcal{S})} \geq 1 - \frac{\bar{p}}{B}$. Third, I prove $U_{comp}^{GSA}(\mathcal{S}) \geq U_{mod}^*(\mathcal{S})$ in Lemma 2.3. Similar to Weitzman (1979), this step is done via induction on the number of closed boxes. Lastly, I combine all previous steps to obtain $\frac{U_{comp}^{GSA}(\mathcal{S})}{U^*(\mathcal{S})} = \frac{U_{mod}^*(\mathcal{S})}{U^*(\mathcal{S})} \times \frac{U_{comp}^{GSA}(\mathcal{S})}{U_{mod}^*(\mathcal{S})} \times \frac{U_{comp}^{GSA}(\mathcal{S})}{U_{comp}^{GSA}(\mathcal{S})} \geq (1 - \frac{\bar{p}}{B})^2$. Therefore, $L(\mathcal{S}) = 1 - \frac{U_{comp}^{GSA}(\mathcal{S})}{U^*(\mathcal{S})} \leq 1 - (1 - \frac{\bar{p}}{B})^2$.

There are simple examples where $L(\mathcal{S})$ is arbitrarily close to $1 - \frac{\bar{p}}{B}$. So the general upper bound on $L(\mathcal{S})$ is at least $1 - \frac{\bar{p}}{B}$. However, I cannot find any examples of a search problem \mathcal{S} , where $L(\mathcal{S}) \in [1 - \frac{\bar{p}}{B}, 1 - (1 - \frac{\bar{p}}{B})^2]$. Thus, I conjecture that $L(\mathcal{S}) < 1 - \frac{\bar{p}}{B}$, $\forall \mathcal{S}$. Nonetheless, $1 - (1 - \frac{\bar{p}}{B})^2$ and $1 - \frac{\bar{p}}{B}$ are of the same order because they both shrink linearly as $\frac{\bar{p}}{B}$ approaches zero.

To illustrate the use of Theorem 2.1, consider a sequence of search problems $\{\mathcal{S}_n\}_{n=1}^\infty$. The search problem $\mathcal{S}_1 = \{\mathbb{B}_1, B\}$ has a purchase budget B and a set of closed boxes \mathbb{B}_1 . Each closed box $i \in \mathbb{B}_1$ has a price $p_i^1 \in (0, \bar{p}]$ and a (independently distributed) value X_i^1 , and can be open at cost c_i^1 . For any $n \geq 2$, the search problem $\mathcal{S}_n = \{\mathbb{B}_n, n \times B\}$ has a purchase budget $n \times B$ and a total number of $n \times |\mathbb{B}_1|$ closed boxes in \mathbb{B}_n . For each box $j \in \mathbb{B}_n$, let $k \in [1, |\mathbb{B}_1|]$ be the unique integer such that $(k - 1)n < j \leq kn$. Box j has a price $p_j^n = p_k^1$ and a (independently distributed) random value $X_j^n = X_k^1$, and can be open at cost $c_j^n = c_k^1$. That is, \mathbb{B}_n can be obtained by multiplying each box




Figure 2.1: Convergence speed of the GSA.

in \mathbb{B}_1 into n independent and identical copies. Theorem 2.1 implies that $L(\mathcal{S}_n)$ shrinks to 0. Furthermore, $L(\mathcal{S}_n)$ is of the same order as $\frac{1}{n}$, so the convergence speed is linear.

Corollary 2.1. *For any sequence of search problems $\{\mathcal{S}_n\}_{n=1}^{\infty}$, the payoff loss $L(\mathcal{S}_n) = O(\frac{1}{n})$ and $\lim_{n \rightarrow \infty} L(\mathcal{S}_n) = 0$.*

Example 2.3. Consider a sequence of search problems $\{\mathcal{S}_n\}_{n=1}^{\infty}$ with $\mathcal{S}_1 = \{\mathbb{B}_1, 10\}$ given as in Example 2 and $\mathcal{S}_n = \{\mathbb{B}_n, 10 \times n\}$ for $n \geq 2$, where \mathbb{B}_n multiplies each box in \mathbb{B}_1 into n independent and identical copies. The following figure plots the loss from applying the GSA, $L(\mathcal{S}_n)$, for $1 \leq n \leq 11$.

◇

2.3.3 When is the GSA optimal?

The GSA can have a positive payoff loss because it may open boxes in a suboptimal order, terminate the search too early, or purchase a suboptimal combination of open boxes, leading to an inefficiently large leftover purchase budget. However, the following result shows that the GSA is optimal in some special cases.

Theorem 2.2. *The GSA is optimal for a search problem $\mathcal{S} = \{\mathbb{B}, B\}$ if one of the following conditions holds.*

- $\sum_{i \in \mathbb{B}, r_i \geq 1} p_i \leq B$.

- [Weitzman, 1979] $p_i = B$ for all $i \in \mathbb{B}$.
- [Olszewski and Weber, 2015b] $p_i = p \in (0, B]$ for all $i \in \mathbb{B}$.

Proof. When $\sum_{i \in \mathbb{B}, r_i > 1} p_i \leq B$, it is optimal to open all closed boxes $i \in \mathbb{B}$ with $r_i > 1$ and then purchase all open boxes $j \in O$ with $r_j > 1$. This is because the expected gain from opening each box with $r_i > 1$ is positive, $-c_i + \int_{-\infty}^{\infty} \max\{X_i, p_i\} dF_i(X_i) - p_i > -c_i + \int \max\{X_i, r_i p_i\} dF_i(X_i) - r_i p_i = 0$; and similarly, the expected gain from opening each box with $r_i \leq 1$ is (weakly) negative.

When $p_i = B$ for all $i \in \mathbb{B}$, \mathcal{S} can be restated as the search problem of Weitzman (1979). This is done by adding one additional open box $j = 0$ with $p_0 = x_0 = B$ to capture the outside value of the leftover budget. Then, the agent’s objective becomes maximize $\mathbb{E}[\max_{i \in O \cup \{0\}} X_i - \sum_{i \in O} c_i]$. By Weitzman (1979), the Pandora’s rule is optimal. Since the GSA opens and purchases the same box(es) as the Pandora’s rule, it is also optimal.

When $p_i = p$ for all $i \in \mathbb{B}$, the search problem \mathcal{S} can be restated as a special case (additive preferences for the top $\lfloor \frac{B}{p} \rfloor$ best alternatives) considered in Olszewski and Weber (2015b).¹ To do so, I capture the outside value of leftover budget using $\lfloor \frac{B}{p} \rfloor$ additional open boxes $i \in O'$, each with $p_j = x_j = \lfloor \frac{B}{p} \rfloor$. Then, the agent’s objective transforms to maximize $B - \lfloor \frac{B}{p} \rfloor + \mathbb{E}[\sum_{i \in P} X_i - \sum_{i \in O} c_i]$ with the purchase constraints $P \subset O \cup O'$ and $|P| = \lfloor \frac{B}{p} \rfloor$. Since the agent’s preference is SPR compatible (Olszewski and Weber, 2015b,a), by Theorem 2 of Olszewski and Weber (2015b), the “generalized Pandora’s rule” is optimal. The GSA coincides with the generalized Pandora’s rule by following the same search and purchase rule. Hence, the GSA is also optimal. \square

2.4 Applications

The search problem can be of interest in multiple different environments.

A consumer shopping in a store. A store offers n different products. A consumer is shopping on a budget of B dollars. This may be an actual monetary constraint or a constraint due to mental accounting (e.g., Thaler, 1990). Retail prices are directly observable but product match values are not. The consumer’s prior belief on the match

¹Here, $\lfloor \frac{B}{p} \rfloor$ denotes the integer k that satisfies $k \leq \frac{B}{p} < k + 1$.

value is the distribution induced by X_i . She can reveal the match value x_i by paying an inspection time/effort cost of c_i . The consumer wants to maximize the payoff of this shopping trip.

A marketer optimizing its advertisement expenditure. A marketer faces n advertising options - different keywords, websites, target consumers, billboard locations and sizes, etc. The company allocates B dollars to ads. The marketer knows the price of posting ad i , p_i . But the advertisement conversion rate x_i is hidden and can only be learned through marketing research, which costs time and effort c_i . The marketer must decide how to conduct her marketing research and allocate the advertisement budget.

A firm interviewing and recruiting new workers. A total of n candidates applied to different job openings at a firm. The firm has a budget of B dollars. The salary of each position, p_i , is set and known by the firm. But the ability, x_i , of each applicant i can only be learned through an interview, which costs the human resources time and effort, c_i . The human resources team needs to decide which candidates to interview, in what order, and which candidates to hire.

A purchasing manager selecting the product assortment of its store. The purchasing manager of a store (e.g., Target or Costco warehouse) is in charge of procuring products for retail. He faces a variety of products from n different vendors. The physical store has a total shelf capacity of B . The manager directly observes the physical size of products, p_i , and can learn the retail profit x_i at cost c_i (e.g. cost of conducting consumer research and testing product samples). The products sold in-store cannot exceed the retailer's total shelf capacity, so $\sum_{i \in P} p_i \leq B$. The retailer must decide which products to inspect and which products to offer in its store.

In this final example, the leftover purchase budget (store shelf space) has no value. So the objective function changes from Equation (1) to $\mathbb{E}[\sum_{i \in P} X_i + b_0(B - \sum_{i \in P} p_i) - \sum_{i \in O} c_i]$, where the outside option of “money” is $b_0 = 0$. The same GSA can be applied to this general case with a tweak of the knapsack function. In general, suppose that $b_0 \in \mathbb{R}$. Redefine the knapsack function to temporarily allocate the first p_1 dollars to box 1, the next p_2 dollars to box 2, and so on, until the budget depletes or the next open box i has $r_i < b_0$. All leftover money in the knapsack is allocated to the outside option of value

b_0 . Following the same arguments, it can be shown that Theorem 2.1 still holds in this more general case.

2.5 Conclusion

This chapter studies a sequential search problem where the decision maker is budget constrained. The search problem nests the classic ‘‘Pandora’s problem’’ (Weitzman, 1979) as a special case. I present the computationally efficient Greedy Search Algorithm and show that it is near optimal when the ratio of the largest box price over the total budget is small.

2.6 Appendix

I present the full proof of Theorem 2.1 in this appendix. In the following, I fix any arbitrary search problem $\mathcal{S} = \{B, \mathbb{B}\}$, and abbreviate $U^*(\mathcal{S})$ and $U^{GSA}(\mathcal{S})$ respectively as U^* and U^{GSA} .

First, I construct the ‘‘modified optimal strategy’’ as follows: Consider the search rule of the optimal strategy and additionally impose the GSA’s search termination rule - let $j = \arg \max_{i \in O^c} r_i$, terminate search whenever $r_j < K_O(B - p_j)$. When the search terminates (regardless of whether it is due to the optimal strategy or the GSA’s termination rule), follow the GSA’s purchase strategy - purchase open boxes in descending ratio index order, until either the next open box i has $r_i < 1$ or is unaffordable.

Denote the expected utility obtained from this modified optimal strategy as U_{mod}^* .

Lemma 2.1. *For any search problem \mathcal{S} , $\frac{U_{mod}^*(\mathcal{S})}{U^*(\mathcal{S})} \geq 1 - \frac{\bar{p}}{B}$.*

Proof. Suppose the GSA’s search termination rule does not apply, then the modified optimal strategy opens the same boxes as the optimal strategy. If the search terminates with $\sum_{i \in O, r_i > 1} p_i \leq B$, then both the optimal algorithm and the modified optimal algorithm purchase all open boxes $i \in O$ with $r_i > 1$. If $\sum_{i \in O, r_i > 1} p_i > B$, suppose open boxes are labeled in descending ratio index order. Let $k \in O$ be the box such that $\sum_{i \in O, i \leq k-1} p_i \leq B$ and $\sum_{i \in O, i \leq k} p_i > B$, i.e., the agent can afford all boxes $i = 1, 2, \dots, k-1$, but cannot afford k on top of that. It is easy to see that $U^* \leq \sum_{i \in O, i \leq k-1} p_i r_i + \left(B - \sum_{i \in O, i \leq k-1} p_i \right) r_k$.

Since the modified optimal algorithm purchases open boxes in descending ratio index order, hence $U_{mod}^* = \sum_{i \in O, i \leq k-1} p_i r_i$. Therefore, $\frac{U_{mod}^*}{U^*} \geq \frac{\sum_{i \in O, i \leq k-1} p_i r_i}{\sum_{i \in O, i \leq k-1} p_i r_i + \left(B - \sum_{i \in O, i \leq k-1} p_i \right) r_k} = 1 - \frac{\left(B - \sum_{i \in O, i \leq k-1} p_i \right) r_k}{\sum_{i \in O, i \leq k-1} p_i r_i + \left(B - \sum_{i \in O, i \leq k-1} p_i \right) r_k} \geq 1 - \frac{\left(B - \sum_{i \in O, i \leq k-1} p_i \right) r_k}{\left(\sum_{i \in O, i \leq k-1} p_i \right) r_k + \left(B - \sum_{i \in O, i \leq k-1} p_i \right) r_k} = 1 - \frac{\bar{p}}{B}$, where the third step is because boxes are labeled in descending ratio index order, and the last step is because $B - \sum_{i \in O, i \leq k-1} p_i \leq \bar{p}$.

Now, suppose the GSA's search termination rule applies, and the modified optimal algorithm terminates search because a box $j = \arg \max_{i \in O^c} r_i$ satisfies $r_j < K_O(B - p_j)$. Let $\bar{U}_{mod}^* = \int_0^B \max\{K_O(b), r_j\} db$. First, it is shown that $\bar{U}_{mod}^* \geq U^*$. Intuitively, this is true because each dollar is assigned to a value that is weakly greater than the indices of all closed boxes. Specifically, suppose the agent has the option to terminate the search and obtain \bar{U}_{mod}^* , and consider the benefit of continue searching to open any closed box $i \in O^c$. The expected gain from opening box i is at most $\Delta = -c_i + \int_{-\infty}^{\infty} \int_{B-p_i}^B \max\{\frac{X_i}{p_i}, \max\{K_O(b), r_j\}\} db dF(X_i) - \int_{B-p_i}^B \max\{K_O(b), r_j\} db$. By definition, the Weitzman index $w_i = r_i p_i$ solves $\int_{-\infty}^{\infty} \max\{X_i, w_i\} dF_i(X_i) - c_i - w_i = 0$. Since $\max\{K_O(B-p_i), r_j\} \geq r_j \geq r_i$, hence $\Delta \leq -c_i + \int_{-\infty}^{\infty} \max\{X_i - r_i p_i, 0\} dF_i(X_i) = 0$. Thus, the expected gain from further opening any box $i \in O^c$ is non-positive, and $\bar{U}_{mod}^* \geq U^*$.

Next, let k be the open box such that $\sum_{i \in O, i \leq k-1} p_i \leq B$ and $\sum_{i \in O, i \leq k} p_i > B$. From the purchase strategy of the GSA, $\frac{U_{mod}^*}{U^*} = \frac{\sum_{i \in O, i \leq k-1} p_i r_i}{\int_0^B \max\{K_O(b), r_j\} db} \geq \frac{\int_0^{B-\bar{p}} K_O(b) db}{\int_0^{B-\bar{p}} K_O(b) db + \int_{B-\bar{p}}^B K_O(B-\bar{p}) db} = 1 - \frac{\int_{B-\bar{p}}^B K_O(B-\bar{p}) db}{\int_0^{B-\bar{p}} K_O(b) db + \int_{B-\bar{p}}^B K_O(B-\bar{p}) db} \geq 1 - \frac{\int_{B-\bar{p}}^B K_O(B-\bar{p}) db}{\int_0^{B-\bar{p}} K_O(B-\bar{p}) db + \int_{B-\bar{p}}^B K_O(B-\bar{p}) db} \geq \frac{B-\bar{p}}{B}$, where second step uses $B - \sum_{1 \leq i \leq k} p_i \leq \bar{p}$, and the second and the fourth step use $K_O(b)$ is decreasing. Since $\bar{U}_{mod}^* \geq U^*$, $\frac{U_{mod}^*}{U^*} \geq \frac{U_{mod}^*}{\bar{U}_{mod}^*} \geq 1 - \frac{\bar{p}}{B}$. \square

It is generally not true that $U^{GSA} \geq U_{mod}^*$. The workaround is to construct a compensated GSA, which artificially increases the payoff of using the GSA to $U_{comp}^{GSA} \geq U^{GSA}$, so that (1) when taking into account this compensation, opening boxes in decreasing ratio index order is better than any other search order that also respects the GSA's search termination condition, and (2) the compensation is negligible in the limit as $\frac{\bar{p}}{B} \rightarrow 0$. The first condition ensures that $U_{comp}^{GSA} > U_{mod}^*$, because the later strategy

respects the GSA's search termination condition and does not get the compensation. The second condition guarantees that, U_{comp}^{GSA} is close to U^{GSA} as $\frac{\bar{p}}{B} \rightarrow 0$. Combining the two conditions, U^{GSA} is close to U_{mod}^* when $\frac{\bar{p}}{B}$ is small. Then, we can use $U_{comp}^{GSA} \geq U^{GSA}$ and Lemma 2.1 to show that U^{GSA} is close to U^* when $\frac{\bar{p}}{B}$ is small.

For the purpose of constructing the compensation for U_{comp}^{GSA} , we keep track of two indices for each open box $i \in O$: one is the ratio index obtained from dividing the actual box value by price, $r_i = \frac{x_i}{p_i}$; the other is the old ratio index assigned to box i when this box was closed - denote the old ratio index as r_i^c . I also keep track of the sequence in which boxes were open: for any $i_{l-1}, i_l \in O$, $l-1$ denotes the box that was open right before box l .

I construct a compensation value, and this compensation can be applied to any strategy at the time the search terminates and the agent stops opening boxes. The “*compensation*” C is constructed as follows. When the search terminates, suppose the set of open boxes is O . Consider a set of blocking boxes $\tilde{O} \subset O$, where for any single box $i \in \tilde{O}$, the search would not have terminated if the realization of r_i is replaced by 1 (the outside option of leftover budget). I say that “ O blocks box k ”, if reassigning $r_i = 1$ to the most recent blocking box $i = \arg \max_{i \in \tilde{O}} l$ leads to opening box k next. If O does not block any boxes (i.e., $\tilde{O} = \emptyset$), then (in this case, the GSA would continue opening boxes if $O^c \neq \emptyset$, and) there is no compensation, and $C = 1$ is the outside option of leftover budget. If O blocks box k and $r_k^c \geq K_O(B - p_k)$, then (in this case, the GSA would continue opening boxes, and) there is also no compensation, and $C = 1$. If O blocks box k and $r_k^c < K_O(B - p_k)$, then (in this case, the GSA would also terminate the search, and) the compensation is given by $C = \max\{\min\{\tilde{r}, r_j^c\}, 1\}$, where $j = \arg \max_{i \in O^c} r_i^c$, $\tilde{r} = r_{\tilde{i}}^c$, and box $\tilde{i} = \arg \max_{i \in O} l$ s.t. $r_k^c \geq K_{O \setminus \{\tilde{i}\}}(B - p_k)$ is the most recent open box satisfying $r_k^c \geq K_{O \setminus \{\tilde{i}\}}(B - p_k)$. Note that $k = \arg \max_{i \in O^c} r_i^c$ always holds when using the GSA, but $k \neq j$ is possible with other search strategies.

The “*compensated GSA*” is defined as follows. Use the same search and termination rule as the GSA. When the search terminates, the payoff of the compensated GSA is given by the compensated knapsack $\tilde{K}_O(b) = \max\{K_O(b), C\}$, i.e.,

$$U_{comp}^{GSA} = \int_0^B \tilde{K}_O(b) db = \int_0^B \max\{K_O(b), C\} db,$$

where the compensation C is defined as in the previous paragraph.

Note that the compensated GSA is constructed only for the purpose of bounding

the loss from using the GSA. The additional steps to derive the compensation have no impact on the computation complexity of the original GSA.

The compensation increases the bang-per-buck of the knapsack to at least $\min\{\tilde{r}, r_j^c\}$. The idea of the compensation is as follows. Following the GSA's search strategy, if the search terminates and O blocks a box k , then it must be that $k = j = \arg \max_{i \in O^c} r_i^c$. Compensating by at most r_j^c and only when $r_j^c < K_O(B - p_j)$ ensures that the compensation for the GSA is small in every realization. So the payoffs from using the GSA and the compensated GSA are close (see Lemma 2.2). Capping the compensation by \tilde{r} punishes the history of opening boxes with small r_i^c . So, when taking into account this compensation (along with the GSA's search termination constraint), it becomes optimal to follow the GSA's search strategy - open boxes in descending ratio index order (see Lemma 2.3).

Lemma 2.2. *For any search problem \mathcal{S} , $\frac{U^{GSA}(\mathcal{S})}{U_{comp}^{GSA}(\mathcal{S})} \geq 1 - \frac{\bar{p}}{B}$.*

Proof. If all boxes are open, then $C = 1$, $\tilde{K}_O = K_O$, and $\frac{U^{GSA}}{U_{comp}^{GSA}} = 1 > 1 - \frac{\bar{p}}{B}$. I only need to consider the case where the search terminates because a closed box $j \in O^c$ is blocked by an open box $\tilde{i} \in O$, and the compensation $C > 1$. Since search terminated, $j = \arg \max_{i \in O^c} r_i$, $r_j^c < K_O(B - p_j)$. Since boxes are open in descending index ratio order, $r_{\tilde{i}}^c > r_j^c$. So the compensated knapsack is given by $\tilde{K}_O(b) = \max\{r_j^c, K_O(b)\}$, and $\frac{U^{GSA}}{U_{comp}^{GSA}} \geq \frac{\int_0^{B-\bar{p}} K_O(b) db}{\int_0^B \tilde{K}_O(b) db} = \frac{\int_0^{B-\bar{p}} K_O(b) db}{\int_0^{B-\bar{p}} K_O(b) db + \int_{B-\bar{p}}^B \tilde{K}_O(b) db} \geq \frac{\int_0^{B-\bar{p}} K_O(b) db}{\int_0^{B-\bar{p}} K_O(b) db + \int_{B-\bar{p}}^B K_O(B-\bar{p}) db} = 1 - \frac{\int_{B-\bar{p}}^B K_O(B-\bar{p}) db}{\int_0^{B-\bar{p}} K_O(b) db + \int_{B-\bar{p}}^B K_O(B-\bar{p}) db} \geq 1 - \frac{\bar{p}}{B}$. Here, the first step is because running the GSA results in at most \bar{p} dollars leftover, the second step is because $K_O(b) \geq K_O(B - p_j) > r_j^c$ for all $b \in [0, B - \bar{p}]$, the third step uses $\tilde{K}_O(b) = \max\{r_j^c, K_O(b)\} < \max\{K_O(B - p_j), K_O(b)\} \leq K_O(B - \bar{p})$ for $b \in [B - \bar{p}, B]$, and the last step uses $K_O(b) \geq K_O(B - \bar{p})$ for $b \in [0, B - \bar{p}]$. \square

Lemma 2.3. *For any search problem \mathcal{S} , $U_{comp}^{GSA}(\mathcal{S}) \geq U_{mod}^*(\mathcal{S})$.*

Proof. Notice that the modified optimal strategy terminates search when box $j = \arg \max_{i \in O^c} r_i^c$ satisfies $r_j^c < K_O(B - p_j)$, purchases open boxes in descending ratio index order, and obtains utility $U_{mod}^* \leq \int_0^B K_O(b) db \leq \int_0^B \tilde{K}_O(b) db$. Thus, it suffices to show that the compensated GSA yields a higher expected utility than any other strategy that (1) terminates search when $r_j^c = \max_{i \in O^c} r_i^c < K_O(B - p_j)$, (note that this does not

prohibit using other additional search termination conditions, so the modified optimal strategy is allowed), and (2) when search terminates, obtains the compensated knapsack. Similar to Weitzman (1979), this is proven via induction on the number of closed boxes. Consider an arbitrary strategy α satisfying the requirements (1) and (2). Denote the current set of open boxes as O . I prove $U_{comp}^{GSA} \geq U_\alpha$ using induction.

Consider the case with only one closed box j remaining. (Case 1) Suppose $r_j^c \geq K_O(B - p_j)$, so the compensated GSA opens j . If strategy α does not open box j , there is no compensation and $U_\alpha = \int_0^B K_O(b) db$. Following the compensated GSA and open j leads to a higher payoff, because the expected gain from opening box j is $\Delta = \int_{-\infty}^{\infty} \int_{B-p_j}^B \max\{X_j - K_O(b), K_O(b)\} db dF(X_j) - c_j - \int_{B-p_j}^B K_O(b) db$, which is greater than $\int \max\{X_j - r_j^c p_j, r_j^c p_j\} dF(X_j) - c_j - r_j^c p_j = 0$. If strategy α also opens j , then strategy α is identical with the compensated GSA, and $U_{comp}^{MGA} = U_\alpha$. (Case 2) If $r_j^c < K_O(B - p_j)$, then both strategies terminate search and $U_{comp}^{GSA} = U_\alpha$.

Suppose $U_{comp}^{GSA} \geq U_\alpha$ still holds when there are m closed boxes. In the rest of this proof, I use induction and show $U_{comp}^{GSA} \geq U_\alpha$ still holds for the case with $m + 1$ closed boxes.

(Case 1) Suppose box $j = \arg \max_{i \in O^c} r_i^c$ satisfies $r_j^c < K_O(B - p_j)$, and both strategies terminate the search. Then, the two strategies are identical and $U_{comp}^{MGA} = U_\alpha$. (Case 2) Suppose box $j = \arg \max_{i \in O^c} r_i^c$ satisfies $r_j^c \geq K_O(B - p_j)$ and both strategies opens box j , then $U_{comp}^{GSA} \geq U_\alpha$ because of the assumption for m closed boxes. (Case 3) Suppose box $j = \arg \max_{i \in O^c} r_i^c$ satisfies $r_j^c \geq K_O(B - p_j)$ and strategy α terminates search. Hence, $U_{comp}^{GSA} \geq U_\alpha$. This is because opening box j and terminating the search immediately would yield a higher expected utility than U_α , and by induction, after opening box j , continuing with the compensated GSA is better than terminating the search. (Case 4) Suppose box $j = \arg \max_{i \in O^c} r_i^c$ satisfies $r_j^c \geq K_O(B - p_j)$ and strategy α opens a box $k \neq j$. To complete the proof, I need to show $U_{comp}^{GSA} \geq U_\alpha$ in this final case.

Denote h as the second highest ratio index among all closed boxes, i.e., $h = \arg \max_{i \in O^c \setminus \{j\}} r_i^c$. Consider two auxiliary search strategies. The first auxiliary search strategy β opens box k first. After opening box k , proceed with the compensated GSA. Denote the expected payoff obtained from using strategy β as U_β . Since there will be m boxes left after k is open, by induction, $U_\beta \geq U_\alpha$. The second auxiliary search strategy γ runs the compensated GSA for one step, and opens box j first. Then, terminate search if either $r_j > r_h^c$ or $r_k \leq K_{O \cup \{j\}}(B - p_k + 1)$. Otherwise, open box k , and terminate search if $r_k \geq r_h^c$. Then, follow the compensated GSA onward. Denote the expected

payoff obtained from using strategy γ as U_γ . Since there are m boxes left after opening j , by induction, deviating from the compensated GSA, while still constrained by the same termination and purchase rules, gives a weakly lower payoff. So, $U_{comp}^{GSA} \geq U_\gamma$. Next, I prove $U_\gamma \geq U_\beta$, which will imply $U_{comp}^{GSA} \geq U_\gamma \geq U_\beta \geq U_\alpha$ in (Case 4).

I introduce some new notations. Recall that O is the set of open boxes prior to opening i, j, h . Let $b_i = \max\{b \mid K_O(b) \geq r_i^c\}$. Since $r_j^c \geq r_h^c \geq r_k^c$, $b_j \leq b_h \leq b_k$. If $b_j + p_k \leq B - p_j$, then both strategies β and γ open j with certainty. If $b_k + p_j \leq B - p_k$ and $b_h + p_j \leq B - p_h$, then both strategy β and γ open box k with certainty. (Case 4.1) Hence, if $b_j \leq b_k \leq B - p_j - p_k$ and $b_h \leq B - p_j - p_h$, then both box j and k are open using either strategy, and the sequence does not matter. (Case 4.2) If $b_j \leq B - p_j - p_k \leq b_k$ or $b_h > B - p_j - p_h$, then box j always gets open in both strategies, while box k always gets open only in strategy β . So, it is better to follow strategy γ . (Case 4.3) Thus, I only need to focus on the case with $B - p_j - p_k \leq b_j \leq b_k$ and $b_h \leq B - p_j - p_h$.

Using strategy β , the agent first opens box k at cost c_k . If $r_k > r_j^c$, then $K_{O \cup \{k\}}(B - p_j) > r_j$ because of $B - p_j - p_k \leq b_j \leq b_k$, and the search terminates. Denote the payoff of this case as U_1 . If $r_h^c < r_k \leq r_j^c$, then strategy β continues to open j . When $r_j > r_j^c$, then $r_j > r_k > r_h^c$ and the search terminates, the agent discards the unaffordable box k , and obtains a payoff denoted as U_2 ; when $r_h^c < r_j \leq r_j^c$, search also terminates and the agent gets a payoff denoted as U_3 ; when $r_j \leq r_h^c$, then depending on the price p_h , the agent might or might not be open box h , and get a continuation value as V_1 . If $r_k < r_h^c$, then strategy β also continues to open j . When $r_j > r_h^c$, the agent open h because $b_h \leq B - p_j - p_h$, and obtains a continuation value denoted as V_2 ; when $r_j < r_h^c$, the strategy β considers box h , and obtains a continuation value denoted as V_3 . In summary, the total expected payoff from using strategy β can be written as

$$\begin{aligned} U_\beta = & -c_k + P[r_k > r_j^c]U_1 \\ & + P[r_h^c < r_k \leq r_j^c] \left\{ -c_j + P[r_j > r_j^c]U_2 + P[r_h^c < r_j \leq r_j^c]U_3 + P[r_j \leq r_h^c]V_1 \right\} \\ & + P[r_k \leq r_h^c] \left\{ -c_j + P[r_j > r_h^c]V_2 + P[r_j \leq r_h^c]V_3 \right\} \end{aligned}$$

Using strategy γ , the agent first opens box j at cost c_j . If $r_j > r_j^c$ or $r_h^c < r_j \leq r_j^c$, then the search terminates and the agent purchases j . Denote the payoff of these two cases as U_4 and U_5 , respectively. If $r_j \leq r_h^c$, then strategy γ requires opening box k . When $r_k > r_j^c$, then strategy γ terminates search by definition, and denote the payoff in this case as U_6 ; when $r_h^c < r_k \leq r_j^c$, the continuation payoff is V_1 ; when $r_k \leq r_h^c$, the

continuation payoff is V_3 . In summary, the total expected payoff from strategy γ is

$$U_\gamma = -c_j + P[r_j > r_j^c]U_4 + P[r_h^c < r_j \leq r_j^c]U_5 \\ + P[r_j \leq r_h^c] \left\{ -c_k + P[r_k > r_j^c]U_6 + P[r_h^c < r_k \leq r_j^c]V_1 + P[r_j \leq r_h^c]V_3 \right\}$$

Hence,

$$U_\gamma - U_\beta = -c_j + P[r_j > r_j^c]U_4 + P[r_h^c < r_j \leq r_j^c]U_5 \\ + P[r_j \leq r_h^c] \left\{ -c_k + P[r_k > r_j^c]U_6 + P[r_h^c < r_k \leq r_j^c]V_1 \right\} \\ + c_k - P[r_k > r_j^c]U_1 \\ - P[r_h^c < r_k \leq r_j^c] \left\{ -c_j + P[r_j > r_j^c]U_2 + P[r_h^c < r_j \leq r_j^c]U_3 + P[r_j \leq r_h^c]V_1 \right\} \\ - P[r_k \leq r_h^c] \left\{ -c_j + P[r_j > r_h^c]V_2 \right\}$$

Rearranging terms,²

$$U_\gamma - U_\beta = P[r_j > r_j^c] \{ U_4 - P[r_h^c < r_k \leq r_j^c]U_2 - P[r_k \leq r_h^c]V_2 \} - c_j P[r_k > r_j^c] \\ - P[r_k > r_j^c] \{ U_1 - P[r_j \leq r_h^c]U_6 \} + c_k P[r_j > r_h^c] \\ + P[r_h^c < r_j \leq r_j^c]U_5 - P[r_h^c < r_k \leq r_j^c]P[r_h^c < r_j \leq r_j^c]U_3 \\ - P[r_k \leq r_h^c]P[r_h^c < r_j \leq r_j^c]V_2$$

By construction of the compensation value, $P[r_j > r_j^c]U_4 = P[r_j > r_j^c] \{ \int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j \mid r_j > r_j^c] + r_h^c(B - b_j - p_j) \}$, $P[r_h^c < r_k \leq r_j^c]U_2 = P[r_h^c < r_k \leq r_j^c] \{ \int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j \mid r_j > r_j^c] + r_h^c(B - b_j - p_j) \}$, the continuation payoff $P[r_k \leq r_h^c]V_2 \leq P[r_k \leq r_h^c] \{ \int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j \mid r_j > r_j^c] + r_h^c(B - b_j - p_j) \}$, and $P[r_h^c < r_j \leq r_j^c]V_2 \leq P[r_j > r_j^c] \{ \int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j \mid r_h^c < r_j \leq r_j^c] + r_h^c(B - b_j - p_j) \}$. Thus,

$$U_\gamma - U_\beta \geq P[r_k > r_j^c]P[r_j > r_j^c] \left\{ \int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j \mid r_j > r_j^c] + r_h^c(B - b_j - p_j) \right\} - P[r_k > r_j^c]c_j \\ - P[r_k > r_j^c] \{ U_1 - P[r_j \leq r_h^c]U_6 \} + c_k P[r_j > r_h^c] \\ + P[r_h^c < r_j \leq r_j^c]U_5 - P[r_h^c < r_k \leq r_j^c]P[r_h^c < r_j \leq r_j^c]U_3 \\ - P[r_k \leq r_h^c]P[r_h^c < r_j \leq r_j^c] \left\{ \int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j \mid r_h^c < r_j \leq r_j^c] + r_h^c(B - b_j - p_j) \right\}$$

²Here and in the following, I slightly abuse notation so that $P[r_j > r_j^c]P[r_k \leq r_h^c]V_2$ and $P[r_h^c < r_j \leq r_j^c]P[r_k \leq r_h^c]V_2$ both represents the continuation payoff of open box h , but the V_2 in the first expression conditions on $r_j > r_j^c$ and $r_k \leq r_h^c$, while the V_2 in the second expression conditions on $r_h^c < r_j \leq r_j^c$ and $r_k \leq r_h^c$.

Notice that, $P[r_j > r_h^c]P[r_k > r_j^c]U_1 = P[r_j > r_h^c]P[r_k > r_j^c]\{\int_0^{b_k} K_O(b) db + \mathbb{E}[r_k p_k | r_k > r_j^c] + r_k^c(B - b_k - p_k)\}$, $P[r_j \leq r_h^c]P[r_k > r_j^c]U_6 = P[r_j \leq r_h^c]P[r_k > r_j^c]U_1$, and $P[r_h^c < r_j \leq r_j^c]U_5 = P[r_h^c < r_j \leq r_j^c]\{\int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j | r_h^c < r_j \leq r_j^c] + r_h^c(B - b_j - p_j)\}$. Hence,

$$\begin{aligned}
U_\gamma - U_\beta &\geq P[r_k > r_j^c]P[r_j > r_j^c]\left\{\int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j | r_j > r_j^c] + r_h^c(B - b_j - p_j)\right\} - P[r_k > r_j^c]c_j \\
&\quad - P[r_k > r_j^c][r_j > r_h^c]\left\{\int_0^{b_k} K_O(b) db + \mathbb{E}[r_k p_k | r_k > r_j^c] + r_k^c(B - b_k - p_k)\right\} + P[r_j > r_h^c]c_k \\
&\quad + P[r_h^c < r_j \leq r_j^c]\left\{\int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j | r_h^c < r_j \leq r_j^c] + r_h^c(B - b_j - p_j)\right\} \\
&\quad - P[r_h^c < r_k \leq r_j^c]P[r_h^c < r_j \leq r_j^c]U_3 \\
&\quad - P[r_k \leq r_h^c]P[r_h^c < r_j \leq r_j^c]\left\{\int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j | r_h^c < r_j \leq r_j^c] + r_h^c(B - b_j - p_j)\right\}
\end{aligned} \tag{2.5}$$

The following steps are analogous to Weitzman (1979). By definition of the ratio index, $c_j = P(X_j \geq w_j)\{\mathbb{E}[X_j | X_j \geq w_j] - w_j\} = P(r_j \geq r_j^c)\{\mathbb{E}[r_j p_j | r_j \geq r_j^c] - r_j^c p_j\}$ and $c_k = P(X_k \geq w_k)\{\mathbb{E}[X_k | X_k \geq w_k] - w_k\} = P(r_k \geq r_j^c)\{\mathbb{E}[r_k p_k | r_k \geq r_j^c] - r_k^c p_k\} + P(r_h^c \leq r_k < r_j^c)\{\mathbb{E}[r_k p_k | r_h^c \leq r_k < r_j^c] - r_k^c p_k\} + P(r_k^c \leq r_k < r_h^c)\{\mathbb{E}[r_k p_k | r_k^c \leq r_k < r_h^c] - r_k^c p_k\}$. Plugging these equalities into Expression 5 and rearrange terms,

$$\begin{aligned}
U_\gamma - U_\beta &\geq P[r_j > r_j^c]P[r_k > r_j^c] \left\{ \begin{array}{l} \int_0^{b_j} K_O(b) db + r_j^c p_j + r_h^c(B - b_j - p_j) \\ - \int_0^{b_k} K_O(b) db - r_k^c p_k - r_k^c(B - b_k - p_k) \end{array} \right\} \\
&\quad + P[r_h^c < r_j \leq r_j^c]P[r_k > r_j^c]\left\{\int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j | r_h^c < r_j \leq r_j^c] + r_h^c(B - b_j - p_j)\right\} \\
&\quad + P[r_h^c < r_j \leq r_j^c]P[r_h^c < r_k \leq r_j^c] \left\{ \begin{array}{l} \int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j | r_h^c < r_j \leq r_j^c] \\ + r_h^c(B - b_j - p_j) \\ + \mathbb{E}[r_k p_k | r_h^c < r_k \leq r_j^c] - r_k^c p_k - U_3 \end{array} \right\} \\
&\quad + P[r_j > r_j^c]P[r_h^c < r_k \leq r_j^c]\{E[r_k p_k | r_h^c < r_k \leq r_j^c] - r_k^c p_k\} \\
&\quad + P[r_j > r_h^c]P[r_k^c < r_k \leq r_h^c]\{E[r_k p_k | r_k^c < r_k \leq r_h^c] - r_k^c p_k\}
\end{aligned} \tag{2.6}$$

Since $b_j < b_k$, $r_j^c \geq r_h^c \geq r_k^c$, and

$$U_3 \leq \begin{cases} \int_0^{b_k} K_O(b) db + \mathbb{E}[r_k p_k \mid r_h^c < r_k \leq r_j^c] + r_h^c (B - b_k - p_k), & r_j \leq r_k \\ \int_0^{b_j} K_O(b) db + \mathbb{E}[r_j p_j \mid r_h^c < r_j \leq r_j^c] + r_h^c (B - b_j - p_j), & r_k \leq r_j \end{cases}$$

each term on the right hand side of Expression 6 is non-negative. Thus, $U_{comp}^{GSA} \geq U_\gamma \geq U_\beta \geq U_\alpha$ also holds in (Case 4.3).

By induction, $U_{comp}^{GSA} \geq U_\alpha$ for any strategy α that respects the GSA's search termination rule and obtains the compensated knapsack. Therefore, $U_{comp}^{GSA} \geq U_{mod}^*$. \square

Now, I am ready to prove Theorem 2.1.

Proof of Theorem 2.1: Consider an arbitrary search problem \mathcal{S} . Notice that

$$\frac{U^{GSA}(\mathcal{S})}{U^*(\mathcal{S})} = \frac{U^{GSA}(\mathcal{S})}{U_{comp}^{GSA}(\mathcal{S})} \times \frac{U_{comp}^{GSA}(\mathcal{S})}{U_{mod}^*(\mathcal{S})} \times \frac{U_{mod}^*(\mathcal{S})}{U^*(\mathcal{S})}$$

Lemma 2.1 shows that $\frac{U_{mod}^*(\mathcal{S})}{U^*(\mathcal{S})} \geq 1 - \frac{\bar{p}}{B}$, and Lemma 2.2 shows that $\frac{U_{comp}^{GSA}(\mathcal{S})}{U_{mod}^*(\mathcal{S})} \geq 1 - \frac{\bar{p}}{B}$. From Lemma 2.3, $\frac{U_{comp}^{GSA}(\mathcal{S})}{U^*(\mathcal{S})} \geq 1$. Therefore, $\frac{U^{GSA}(\mathcal{S})}{U^*(\mathcal{S})} \geq (1 - \frac{\bar{p}}{B})^2$. \square

2.7 Bibliography

- AHARON, R. AND E. VEENDORP (1983): "Sequential search with a budget constraint," *Economics Letters*, 11, 81–85.
- MANNING, R. AND J. R. MANNING (1997): "Budget-constrained search," *European Economic Review*, 41, 1817–1834.
- MANNING, R. AND P. B. MORGAN (1982): "Search and consumer theory," *The Review of Economic Studies*, 49, 203–216.
- OLSZEWSKI, W. AND R. WEBER (2015a): "A more general Pandora rule?" *Journal of Economic Theory*, 160, 429–437.
- (2015b): "A more general Pandora's rule?" *arXiv preprint arXiv:1509.09095*.

THALER, R. H. (1990): "Anomalies: Saving, fungibility, and mental accounts," *Journal of economic perspectives*, 4, 193–205.

VEENDORP, E. (1984): "Sequential search without reservation price," *Economics Letters*, 16, 53–57.

WEITZMAN, M. L. (1979): "Optimal search for the best alternative," *Econometrica: Journal of the Econometric Society*, 641–654.

CHAPTER 3

PRODUCT RETURNS AND ASSORTMENT DECISIONS: A STRATEGIC ANALYSIS OF ONLINE AND OFFLINE COMPETITION (WITH RAGHUNATH RAO AND PAOLA MALLUCCI)

3.1 Introduction

The share of online commerce in retail sales has been rising fairly rapidly in the past decade. For example, online sales in 2011 accounted for 6.4% of total retail sales, and by 2021, their share had increased to almost 20% (Young, 2020; Ali and Young, 2021), attributable in part to the ongoing pandemic-related concerns and the rapid advancement in logistics, including single-day shipping offered by prominent retailers. This trend is likely to accelerate in the coming years. In general, the allure of online shopping for consumers is quite obvious: A large online marketplace can make almost limitless inventory available to consumers in the comfort of their homes. On the other hand, shopping at physical stores requires a costly trip to a store, and the product choices available to the customer are limited to the inventory on the floor of the store at a given point in time.

The rapid rise of online retailing also has led many observers to worry about the traditional retail format's fate, with some forecasters predicting the impending demise

of the physical stores (The Economist, 2017). In contrast, others have called such doomsday scenarios “premature” (Dennis, 2018). While many big box retailers struggle in the digital era, the profits of specialty stores have grown steadily at an average rate of around 2.8% per year (Daly, 2021). Retail industry analysts and managerial works have also suggested that specialization could be an effective way to shield against online retail platforms (Reinartz et al., 2019; Johansson, 2020).

From the online retailers’ perspective, two main market frictions are considered to be the primary hindrance toward their more widespread adoption. These frictions are (1) delivery times, and (2) product mismatch. Unlike the purchases made at a brick-and-mortar retailer (BMR), where products can be (almost) instantly acquired, purchases from an online retailer (OR) include a lead time that is unappealing to many consumers. In addition, a purchase made at an OR also precludes a close inspection, which can result in more significant instances of mismatch between the product and customer requirements. This risk is especially present for products that have fewer so-called “digital attributes” (Bell et al., 2012). For example, when a customer buys a skirt from an OR, she has to rely on images and size information and has no access to an actual fitting room that would allow for checking the exact fit, texture, and feel. Such uncertainty for fit can be a significant hurdle in consumers’ adoption of Ecommerce, especially in cases where product fit entails greater uncertainty. Industry surveys have indicated that 62% of shoppers have stated that they do not shop online because of the inability to physically inspect products (Skrovan, 2017).

ORs have responded to the first issue (delivery times) by investing in shipping technology that allows for faster shipments, sometimes within a few hours. They have reacted to the second issue (product mismatch) by investing in augmented and virtual reality tools that allow customers to try products (e.g., make-up and apparel) on their own images (Okoli, 2017). Although such tools can help, they are still quite removed from an actual in-store experience. To alleviate the concerns arising from such mismatch issues, most ORs have instituted more generous return policies.

As a consequence of both greater mismatch and generous return policies, the return rates for online purchases are skyrocketing. Industry reports suggest that the return rate could vary between 20 and 40 percent for online purchases, compared with about 5 percent for BMR purchases (Reagan, 2019). Despite the high return rates, consumers often complain about the hassles of online return because it requires repackaging, printing out shipping labels, traveling to a physical location (e.g., the post office), and waiting

to receive a refund for the return (Selyukh, 2018). In addition, consumers returning online purchases worry about the return packages being lost in the mail. Meanwhile, the return process at a physical store is significantly easier because customers just have to drop off merchandise at a customer service desk and almost instantly receive the refund.

The purpose of our research is twofold. First, we study how online retailing impacts the store positioning of BMRs. Second, we want to understand the effect of consumer returns on the product assortment decisions of online and brick-and-mortar stores.¹ To better understand the role of product returns within a competitive setting, we model competition between a physical store and an online store, each of which can carry an assortment of product choices for consumers. The first part of our paper builds a framework to analyze how consumer returns affect the respective retailers' product assortment choices (i.e., the type of products that each retailer carries). Our results indicate that both the consumers' travel costs to the BMR and product return costs have important bearings on retailers' assortment choices. Specifically, we find that the BMR ends up stocking relatively higher value products that might exhibit ex-ante a high or low probability of fit with consumer needs. The OR ends up stocking both higher and lower value products with an ex-ante higher probability of fit. The second part of our paper applies the model to analyze an important new development in the consumer returns infrastructure: Many ORs have started allowing their customers to return their purchased merchandise to (competing) BMRs (e.g., Thomas, 2021). In the short run, this development alters consumers' decisions to visit the OR and BMR because of the reduction in return costs associated with their online purchases. But beyond this first-order effect on consumers, the decision fundamentally alters the equilibrium assortment choices in the long run. In contrast to the main model, the retailers under such a returns regime will eventually differentiate and sell exclusive products. In the long run, the OR expands assortment to stock more goods with a lower fit, while the BMR becomes a more niche player, stocking only high-value products. Overall, under such a return agreement, the product variety increases, consumers shop online more and return more online purchases. The OR always benefits and the BMR can also benefit within such a return agreement under certain conditions.

The rest of the paper proceeds as follows. In Section 3.2, we provide a brief review of the related literature. In Section 3.3, we set up our baseline theoretical apparatus of

¹The product returns are an economically significant phenomenon; for example, in the year 2021, consumers returned products worth more than US\$700 billion—a figure that represents more than 16% of the retail sales (Repko, 2021).

product assortment when product returns are allowed. We analyze the main model in Section 3.4 and analyze the effects when online purchases can be returned to a physical store in Section 3.5. In Section 3.6, we present two extension of our model. We conclude the paper with the implications and limitations of our research in Section 3.7.

3.2 Related Literature

At a broader level, our paper fits into the extensive literature on the effects of e-tailing. This strand of literature has studied a variety of issues including the reduction in price dispersion (Cavallo, 2017), the effects of “showrooming” (Jing, 2018; Kuksov and Liao, 2018), the design of online marketing channels (Dukes and Liu, 2016), advertising efficacy (Goldfarb and Tucker, 2011; Mayzlin and Shin, 2011), and the effect of online word-of-mouth (Godes and Mayzlin, 2004), including reviews and blogs (Mayzlin and Yoganarasimhan, 2012).

Within this broad theme, our paper specifically makes both substantive and methodological contributions to the literature on product returns, as well as the literature on online-offline competition. The first stream of work in this area studies the causes of consumer returns. Researchers have studied a multitude of factors that are related to consumer returns, including: a) the extent of pre-purchase information provided by the firms (Shulman et al., 2015), b) salesperson competence and the in-store environment (Ertekin et al., 2020), c) delivery times (Rao et al., 2014), d) the availability of broader assortments (Rabinovich et al., 2011), and e) the duration of the return deadline (Janakiraman and Ordóñez, 2012).

Because returns are costly for retailers, the second stream of literature has studied policies associated with the management of consumer returns. Shulman et al. (2010) examine the product returns within a supply chain and demonstrate that, despite higher inefficiencies, the optimal approach, at times, is for a manufacturer to accept the returns directly, rather than relying on the retailers. In another study, Shulman et al. (2011) show how restocking fees might endogenously arise in competitive markets, allowing firms to pass the costs of the returns on to consumers. Huang and Zhang (2020) study how a firm’s product design and return policies are affected by consumer valuation and uncertainty. Specifically, they show that when valuation heterogeneity is low, the firm optimally offers two products (one for high types and one for low types) and does not allow returns.

The third stream of work is closer in spirit to our work. This work has studied the competitive interaction of online and offline retailing in conjunction with consumer returns. Ofek et al. (2011) examine competing firms that can start an online channel and that need to manage product returns and to decide on the level of sales assistance in the physical stores. They show, somewhat counter-intuitively, that having an online channel can improve the sales assistance offered at the physical stores. Nageswaran et al. (2020) recently studied the return policies in the context of a single omnichannel seller; they find that full or partial refunds might be optimal, depending on a product's in-store salvage value. Finally, a recent working paper Nageswaran et al. (2021) studies a specific returns policy, similar to the one studied by us, whereby the purchases from an online seller could be returned to competing physical stores. The authors show that the efficacy of this policy depends on the (exogenous) product overlap between the two types of sellers, and the retailers that have lower overlap are more likely to enter into such partnerships.

We contribute to this literature by focusing on how product returns affect retailers' assortment choices by building a parsimonious yet more general framework. For example, previous work in this domain has assumed that firms carry either a single product or a single product category with multiple ex-ante identical variants. Our model allows for a wide product assortment that is heterogeneous in terms of value and of ex-ante consumer fit, allowing us to better understand how online-offline competition might affect assortment decisions. That consumer returns are costly is well understood; thus, because returns have substantial implications on profitability, retailers' assortment decisions are likely to take consumer returns into account. By allowing endogenous assortment decisions within a competitive setting, our model enables us to account for this critical yet missing feature in the current models and thus provide valuable and practical insights. Finally, our substantive application of the emerging practice that enables consumers to buy online and return to a competing physical store points toward how the new ways of processing returns might result in greater consumer access and consequently alter the retail landscape.

The novel "buy online and return offline" policy is likely to grow in the post-Covid era because of the fundamental challenges facing BMRs and their desire to find new sources of revenue. This practice shares similarities with other policies, including ones that allow consumers to buy online and pick up in-store, which may alleviate conflicts between retailers (Gao and Su, 2017). However, the practice also may induce a "research

online, purchase offline" shopping strategy, thus decreasing online sales (Gallino and Moreno, 2014) by diverting traffic to the stores. Existing empirical evidence shows that "buy online, pick up in-store" increases the sale of low-demand products (Gallino et al., 2017). Our game-theoretic model shows that "buy online and return in-store" not only influences consumers' shopping strategy but also changes retail assortment strategies - a crucial ingredient missing from the current literature on consumer returns.

Our work also relates to the limited literature on multi-product retailing. Rhodes (2015) recently built a model that allows for a greater variety in assortment, and showed that a single multi-product retailer can build a low-price image by stocking many products or by advertising low prices for one product in a consumer search framework. In a follow up research, Rhodes et al. (2020) show that a single retailer optimally stocks products that offer a high consumer surplus to attract visits, as well as products with high margins. Earlier works studied the impact of competition on product variety, with simplifying restrictions of a single product category (Kök and Xu, 2011) or two available products (Dukes et al., 2009). We extend this literature by analyzing how competition and return policies affect multi-product retailers' assortment strategies. The operations literature has studied the newsvendor problem regarding the timing of ordering and optimal stocking (Arrow et al., 1951; Holmes, 2001; Kök et al., 2015); we abstract away from these issues to focus on assortment composition. Finally, and more broadly, our work relates to the literature on consumer preferences for product variety (Spence, 1976; Dixit and Stiglitz, 1977; Bronnenberg, 2015, 2020). Bronnenberg (2020) uses a Salop circle model to analyze retailer entry. We differ from this literature in two ways. First, products in our setup are ex-ante heterogeneous in terms of valuation and probability of fit. Second, the extant literature focuses on a competitive equilibrium (with zero firm profits), while we study imperfect competition in which one OR and one BMR compete on assortment choices, with product returns as the main focus.

3.3 Model Setup

3.3.1 Products

Products in a marketplace differ in a number of dimensions. We focus on two orthogonal dimensions - fit and value - which are of particular relevance for studying traditional retailers with shelf capacity constraints and online retailers facing high return rates. Fit

Table 3.1: An example of the product space, $[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]$.

| | low θ | high θ |
|----------|--------------------|------------------------------------|
| high v | jewelry, furniture | motherboard, graphics card |
| low v | footwear, pillow | cleaning product, building tool |

captures product attributes that need to be inspected physically to determine whether the product is a good match for a consumer. One example of such an attribute is a T-shirt cut. Actual fit can only be assessed with a physical inspection of the product, but some products provide a higher ex-ante likelihood of a good match than others. For example, the fit of computer parts is generally easier to infer than to the fit of a furniture item. The value dimension captures vertical quality attributes that can be objectively evaluated and ranked based on a simple description (e.g., the resolution of a computer monitor or the thread count of a T-shirt). These two dimensions - fit and value - parsimoniously capture a wide range of products that could differ among multiple attributes. (See Chapter 3.1 for some examples.)

The market consists of a unit mass of heterogeneous products. Each product has a fit parameter, $\theta \in [\underline{\theta}, \bar{\theta}] \subset [0, 1]$, that captures the ex-ante probability that the product is a good match for the consumer who buys it, and a value parameter, $v \in [\underline{v}, \bar{v}] \subset [0, \infty)$, that captures the value of a matched product. The set of all products forms a joint distribution with a cumulative distribution function (CDF) that is denoted by $F(v, \theta)$. If the product (v, θ) fits a consumer and has price p , the consumer demands $D(v, p)$ units. Otherwise, the product provides no value to the consumer, and she demands zero units.

The demand function $D(v, p)$ is assumed to be increasing in v and decreasing in p , and can also be interpreted as a purchase probability. For example, a special case of our setup is unit demand consumers with consumption utility $v - p - x$, where x is a privately observable individual taste or outside option.² For any $v \in [\underline{v}, \bar{v}]$, we assume there exists a unique profit-maximizing price $p^*(v) = \arg \max_{p \geq 0} D(v, p) \times p$.

²Suppose $x \stackrel{i.i.d.}{\sim} U[0, 1]$, the demand function (purchase probability) is given by $Pr[v - p - x \geq 0] = v - p := D_1(v, p)$. Alternatively, we can also let consumption utility be $xv - p$. Then, the consumer buys the product with probability $Pr[xv - p \geq 0] = 1 - \frac{p}{v} := D_2(v, p)$. Both $D_1(v, p)$ and $D_2(v, p)$ satisfy our assumptions on the demand function.

3.3.2 Retailers

There are two profit-maximizing retailers. One retailer is the BMR (e.g., Kohl's), and the other is the OR (e.g., Amazon.com). The BMR has limited shelf space and can display at most a fraction of $K \in (0, 1)$ of the products. The OR does not face a shelf space constraint and can display a very large number of products. We abstract away from inventory capacity constraints for each product but capture that the BMR is limited by its store size.

Each retailer chooses its own product assortment. Denote $\mathcal{B} \subset [\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]$ as the set of products sold by the BMR and $\mathcal{O} \subset [\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]$ as the set of products sold by the OR.³ Each retailer i also sets price $p_i(v, \theta)$ for every product it offers. The marginal costs of stocking and selling are assumed to be zero.

3.3.3 Consumers

There is a unit mass of consumers. Each consumer has a potential demand for all products that fit but can learn whether a product fits only through physical inspection (at the BMR). The consumers can choose to visit only one retailer, or visit two retailers sequentially, or visit neither retailer and take their outside option 0.

To visit the BMR, a consumer has to incur a privately known travel cost $s \in [\underline{s}, \bar{s}] \subset [0, \infty)$. The travel cost is heterogeneous across consumers, and s follows a distribution with CDF $G(s)$. Upon arriving at the BMR, a consumer perfectly learns (through costless physical inspection) whether each product $(v, \theta) \in \mathcal{B}$ fits. Consumers can also visit the OR and search online for a subset of products in \mathcal{O} to learn prices and make purchase decisions. The travel cost and search cost of online shopping are zero, but product fit can be observed only after purchase. If a product purchased online does not fit, the consumer can return it for a full refund at a hassle cost of $r > 0$ per product. The OR also incurs a cost $R > 0$ for each return.

The travel cost s is a proxy for opportunity cost that results from the travel time and the time spent within a store (e.g., searching and inspecting products and waiting in checkout lines). This cost is likely to be a function of the distance of individuals' homes to the store, as well as the availability and accessibility of various transportation methods. This cost is privately known by consumers, and retailers observe only the aggregate population statistics $G(s)$ (e.g., learning through marketing research). We

³We require the sets \mathcal{O} and \mathcal{B} to be measurable with respect to the σ -algebra generated by (v, θ) .

make two simplifying assumptions about consumers' return hassle cost r : (1) We assume the hassle cost of returning is not too small, so if a consumer gains non-negative expected utility from buying a product online, the OR also obtains non-negative expected profit from selling the product to her.⁴ (2) Consumers return all online purchases that do not fit. For this to be optimal, it suffices to assume that the hassle cost of returning is not too large.⁵

In summary, a consumer's expected utility can be written as

$$U = \int_{\text{purchases}} u(v, \theta, p) dF(v, \theta) - \int_{\text{returns}} r dF(v, \theta) - s \times \mathbb{1}_{\text{visit BMR}},$$

where $u(v, \theta, p) = \theta \int_p^\infty D(v, \tilde{p}) d\tilde{p}$ is the expected surplus of purchasing a product (v, θ) at price p . In the expected utility equation, the first integration represents the total expected surplus from consumption, the second integration represents the total cost of returning all products that do not fit, and the third term is a travel cost that is incurred only if the consumer visits the BMR.

3.3.4 Timing

First, the BMR and the OR choose their product assortment, \mathcal{B} and \mathcal{O} , respectively. Second, the two retailers decide prices $p(v, \theta)$ for each product (v, θ) that they offer. Third, consumers observe the retailers' assortment decisions and their private travel cost s , and they choose to visit only one retailer or to visit two retailers sequentially, or to take their outside option 0.

This sequence of timing reflects the fact that retailers' brand positioning is a long-term strategy, and prices change relatively more frequently than assorted SKUs. We assume that consumers observe the product assortments of each retailer but not the prices of each product.⁶ This assumption captures the idea that consumers learn about retailers' assortment and general price positioning over time but do not directly observe the daily fluctuating prices before arriving at stores.

⁴The first condition allows us to rule out a complicated case where the OR may sell some products with low θ and lose money (due to high return probability) to attract more consumers.

⁵Specifically, we impose the assumption that $r \in [\underline{r}, \bar{r}]$, where $\underline{r} = \max_v \frac{\int_{p^*(v)}^\infty D(v, p) dp}{D(v, p^*(v)) \times p^*(v)} R$ and $\bar{r} = p^*(v) = \arg \max_{p \geq 0} D(v, p) \times p$.

⁶We assume that consumers' price beliefs are constant on and off the equilibrium path. This passive beliefs assumption is common in the sequential search literature and implies that when consumers are surprised by one retailer's deviation, their beliefs about the other retailer's prices do not change.

Table 3.2: Table of notations.

| Notation | Definition |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| v | The value of a product. |
| θ | The ex-ante probability of fit of a product. |
| $F(v, \theta)$ | The cumulative distribution function of (v, θ) |
| $p(v, \theta)$ | The price of a product (v, θ) . |
| p^e | Consumers' anticipated price. |
| $D(v, p)$ | The demand for a product, given value v and price p . |
| $p^*(v)$ | The price that solves $\max_p D(p, v)p$. |
| $u(\theta, v, p)$ | The consumer surplus for a product (v, θ) that is priced at p . |
| $u(\theta, v)$ | The consumer surplus for a product (v, θ) that is priced at $p^*(v)$. |
| s | Consumers' cost of traveling to the BMR. |
| $G(s)$ | The cumulative distribution function of s . |
| r | Consumers' hassle cost of returning a product purchased online. |
| R | The OR's hassle cost of handling a product return. |
| \bar{R} | The OR's per-unit hassle cost of handling returns to the BMR. |
| \mathcal{O} | The set of all products stocked in the online store. |
| \mathcal{B} | The set of all products stocked in the BMR. |
| K | The store size capacity of the BMR. |
| \hat{s} | $\hat{s} = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta)$ |
| \hat{s}_{BORS} | $\hat{s}_{BORS} = \int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta)$ |

A summary of the main parameters is available in Chapter 3.2 for easy reference.

3.4 Model analysis

Before analyzing the main model, we use a simple example to illustrate the key insights of the model.

3.4.1 An illustrative example

To understand the key insights of our model, consider a scenario with only four different products available to the retailers, $(v, \theta) \in \{v_l, v_h\} \times \{\theta_l, \theta_h\}$. Assume $v_h = \theta_h = 1$ and $v_l, \theta_l \in (0, 1)$. Due to the store size constraint $K = \frac{1}{2}$, the BMR retailer can sell no more than $4K = 2$ products. Consumers' travel cost to the BMR follows the uniform distribution $s \sim U[0, 1]$. If a product (v, θ) fits a consumer, she purchases with

probability $D(v, p) = v - p \in [0, 1]$. For simplicity of the example, we directly impose the assumption that a product (v, θ) is always priced at the “monopoly” price $\frac{v}{2}$.

Case 1: Only BMR present. Suppose that the BMR offers two products i and j , then its profits are given by $\pi(i, j) = \frac{\theta_i v_i^2 + \theta_j v_j^2}{4} \hat{s}$, where the first term is the BMR’s expected profit obtained from a single consumer, and the second term $\hat{s} = \frac{\theta_i v_i^2 + \theta_j v_j^2}{8}$ represents the store foot traffic. The BMR should optimally offer products with high value (fit) if $\theta_l (v_l)$ is large enough,

$$\mathcal{B}_1 = \begin{cases} \{(v_h, \theta_h), (v_h, \theta_l)\}, & \theta_l \geq \sqrt{1 + v_l^2} - 1 \\ \{(v_h, \theta_h), (v_l, \theta_h)\}, & \theta_l < \sqrt{1 + v_l^2} - 1 \end{cases}$$

We can see that the BMR always prefers products with high value and high fit, as it always sells (v_h, θ_h) but never sells (v_l, θ_l) . The reason is straightforward - the high value and high fit both increase the expected profit (per unit of shelf space), and also attract more consumers to the store. However, as shown below, this intuition does not work when there is competition.

Case 2: Online-offline competition. When the OR enters the market, it avoids selling products with low fit probability due to high costs of processing returns, so $\mathcal{O} = \{(v_h, \theta_h), (v_l, \theta_h)\}$. Consider what would happen if the BMR offers $\mathcal{B} = \mathcal{O}$. This strategy was previously optimal when v_l is large enough. But now, since consumers can conveniently shop online and not worry about returns (due to the assumption $\theta_h = 1$), the BMR would get zero foot traffic! Thus, the BMR would want to differentiate from the OR and optimally select the assortment $\mathcal{B}_2 = \{(v_h, \theta_h), (v_h, \theta_l)\}$ regardless of the values of v_l and θ_l . Perhaps more surprisingly, when $\theta_h < 1$, the BMR may want to further increase consumer foot traffic by full differentiating and specialize in only selling niche products, i.e., $\mathcal{B}_2 = \{(v_h, \theta_l), (v_l, \theta_l)\}$.

In general, online competition incentivizes the BMR to specialize and sell niche (low θ) products. At the cost of lower sales volume, the specialization by BMR increases foot traffic for two reasons. First, offering niche products helps it to differentiate from the OR. Second, from a consumer’s perspective, the value of in-store inspection is larger for niche products, because the expected return cost $(1 - \theta)r$ is higher at OR. As shown in the ensuing main model, this insight is robust within a more general framework.

3.4.2 Benchmark: Only BMR present

We start with the analysis of a scenario wherein there is no OR. So, consumers can choose either to visit the BMR or opt for the outside option. We solve the game using backward induction. Since the prices are not observable prior to search, the smallest subgame consists of the last two stages, where the BMR sets its prices and consumers choose their shopping strategy.

Lemma 3.1. *Suppose there is only BMR in the market. In the unique perfect Bayesian equilibrium of the last two stages, the BMR sets price $p^*(v)$ for product (v, θ) . A consumer visits the BMR if*

$$s \leq \hat{s}_{bench} = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta).$$

Otherwise, the consumer picks her outside option.


Now we analyze the first stage of the game, where the BMR chooses its product assortment. The BMR's profit maximization can be expressed as,

$$\begin{aligned} \max_{\mathcal{B} \subset [\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} & \int_{\mathcal{B}} \theta D(v, p^*(v)) p^*(v) dF(v, \theta) \times G(\hat{s}_{bench}) \\ \text{s.t. } & \hat{s}_{bench} = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) \\ & \int_{\mathcal{B}} dF(v, \theta) \leq K \end{aligned}$$

Here, the objective function is the profit obtained per consumer ($\int_{\mathcal{B}} \theta D(v, p^*(v)) p^*(v) dF(v, \theta)$) multiplied by the total foot traffic ($G(\hat{s}_{bench})$). This expected profit is maximized over all the possible product assortments $\mathcal{B} \subset [\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]$, and the two constraints reflect the consumer behavior and BMR's store capacity.

Lemma 3.2. *Suppose there is only one BMR. In equilibrium, the BMR stocks a product (v, θ) if and only if*

$$\theta D(v, p^*(v)) p^*(v) G(\hat{s}_{bench}) + \lambda_2 u(v, \theta) \geq \lambda_1,$$



Note. This figure plots the equilibrium product assortment when only BMR is in the market. Here, the set of all products (v, θ) is jointly uniform on $[0, 1] \times [0, 1]$, consumers' travel cost is $s \stackrel{iid}{\sim} U[0, 1]$, demand function is $D(v, p) = 8v(1 - p)$, consumers' marginal cost of returning is $r = \frac{1}{2}$, and the BMR's store capacity is $K = \frac{1}{4}$.

Figure 3.1: Optimal product assortment without online competition.

where $\{\hat{s}_{bench}, \lambda_1, \lambda_2\}$ is the solution to the system of equations

$$\begin{cases} \int_{\mathcal{B}} \theta D(v, p^*(v)) p^*(v) dF(v, \theta) \times g(\hat{s}_{bench}) = \lambda_2 \\ \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) = \hat{s}_{bench} \\ \int_{\mathcal{B}} dF(v, \theta) = K \end{cases}$$

Lemma 3.2 implies that the capacity constrained BMR offers products with sufficiently high v and θ (see Figure 3.1 for an example). The intuition is simple. Products with high v yield higher profit for each unit of sales, and products with high θ have a higher chance of being sold. In what follows, we analyze the main model and show that this seemingly simple intuition breaks down in the presence of online retail competition.

3.4.3 The main model

Now we come to the main model where a BMR competes with an OR. We solve for the subgame perfect equilibrium using backward induction. Because consumers cannot directly observe prices, the smallest proper subgame comprises the last two stages, where the two retailers set prices and the consumers form their shopping strategies.

3.4.3.1 Consumers' shopping strategies and retailers' pricing strategy

We seek a simple perfect Bayesian equilibrium where the two retailers do not compete on prices. This allows us to abstract away from pricing and thus to focus on the retailers' assortment decisions - the main focus of our paper. The comparison with the benchmark is also cleaner because only assortment changes.

First, we analyze consumers' optimal shopping strategies, assuming their belief that the two retailers set equal prices for the same product. Then, we show that, under this consumer belief, it is indeed optimal for the two retailers to set equal prices for the same product (if they both offer it). So consumers' price beliefs are correct on the equilibrium path.

Suppose a consumer anticipates that a product (v, θ) is priced at $p^e(v, \theta)$. The consumer chooses between visiting only the BMR, visiting only the OR, or visiting both retailers in a sequential order. If the consumer visits only the BMR, she pays travel cost s , buys all products in \mathcal{B} that fit, and obtains $\int_{\mathcal{B}} u(v, \theta, p^e) dF(v, \theta) - s$. If consumer visits only the OR, she searches and buys all products that have expected consumer surplus greater than the expected hassle cost of returning. So the payoff obtained from only online shopping is $\int_{\mathcal{O}} [u(v, \theta, p^e) - (1 - \theta)r]^+ dF(v, \theta)$. Here, $[u(v, \theta, p^e) - (1 - \theta)r]^+$ takes value $u(v, \theta, p^e) - (1 - \theta)r$ if $u(v, \theta, p^e) \geq (1 - \theta)r$, and takes value 0 otherwise.

If a consumer plans to visit both stores, then we show that visiting the BMR first is always optimal (see the proof of Proposition 3.1 in the Appendix). Intuitively, this is because if a product is available in both stores, then the consumer can save the return hassle by inspecting the product in the BMR. Using this shopping strategy, the expected utility from shopping in the BMR is given by $\int_{\mathcal{B}} u(v, \theta, p^e) dF(v, \theta) - s$. The expected utility he obtains from online shopping is the expected consumer surplus minus the expected cost of returning, $\int_{\mathcal{O} \cap \mathcal{B}^c} [u(v, \theta, p^e) - (1 - \theta)r]^+ dF(v, \theta)$, where \mathcal{B}^c denotes the complement of set \mathcal{B} . In aggregate, the payoff of visiting the BMR first and the OR second is equal to

$$\int_{\mathcal{B}} u(v, \theta, p^e) dF(v, \theta) - s + \int_{\mathcal{O} \cap \mathcal{B}^c} [u(v, \theta, p^e) - (1 - \theta)r]^+ dF(v, \theta).$$

In the perfect Bayesian equilibrium of this subgame, consumers choose their shopping strategies to maximize expected utility, given their travel costs and price beliefs; retailers set prices, taking as given consumers' shopping strategies; and the consumers' beliefs are correct on the equilibrium path (i.e., $p^e(v, \theta) = p(v, \theta)$). Let $u(v, \theta) = u(v, \theta, p^*(v))$

for simplicity of notations, the following proposition characterizes the equilibrium for the last two stages.

Proposition 3.1. *In the last two stages, there exists a perfect Bayesian equilibrium where both retailers set price $p(v, \theta) = p^*(v)$ for each product (v, θ) in their respective assortment. Furthermore, in this equilibrium, there exists a cut-off travel cost,*

$$\hat{s} = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta),$$

such that:

1. *A consumer with $s \leq \hat{s}$ visits the BMR first to buy all products that fit, and then visit the OR to search and buy easy-to-fit products that are unavailable in the BMR, $\{(\theta, v) \in \mathcal{O} \cap \mathcal{B}^c \mid u(v, \theta) - (1 - \theta)r \geq 0\}$.*
2. *A consumer with $s > \hat{s}$ searches and buys products $\{(\theta, v) \in \mathcal{O} \mid u(v, \theta) - (1 - \theta)r \geq 0\}$ only from the OR.*

Proposition 3.1 establishes consumers' optimal shopping strategy and retailers' equilibrium pricing. While the consumer's part is easy to understand, the pricing part deserves more explanation. Intuitively speaking, if prices are unobservable prior to consumer search, then retailers might as well set equally high prices, as they cannot attract additional consumers by (deviating from equilibrium and further) lowering prices. Consumers understand this, and expect no price undercutting in equilibrium. This forms a best response cycle. Note that the resulting equilibrium pricing outcome is not one where firms are making zero profits. In fact, it is an intuitive pricing outcome where products with a higher value are priced higher across all the channels (i.e., p^* is increasing in v).

Due to consumers' zero cost of browsing online, the indifference between searching online or not can lead to other perfect Bayesian equilibria where retailers do compete on prices. This possibility is analyzed in the extension section. It is shown that price competition does not qualitatively change a consumer's shopping strategy and the retailers' assortment decisions. Furthermore, while price competition complicates the analysis, it strengthens our main results. For now, we select the simple equilibrium where the "law of one price" holds to focus attention on the product assortment choices. This contrasts the literature, which abstracts away from product assortment choices to study specific pricing policies.

While simple, our equilibrium selection is also plausible in two ways. First, this equilibrium outcome mirrors the common practice of retailers following manufacturers' suggested retail prices (MSRP). Second, from a theoretical standpoint, consider a perturbed model where consumers face a small marginal cost of browsing online. With this perturbation, searching online for a product exposes consumers to a well-known holdup problem, similar to the Diamond paradox (Diamond, 1987, see also Armstrong, 2017; Rhodes et al., 2020 for similar examples). Even when the perturbation is arbitrarily small, the equilibrium characterized in Proposition 3.1 is unique, in the sense that any other equilibria can only differ from Proposition 3.1 on a set of products that has measure zero. The formal statement and proof can be found in Proposition 3.10 of the Appendix.

3.4.3.2 Equilibrium product assortment

Next, we analyze the first stage of the game, where the two retailers choose their product assortments.

The OR takes into account two factors when choosing its assortment. First, it accounts for the cost of returns and wants to stock products that yield the highest expected profit per consumer (i.e., products with high $\theta D(v, p^*(v))p^*(v) - (1 - \theta)R$). Second, it wants to attract more consumers to use the OR as their one-stop shopping destination by stocking products that give a higher expected surplus to consumers (i.e., products with high $u(v, \theta) - (1 - \theta)r$). In the Appendix, we elaborate on the OR's optimization problem and show that the two incentives are aligned under our assumption $r \in [\underline{r}, \bar{r}]$. Thus, the OR has a (weakly) dominant strategy: A product (v, θ) is stocked by the OR if and only if consumers obtain non-negative expected utility from buying the product online (i.e., $\mathcal{O} = \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}$).

The BMR chooses its product assortment to maximize expected profit, taking as given consumers' shopping strategy and its own store size constraint. Similar to the benchmark, the profit maximization can be stated as

$$\begin{aligned} & \max_{\mathcal{B} \subset [\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} \int_{\mathcal{B}} \theta D(v, p^*(v)) p^*(v) dF(v, \theta) \times G(\hat{s}) \\ & \text{s.t. } \int_{\mathcal{B}} dF(v, \theta) \leq K \\ & \hat{s} = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta) \end{aligned}$$

Here, the objective function is the profit obtained per consumer ($\int_{\mathcal{B}} \theta D(v, p^*(v)) p^*(v) dF(v, \theta)$) multiplied by the total number of consumers who visit the BMR ($G(\hat{s})$). This expected profit is maximized over all possible product assortments $\mathcal{B} \subset [v, \bar{v}] \times [\underline{\theta}, \bar{\theta}]$, while taking into account two constraints. The first constraint says that the total number of products in stock cannot exceed the store capacity, and the second constraint specifies the consumers' optimal shopping strategies.

This objective leads to an infinite dimensional optimization problem that requires a constrained maximization over the space of all measurable subsets of $[v, \bar{v}] \times [\underline{\theta}, \bar{\theta}]$. To make the problem tractable, we break the problem into five steps, and use a technical trick first introduced by Rhodes et al. (2020). First, we transform the optimization into the constrained maximization of a functional by replacing the integration domain \mathcal{B} with a binary stocking function that takes value 1 when the BMR stocks a product and 0 otherwise (i.e., $q(v, \theta) = \mathbb{1}_{(\theta, v) \in \mathcal{B}}$).⁷ Second, we prove that a solution to the optimization problem exists. Third, we show that the capacity constraint must hold with equality at the optimum. Fourth, we write a Lagrangian with one objective and two equality constraints. After some manipulation, we get

$$\begin{aligned} \mathcal{L} = & \int_{\mathcal{O}} [\theta D(v, p^*(v)) p^*(v) G(\hat{s}) - \lambda_1 + \lambda_2(1 - \theta)r] q(v, \theta) dF(v, \theta) \\ & + \int_{\mathcal{O}^c} [\theta D(v, p^*(v)) p^*(v) G(\hat{s}) - \lambda_1 + \lambda_2 u(v, \theta)] q(v, \theta) dF(v, \theta) + \lambda_1 K - \lambda_2 \hat{s}, \end{aligned}$$

where \mathcal{O}^c denotes the complement of set \mathcal{O} . Finally, we exploit the linearity (in the function $q(v, \theta)$) of the Lagrangian to obtain the optimal product assortment. The following proposition characterizes the equilibrium product assortment decisions.

Proposition 3.2. *In equilibrium, the OR stocks products $\mathcal{O} = \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}$. The BMR stocks a product (v, θ) if and only if*

$$\begin{cases} \theta D(v, p^*(v)) p^*(v) G(\hat{s}) + \lambda_2(1 - \theta)r \geq \lambda_1, & (v, \theta) \in \mathcal{O} \\ \theta D(v, p^*(v)) p^*(v) G(\hat{s}) + \lambda_2 u(v, \theta) \geq \lambda_1, & (v, \theta) \notin \mathcal{O} \end{cases}$$

⁷Here, $\mathbb{1}_{(\theta, v) \in \mathcal{B}}$ denotes the indicator function that takes value 1 when $(\theta, v) \in \mathcal{B}$ and 0 otherwise.

where $\{\hat{s}, \lambda_1, \lambda_2\}$ is a solution to the system of equations

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \hat{s}} = 0 \\ \int_{\mathcal{B}} dF(v, \theta) = K \\ \hat{s} = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [u(v, \theta) - (1 - \theta)r] dF(v, \theta) \end{cases}$$

The intuition behind the BMR's equilibrium product assortment selections involves an analysis of marginal cost and benefit. First, by the Envelope Theorem, λ_1 is the BMR's marginal benefit if it can marginally increase store capacity K , and it represents the shadow value of store shelves. Hence, the marginal (opportunity) cost of stocking a product (v, θ) (and taking up an additional unit of store shelf) is λ_1 . Second, the marginal benefit of stocking a product (v, θ) can be decomposed into two components: the *direct effect* of profiting from this product, and the *indirect effect* of making the BMR slightly more attractive by carrying this product. The direct effect is given by $\theta D(v, p^*(v)) p^*(v) G(\hat{s})$, where the product fits each consumer with probability θ , each unit of sales leads to profit $D(v, p^*(v)) p^*(v)$, and a total of $G(\hat{s})$ consumers visits the BMR. The indirect effect is the profit of attracting one additional consumer multiplied by the increase in foot traffic by offering the product (v, θ) . By the Envelope Theorem, the marginal profit of attracting consumers is λ_2 . The increase in foot traffic depends on whether this product is available online: When this product is available online (i.e., $(v, \theta) \in \mathcal{O}$), consumers save the expected return cost $(1 - \theta)r$ by visiting and inspecting the product in the BMR; when this product is not available online (i.e., $(v, \theta) \notin \mathcal{O}$), consumers gain an additional consumption surplus of $u(v, \theta)$ when visiting the BMR. Finally, Proposition 3.2 states that the BMR should stock a product (v, θ) if the marginal benefit (sum of the direct and indirect effects) exceeds the marginal cost.

With specific function forms, the equilibrium stocking decisions often are complicated and hard to express in closed forms. However, we can obtain closed-form solutions in a simpler setting using a degenerating, one-dimensional product variety space. See the Appendix for an explicitly solved example.

It is useful to briefly contrast this equilibrium product assortment with the case where the BMR is a monopoly seller. Absent online competition, the BMR prefers to sell products with high v and high θ (see the benchmark section for details). This is because a higher v implies higher profit for each unit of sales, and a higher θ implies a higher chance of selling each product. However, this seemingly simple intuition does not

extend to online-offline competition.

Proposition 3.3. *With online-offline competition,*

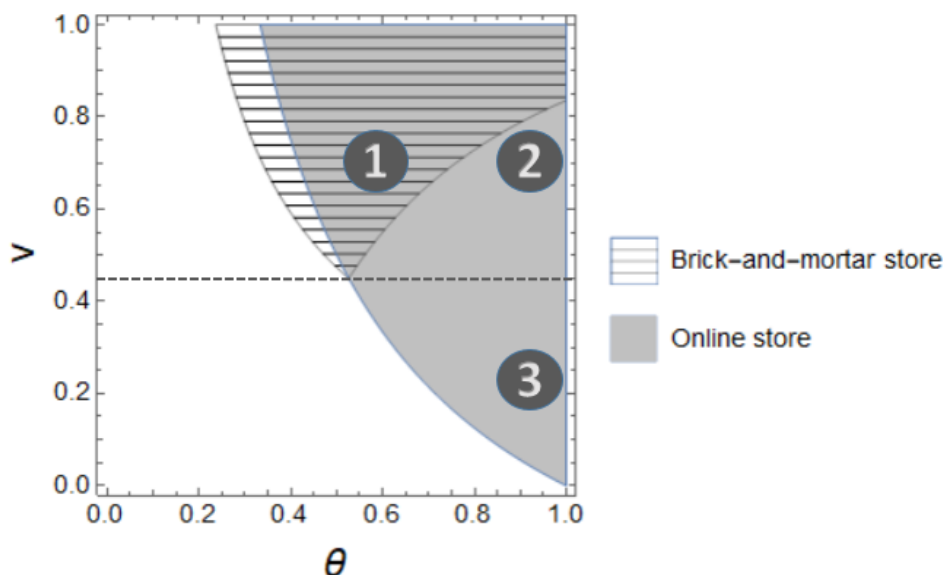
1. *the OR sells products with high fit probability, but products can have high or low value;*
2. *the BMR sells products with high value, but products can have high or low fit probability.*

Proposition 3.3 strikingly shows that, in a competitive setting, a BMR may choose to “waste” its limited store shelf on products with ex-ante *low* fit probability (low θ). The reason involves a subtle trade-off between sales volume and store foot traffic. Intuitively, if a product is sold by both of the retailers, then a consumer obtains an expected consumption profit $u(v, \theta)$ regardless of where the product is purchased. But inspecting the product at the BMR before purchasing can save an expected return hassle cost of $(1 - \theta)r$. Since the return probability $1 - \theta$ decreases in θ , showrooming is more valuable for products with low θ . Thus, the BMR may want to leverage its showrooming capability and sell niche products to attract more foot traffic, only when there is a competing OR. In the language of marginal cost and benefit analysis, for products that are not sold online, both the direct and indirect effects increases in θ ; while for products that are available online, the direct effect still increases in θ , but the indirect effect decreases in θ . For a product $(v, \theta) \in \mathcal{O}$, the BMR prefers lower θ when the indirect effect (benefit of attracting consumers) exceeds the direct effect (loss of lower sales volume).

Figure 3.2 provides a numerical example where the indirect effect of attracting consumers dominates.⁸ In this figure, products ① and ② have the same value, but the BMR only sells product ①, which has a *lower* fit probability. Note that this specialization of the BMR is a consequence of online competition - it would never be optimal without a competing OR (see the benchmark and Figure 3.1).

Our analysis predicts that the prevalence of online retailing may eventually push some traditional retailers to transform into specialty stores. This prediction is supported by the steady growth of specialty stores in the downfall of traditional BMR (Daly, 2021),

⁸To calculate the system of (integral) equations in Proposition 3.2, we first use the interpolation method to approximate the equations. Then, we search for the roots of the system of approximated equations. See the Appendix for details.



Note. This figure plots the equilibrium product assortments of the main model. The BMR gives up selling ②, and instead, sells product ① which has the same value but lower fit probability. Only the OR stocks low value products such as ③.

Figure 3.2: Optimal product assortment in the main model.

and also reinforces the hypothesis that “Specialization may ... effectively shield against the power of online retailer platforms” (Reinartz et al., 2019).

Furthermore, a direct consequence of Proposition 3.3 is that some low value products are sold only online but not in-store. As shown in Figure 3.2, only the OR sells products (e.g. product ③) below the cut-off value \hat{v} (indicated by the horizontal dotted line). The reason is that the OR can sell low value products as long as consumers do not return too often. However, the BMR cannot “afford” to sell these products, because its physical store capacity constraint implies a forbiddingly high *opportunity* cost (i.e., to stock product ③, the BMR needs to give up selling a high value product such as ①). This result is summarized below.

Proposition 3.4. *There exists a cut-off $\hat{v} > 0$ such that only the OR sells products with $v \leq \hat{v}$.*

3.5 Application: Buy Online and Return in Store Agreements

The inability of a consumer to inspect a product before purchase drives the OR to carry products that have mass market appeal, some of which have relatively low value. In contrast, the BMR's capacity constrains it to forgo some of the low value products, which makes visiting its store less attractive and lowers foot traffic.

Possibly to overcome these limits, some ORs and BMRs have reached agreements that allow shoppers to buy online and return at a competitor's physical store (i.e., buy-online-return-in-store, or BORS). For example, Amazon.com recently reached an agreement with Kohl's, allowing products purchased from Amazon.com to be returned to any Kohl's store (Thomas, 2021). Similarly, some convenience stores in Japan offer consumers the option to return online purchases to local convenience stores. In addition, new e-commerce intermediaries have emerged to help facilitate the BORS option. For example, a firm called Navar manages the returns of many online retailers at physical locations of retailers like Walgreens and Nordstrom (Nageswaran et al., 2021). The existence of BORS suggests that ORs and BMRs can mutually benefit in the short run. However, the long-run consequences of such agreements on product assortment choices are unclear. In this section, we analyze how these agreements affect both consumers and the (long-run) assortment choices of retailers, as well as the conditions for BORS to be sustained in the long run.

With BORS, we assume that consumers who want to return products purchased online can either return by incurring hassle cost r or can incur travel cost s to visit the BMR and return these purchases without hassle. If a product is returned via BMR, the OR incurs handling cost $\tilde{R} \geq 0$. We assume this cost is no larger than the cost of handling a direct return, $\tilde{R} \leq R$, which could be seen as the effect of economies of scale on the cost of handling returns. Other details of the model setup are the same as the main model.

In what follows, we first solve for the equilibrium in the case with BORS. Then, we analyze the short-run and long-run consequences of BORS by comparing this equilibrium with the main model.

3.5.1 Model analysis

Similar to the main model, we use backward induction and start from the last two stages, where retailers set prices and consumers form their shopping strategy. As in the main model, we show that both retailers optimally charge an equilibrium price $p^*(v)$ for product (v, θ) . So rational consumers have no incentive to cherry pick for the lowest price. In the following, we analyze the consumer shopping strategy in the case of a BORS option.

To do so, we compare two shopping strategies: visiting only the OR and visiting both retailers sequentially. The reason this approach suffices is identical to that in the main model: Visiting the OR is costless, so visiting only the BMR is suboptimal. In contrast to the main model, we prove that, if a consumer visits both the retailers, then visiting the OR first is optimal (see the proof of Lemma 3.4 in the Appendix). The intuition is that, if a consumer plans to visit the BMR, visiting it second can avoid the return hassle for online purchases. Directly comparing the expected utility of the two shopping strategies, we get the following result.

Lemma 3.3. *Suppose there is a BORS agreement. In equilibrium, both the retailers set price $p(v, \theta) = p^*(v)$ for each product (v, θ) in their respective stock. Furthermore, there exists a cut-off travel cost*

$$\hat{s}_{BORS} = \int_{B \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta)$$

such that:

1. *If travel cost $s \leq \hat{s}_{BORS}$, a consumer searches and buys all products in \mathcal{O} , and then visits the BMR to return unwanted products and to purchase offline exclusive products.*
2. *If travel cost $s > \hat{s}_{BORS}$, the consumer only searches and buys online products with $u(v, \theta) - r(1 - \theta) \geq 0$. Returns are made by incurring the hassle cost r .*

Now, we derive the stocking decisions starting with the OR. First, consider a product with $u(v, \theta) - (1 - \theta)r \geq 0$. Consumers will purchase this product regardless of how they plan to return unwanted products. The OR obtains a positive profit from selling this product. So it should offer all products with $u(v, \theta) - (1 - \theta)r \geq 0$. Second, consider a product with $u(v, \theta) - (1 - \theta)r < 0$. Consumers will purchase this product only if they

plan to return unwanted products to the BMR retailer. The OR is willing to sell these products if the expected revenue minus the cost of handling returns to BMR is positive (i.e. $u(v, \theta) - (1 - \theta)\tilde{R} \geq 0$). In the Appendix, we combine the two cases and show that the OR sells

$$\mathcal{O} = \{(v, \theta) \mid \theta D(v, p^*(v))p^*(v) - (1 - \theta)\tilde{R} \geq 0\}.$$

By Lemma 3.3, with the BORS option, the BMR never gets visited first. Hence, the BMR optimally avoids stocking products that are available online. The BMR maximizes its total profit, which is the profit obtained from each consumer ($\int_{\mathcal{B}} \theta D(v, p^*(v))p^*(v) dF(v, \theta)$) multiplied by the number of consumers that travel to its store ($G(\hat{s}_{BORS})$). Thus, the BMR's profit maximization problem can be written as

$$\begin{aligned} & \max_{\mathcal{B} \subset \mathcal{O}^c} \int_{\mathcal{B}} \theta D(v, p^*(v))p^*(v) dF(v, \theta) G(\hat{s}_{BORS}) \\ & \text{s.t.} \int_{\mathcal{B}} dF(v, \theta) \leq K \\ & \hat{s}_{BORS} = \int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) \end{aligned}$$

Here, the first constraint represents the BMR's store size capacity, and the second constraint characterizes consumers' optimal shopping strategies when returning online purchases to a BMR is feasible.


Similar to the main model, we transform the BMR's profit maximization into a constrained maximization of a functional. Then, we exploit the linearity of the Lagrangian to characterize the BMR's optimal product assortment.

The next lemma summarizes the equilibrium product assortments.

Lemma 3.4. *Suppose there is a BORS agreement. In equilibrium, the OR stocks products $\mathcal{O} = \{(v, \theta) \mid \theta D(v, p^*(v))p^*(v) - (1 - \theta)\tilde{R} \geq 0\}$.*

1. *If $K \geq \int_{\mathcal{O}^c} dF(v, \theta)$, then the BMR's capacity is not binding, and it stocks $\mathcal{B} = \mathcal{O}^c$.*
2. *If $K < \int_{\mathcal{O}^c} dF(v, \theta)$, then the BMR's capacity is binding, and it stocks*

$$\mathcal{B} = \{(v, \theta) \notin \mathcal{O} \mid \theta D(v, p^*(v))p^*(v)G(\hat{s}_{BORS}) + \lambda_2 u(v, \theta) \geq \lambda_1\},$$



Note. The left figure plots an equilibrium stocking when the BMR's capacity is binding, and the right plots a case where capacity is not binding. Here, product variety (v, θ) is jointly uniform on $[0, 1] \times [0, 1]$, consumers' travel cost is $s \stackrel{iid}{\sim} U[0, 1]$, demand function is $D(v, p) = 8v(1 - p)$, consumers' marginal cost of returning is $r = \frac{1}{2}$, and the OR's cost of handling returns to BMR is $\tilde{R} = \frac{1}{3}$.

Figure 3.3: Optimal product assortment in the presence of a BORS agreement.

where $\{\hat{s}_{BORS}, \lambda_1, \lambda_2\}$ is a solution to the system of equations

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \hat{s}_{BORS}} = 0 \\ \hat{s}_{BORS} = \int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) \\ \int_{\mathcal{B}} dF(v, \theta) = K \end{cases}$$

Figure 3.3 shows a graphical illustration of the equilibrium product assortment with BORS. This result can be interpreted in a manner similar to the main model. The BMR stocks a product if the marginal benefit exceeds the marginal cost. With the BORS agreement, consumers never visit the BMR first. So the marginal benefit of stocking products that are available online is 0 because consumers purchase these products online. The marginal benefit of stocking a product (v, θ) that is not available online also can be decomposed: The first *direct effect* captures the profit gain from selling the product $(\theta D(v, p^*(v))p^*(v)G(\hat{s}_{BORS}))$, and the second *indirect effect* captures the benefit from attracting more consumers to the store $(\lambda_2 u(v, \theta))$. If the shelf space capacity is not binding, then the opportunity cost of stocking drops to zero, so all products that are unavailable online get stocked in store. Otherwise, the BMR has an opportunity cost of $\lambda_1 > 0$ for stocking an additional product.

Similar to the main model, the equilibrium is too complicated to express in closed

forms, even with simple function forms on the demand function and search cost distribution. We obtain closed-form solutions by using a degenerating, one-dimensional product variety space. See the Appendix for the explicitly solvable example.

3.5.2 The impact and likelihood of a BORS agreement

When there is a BORS agreement, handling returns becomes less costly for the OR. So some products that have a higher return probability (lower θ) will become available online. By Lemma 3.3 and Lemma 3.4, with a BORS agreement, all consumers prioritize shopping online because visiting the BMR second can save the hassle cost of returning the product to the OR. Thus, the BMR no longer competes to be consumers' first-stop shopping destination and shifts to exclusively selling products that are not available online. Comparing the optimal product assortments in Proposition 3.3 and Lemma 3.4, we have the following result.

Proposition 3.5. *Suppose the two retailers reach a BORS agreement. In the long run, the two retailers fully differentiate and sell exclusive products. The OR increases product variety to sell products that have lower fit probability, and the BMR moves to stock only niche products that have high value. In aggregate, a larger variety of products is available.*

We predict that the agreements such as BORS would hasten BMRs' transformation into specialty stores. But the fully differentiated outcome may seem somewhat implausible. To reconcile this with a more plausible outcome, it is useful to embed in the model a small fraction of consumers who never shop online (see the extension section). As suggested by the Pew Research Center in 2019, about 10% of the U.S. population did not use the Internet at all. The existence of consumers who never buy online could give an upper hand to the BMR and somewhat dampen this full differentiation result. This has a moderating effect on the equilibrium assortment outcomes of BORS. However, in many industries, COVID-19 was a shock that has transformed the retail landscape; as such, the number of consumers who buy exclusively in physical stores is likely to diminish, and our results on assortment decisions are likely to be of greater practical significance.

When the two retailers offer BORS, in the short-run (fixing the equilibrium assortments without BORS), the BMR gains more consumer foot traffic by offering returns. However, the BMR faces a stronger online competitor in the long run, because the OR

will expand its product variety. Therefore, in the long run, it is not obvious how the BMR's store foot traffic will change and whether it still benefits from offering BORS. The next result shows that, in the long-run, BORS increases the BMR's store foot traffic if its store capacity is small.

Proposition 3.6. *Suppose the two retailers reach a BORS agreement. Then consumers purchase more products online and return more online purchases.*

1. *In the short-run (fixing the equilibrium assortments without BORS), BORS increases the BMR's store foot traffic.*
2. *In the long-run (after product assortments adjust), BORS increases the BMR's store foot traffic if its store capacity is small.*

Finally, we consider retailer profits and lay out the conditions under which a BORS agreement is likely to be reached between the two retailers. The OR always benefits from a BORS agreement because handling in-store returns is less costly, and all consumers shop online and buy (weakly) more products when BORS is available. The BMR is harmed by consumers shifting to OR as their first-stop shopping destination, but it can benefit from an increase in consumer traffic. In particular, if the store capacity is small, the BMR alone would attract only a small fraction of the consumers that it attracts with a BORS agreement in place. So the profit-enhancing effect (by increasing foot traffic) of the BORS agreement dominates. In this case, allowing returns to the BMR is mutually beneficial for the two retailers. Alternatively, if the hassle of returning products directly to ORs is large for consumers and an OR's cost of handling returns via a BMR is small, then the BORS agreement significantly improves overall market efficiency. Under such a scenario, the OR could benefit significantly and strongly favors reaching a BORS agreement and would offer to even monetarily compensate the BMR for its potential losses.

Proposition 3.7.

1. *If the BMR's store capacity is small, then the BORS agreement is mutually beneficial for the two retailers.*
2. *If r is large enough and \tilde{R} is small enough, then the aggregate retail profit is higher with the BORS agreement. So the OR should be willing to compensate the BMR for allowing returns to the store.*

As anecdotal evidence to the first part of Proposition 3.7, we often observe ORs collaborating with non-dominant BMRs: Amazon.com purchases can be returned to the smaller department stores Kohl's, not Target; Amazon Japan purchases can be returned with Yu-Pack service, which is available at smaller convenience stores Lawson and Ministop but not at larger ones such as 7-Eleven and FamilyMart; Taobao.com courier stations are often located in mom and pop stores instead of in larger supermarkets.

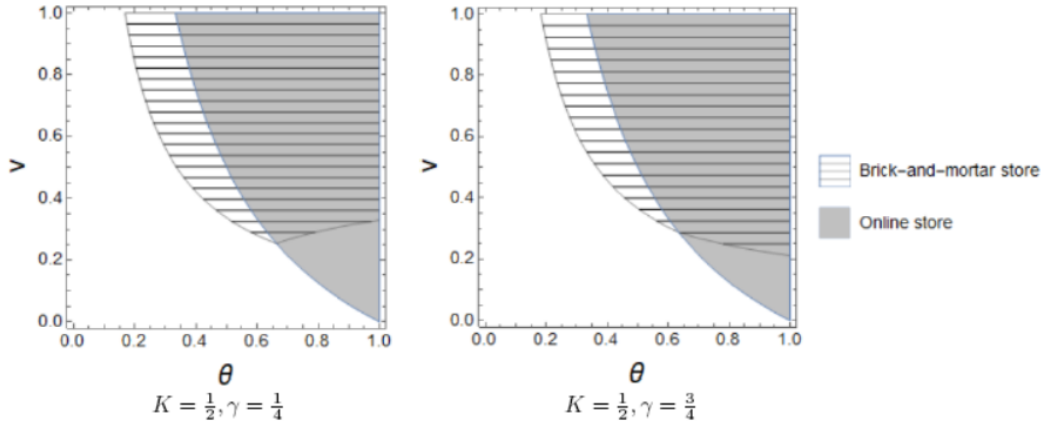
3.6 Extensions

3.6.1 Consumers without access to the OR

In our model so far, we have assumed all consumers have equal access to and preferences for BMRs and ORs. Even though online purchasing is becoming increasingly popular, as of 2019, 10% of the U.S. population did not use the internet at all. A significant literature in marketing has studied the issues related to state dependence, inertia, and loyalty on marketplace outcomes (Seetharaman et al., 1999; Chintagunta, 1998; Shin and Sudhir, 2010), and as such, our assumption of consumers being indifferent to the channel preference seems somewhat restrictive. Would the existence of a segment of consumers loyal to BMRs invalidate our analysis? In the short-run, the BMR's loyal consumers do not shop online and are unaffected by the BORS collaboration. But happens in the long-run, after the BORS collaboration changes the retail landscape?

In this extension, we model a segment of consumers who never shop online. We assume that a fraction of γ of the consumers can only choose between visiting the BMR or taking their outside option, while the remaining $1 - \gamma$ consumers behave the same as in our main model.

We first consider the case without a BORS agreement. The non-loyal consumers make the same shopping decisions as in Proposition 3.1. In contrast, the loyal consumers visit the BMR if their expected utility from shopping exceeds travel costs, and otherwise, they take their outside option. Furthermore, loyal consumers have higher incentives to visit the BMR (i.e., $\hat{s}_l \geq \hat{s}$). Without BORS, the OR sells the same products as in the main model. The BMR behaves more like a monopolist in the presence of more loyal consumers and has a higher incentive to stock products that have high v and high θ if they are available online. However, as its loyal consumers decrease, the BMR assortment decision reverts back to the main model and displays a stronger preference for products



Note. This figure plots the equilibrium product assortments in the presence of loyal consumers. As the fraction of loyal consumers (γ) decreases, the equilibrium converges back to the main model (Figure 3.2). In this plot, product variety (v, θ) is jointly uniform on $[0, 1] \times [0, 1]$, consumers' travel cost $s \stackrel{iid}{\sim} U[0, 1]$, demand function $D(v, p) = 8v(1 - p)$, and consumers' marginal cost of returning $r = \frac{1}{2}$.

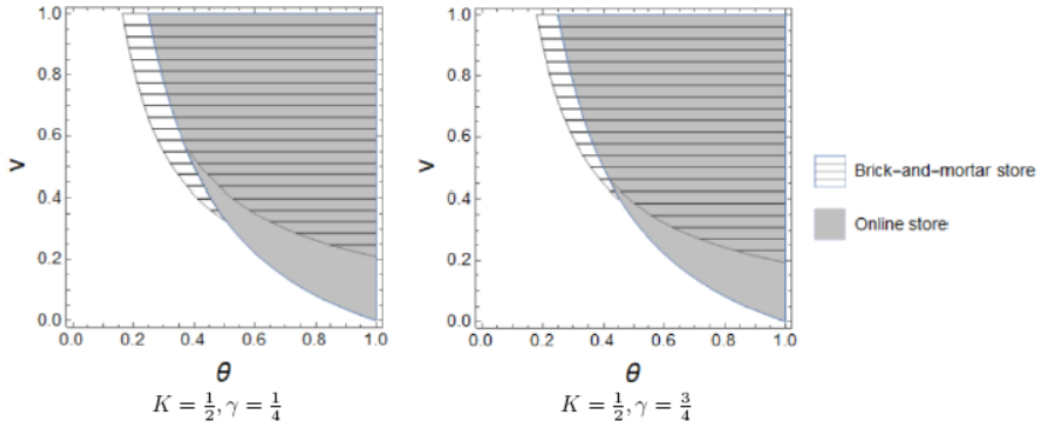
Figure 3.4: Optimal product assortment in the presence of loyal consumers.

that have low θ if the products are available online.

Then, we analyze the effect of offering consumers the BORS option. In the short-run (before assortment decisions adjust), non-loyal consumers use the same shopping strategy as in Lemma 3.3, and loyal consumers are not affected by the return policy and use the same shopping strategy as before. In the long run, BORS will cause differentiation, but we should not expect to see full differentiation of the two retailers if the fraction of consumers who are loyal to the BMR is significant. After the store positionings adjust, loyal consumers will be harmed due to the BMR's niche selection of products.

Proposition 3.8. *BORS collaboration leads to product differentiation, but the two retailers will not fully differentiate when the BMR has loyal consumers. In the long run, BORS hurts the BMR's loyal consumers. Furthermore, the smaller the loyal segment, the more severe is the negative impact on each loyal consumer.*

This has important implications on financial inclusion. The loyal consumers of the BMR can be interpreted as low income or rural consumers without access to the internet, electronic devices, and/or bank accounts. BORS causes the BMR to specialize, and the niche product selection can hurt these minority groups without access to the OR. When the minority group size shrinks, the BMR specializes even more, and each consumer without access to the OR becomes worse-off. Therefore, BORS may lead to concerns



Note. This figure plots the equilibrium product assortments in the presence of loyal consumers and a BORS agreement. As the fraction of loyal consumers (γ) decreases, the equilibrium converges to that of Figure 3.3. In this plot, product variety (v, θ) is jointly uniform on $[0, 1] \times [0, 1]$, consumers' travel cost $s \stackrel{iid}{\sim} U[0, 1]$, demand function $D(v, p) = 8v(1 - p)$, consumers' marginal cost of returning $r = \frac{1}{2}$, and the OR's cost of handling returns to the BMR is $\tilde{R} = \frac{1}{3}$.

Figure 3.5: Optimal product assortment in the presence of loyal consumers and a BORS agreement.

from the standpoint of equity and inclusion, especially when the population of consumers without internet access is small.

3.6.2 Price competition

In the main model, we focused on the perfect Bayesian equilibrium where the two retailers do not compete on prices. In this subsection, we reexamine the equilibrium selection and investigate whether there is an outcome with price competition and how it would alter our results.

As suggested in Proposition 3.10 (see the Appendix), in equilibrium, price differences across retailers are only possible when consumers have a strictly 0 cost of searching online. Technically, the game displays a discontinuity in consumers' cost of going online - any perturbations of the game (so that searching online has a small positive cost) leads to a unique equilibrium with monopoly pricing by both firms, but the limiting case (zero cost of searching online) has multiple equilibria. In the following extension, we investigate the limit case where searching online has cost exactly equal to 0 and show our BMR's specialization result is robust.

Let p_{BMR} and p_{OR} be the respective prices set by the BMR and OR. We show that

there is a unique cut-off \tilde{v} such that the OR never sells a product $(v, \theta) \in \mathcal{B}$ if $v \leq \tilde{v}$. Thus, on equilibrium path, we can restrict attention to $v > \tilde{v}$ for overlapping products, $(v, \theta) \in \mathcal{B} \cap \mathcal{O}$.⁹ Lemma 3.7 in the Appendix shows that overlapping products have equilibrium prices that are random and share the same support, $p_{OR}(v, \theta), p_{BMR}(v, \theta) \in [\underline{p}(v, \theta), \bar{p}(v, \theta)]$. Here, $\bar{p}(v, \theta) = \min\{p^*(v), \hat{p}(v, \theta)\}$, $\hat{p}(v, \theta)$ is the solution to $u(v, \theta, \hat{p}) - (1 - \theta)r = 0$, $p^*(v) = \arg \max_p \theta D(v, p)p$ is the monopoly price, and $\underline{p}(v, \theta)$ as the smaller solution to $\theta D(v, \underline{p}(v, \theta))\underline{p}(v, \theta) = [1 - G(\tilde{s})]\theta D(v, \bar{p}(v, \theta))\bar{p}(v, \theta)$.¹⁰

The next result presents the equilibrium prices.

Proposition 3.9. *There is a unique mixed strategy equilibrium where two retailers compete on prices.*

1. *BMR exclusive products $((v, \theta) \in \mathcal{B} \cap \mathcal{O}^c)$ and OR exclusive products $((v, \theta) \in \mathcal{O} \cap \mathcal{B}^c)$ are priced at the monopoly level $p(v, \theta) = p^*(v)$.*
2. *Overlapping products $((v, \theta) \in \mathcal{B} \cap \mathcal{O})$ have random prices with CDFs given by*

$$H_{BMR}(p) = \begin{cases} 1, & p \geq \bar{p}(v, \theta) \\ \frac{D(v, p)p - D(v, \underline{p}(v, \theta))\underline{p}(v, \theta)}{G(\tilde{s})D(v, p)p}, & \underline{p}(v, \theta) \leq p < \bar{p}(v, \theta) \\ 0, & p \leq \underline{p}(v, \theta) \end{cases}$$

and

$$H_{OR}(p) = \begin{cases} 1, & p \geq \bar{p}(v, \theta) \\ \frac{D(v, p)p - D(v, \underline{p}(v, \theta))\underline{p}(v, \theta)}{D(v, p)p}, & \underline{p}(v, \theta) \leq p < \bar{p}(v, \theta) \\ 0, & p \leq \underline{p}(v, \theta) \end{cases}$$

In this equilibrium with price competition, consumer behavior is similar to Proposition

⁹Off equilibrium path, if there is a product with $v \leq \tilde{v}$ that are sold by both retailers (i.e., $(v, \theta) \in \mathcal{B} \cap \mathcal{O}$ and $v \leq \tilde{v}$), then $p_{OR}(v, \theta) = p_{BMR}(v, \theta) = p^*(v)$ is the unique equilibrium price. Intuitively, the OR cannot compete on prices, because of a net returns loss from its captive consumers. So, the BMR does not face a downwards pressure on pricing.

¹⁰This equilibrium with price competition is similar to Varian (1980); Armstrong and Vickers (2021). Varian (1980) considers a symmetric environment where the two retailers have the same number of captive consumers. Due to the asymmetric nature of online-offline competition, we need to extend the framework of Varian (1980). Armstrong and Vickers (2021) also allow for asymmetric environments, but they allow an arbitrary number of firms and only consider unit-demand consumers. In contrast, this model has only two retailers and works with a more general demand function.

3.1. We have a slightly different cut-off consumer given by

$$\tilde{s} = \int_{\mathcal{B}} \mathbb{E}[u(v, \theta, p_{BMR})] dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [\mathbb{E}[u(v, \theta, p_{\min})] - (1 - \theta)r]^+ dF(v, \theta),$$

where expectations are taken over random prices, and $p_{\min} = \min\{p_{BMR}, p_{OR}\}$. Consumers with $s > \tilde{s}$ optimally choose to only shop online, while consumers with $s \leq \tilde{s}$ shops in-store first and online second.

The two retailers randomize over the prices of overlapping products, $(v, \theta) \in \mathcal{B} \cap \mathcal{O}$. Compared to Proposition 3.2, price competition decreases the direct effect of selling a product, because sales profit decreases from the monopoly level $\theta D(v, p^*)p^*$ to $\mathbb{E}_{p_{BMR}}[\theta D(v, p_{BMR})p_{BMR}]$. However, the indirect effect of selling a product does not change - consumers save the same expected return hassle $(1 - \theta)r$ from in-store inspection. Therefore, price competition strengthens the key finding of Proposition 3.3 - even with price competition, in equilibrium, the BMR has a greater incentive to specialize and sell products with *low* fit probability, θ .

3.7 Conclusion

In this paper, we investigate the competitive product assortment decisions of online and offline retailers. In the main model, we find that the OR sells high value and easy-to-fit products to minimize costly returns and to maximize profit. More interestingly, the BMR might use its limited store size to stock products that have a lower fit probability. This counterintuitive result arises because of a subtle trade-off between generating higher sales and attracting more foot traffic. The results support analysts' view of using specialization as a shield against online platforms, and shed light on the steady growth of specialty stores in the downfall of traditional BMR.

Our model provides new insights into the nascent BORS policy. This policy might alter consumers' shopping strategy, and we show that a broader implication of such policies is that they can push retailers toward product differentiation. Consequently, more products will be made available in the market. We show that the OR always benefits from allowing returns to a physical store, and it may want to compensate the BMR for the collaboration. However, in the absence of a monetary transfer, the return agreement is still mutually beneficial for both the retailers if the BMR's store size is small. This provides a rationale as to why some ORs (e.g., Amazon.com and Amazon

Japan) partner with smaller BMRs instead of larger ones. While BORS can benefit both retailers and prevail in the long run, it can also hurt low income consumers who have limited access to the internet and/or digital payment methods. This may lead to concerns from the standpoint of financial equity and inclusion.

This paper also makes a methodological contribution to modeling product assortment decisions in a competitive setting. Previous work on product assortment decisions typically assumes homogeneous goods and perfect competition (Spence, 1976; Dixit and Stiglitz, 1977; Bronnenberg, 2015, 2020), a single retailer (Rhodes, 2015; Rhodes et al., 2020), or two retailers facing two different products (e.g. Dukes et al., 2009). We contribute to this literature by studying two competing retailers facing a continuum of products that are differentiated in two dimensions: value and fit probability. Our method can also be applied to study other duopoly markets where products display (multiple) other characteristics.

This chapter of the dissertation is directly taken from the Fall 2021 edition of our working paper. Our paper represents the first attempt to model the impact of consumer returns on product assortment decisions in online and offline competition. Two new components will be added to this paper in the near future. First, many BMRs have online stores. Although the revenue generated from these new channels is still insignificant,¹¹ many traditional BMRs are actively expanding their online presence. We will incorporate an analysis of this into our paper. Second, we will also add an extension of price competition when retailers face heterogeneous marginal costs of selling. However, we still leave many questions open. First, including manufacturers within our setup and analyzing how they strategically affect assortment decisions may be an interesting extension of our work. Second, one can further study how online-offline competition and BORS agreements impact the logistics and reverse logistics in the supply chain. Third, we assume that the hassle cost of returning online purchases is moderate. This leads to a dominant strategy for the OR's product assortment choice. Relaxing this assumption is interesting because the OR may attract consumers by offering products with negative expected profit. But showing the existence of an equilibrium and solving for the equilibrium product assortment with these generalizations can be quite challenging.

¹¹For example, online sales account for less than 2% of revenue at Walmart and Target.

3.8 Appendix

3.8.1 Proofs for the main model

3.8.1.1 Proofs for the benchmark: Only BMR present

Proof of Lemma 3.1. Given the BMR's pricing strategy, a consumer's benefit of visiting the BMR is $\int_{\mathcal{B}} u(v, \theta) dF(v, \theta)$. Therefore, if $s \leq \int_{\mathcal{B}} u(v, \theta) dF(v, \theta)$, the consumer obtains non-negative utility from visiting the BMR. Otherwise, the consumer takes her outside option 0.

Given consumers' shopping strategy, it is optimal to set price $p^*(v)$ for product (v, θ) . This is because, by deviating, the BMR does not change foot traffic and generates less profit from each unit of sale. □

Proof of Lemma 3.2. The proof of this lemma follows the same steps as in the proof of Proposition 3.2.

(Step 1) Let $q(v, \theta) \in \{0, 1\}$ denote the brick-and-mortar store's stocking decision for product (v, θ) . Then, the optimization can be rewritten as

$$\begin{aligned} \max_{q(v, \theta) \in \{0, 1\}} & \int \theta D(v, p^*(v)) p^*(v) q(v, \theta) dF(v, \theta) \times G(\hat{s}_{bench}) \\ \text{s.t. } & \hat{s}_{bench} = \int u(v, \theta) q(v, \theta) dF(v, \theta) \\ & \int q(v, \theta) dF(v, \theta) \leq K \end{aligned}$$

(Step 2) We show that a solution to the BMR's profit maximization problem exists. This existence proof is analogous to that in the proof of Proposition 3.2.

(Step 3) By the same arguments as in the proof of Proposition 3.2, the brick-and-mortar retailer's store capacity must bind in equilibrium.

(Step 4) We can write a "Lagrangian" for the optimization problem:

$$\begin{aligned} \mathcal{L} = & \int \theta D(v, p^*(v)) p^*(v) q(v, \theta) dF(v, \theta) \times G(\hat{s}_{bench}) + \lambda_1 [K - \int q(v, \theta) dF(v, \theta)] \\ & + \lambda_2 [\int u(v, \theta) q(v, \theta) dF(v, \theta) - \hat{s}_{bench}] \end{aligned}$$

Collect all terms with $q(v, \theta)$ and separate it out with other terms, we get,

$$\begin{aligned} \mathcal{L} = & \int [\theta D(v, p^*(v)) p^*(v) G(\hat{s}_{bench}) - \lambda_1 + \lambda_2 u(v, \theta)] q(v, \theta) dF(v, \theta) \\ & + \lambda_1 K - \lambda_2 \hat{s}_{bench} \end{aligned}$$

(Step 5) As in Rhodes et al. (2020), since the Lagrangian is linear in $q(v, \theta)$, the optimum product assortment should have $q(v, \theta) = 1$ if the coefficient is positive, and $q(v, \theta) = 0$ if negative. Thus, brick-and-mortar store stocks a product (i.e., $q(v, \theta) = 1$) iff

$$\theta D(v, p^*(v)) p^*(v) G(\hat{s}_{bench}) + \lambda_2 u(v, \theta) \geq \lambda_1$$

Plugging in the optimal $q(v, \theta)$, we can solve $\{\hat{s}_{bench}, \lambda_1, \lambda_2\}$ from the system of equations

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \hat{s}_{bench}} = 0 \\ \hat{s}_{bench} = \int u(v, \theta) q(v, \theta) dF(v, \theta) \\ \int_{[v, \bar{v}] \times [\theta, \bar{\theta}]} q(v, \theta) dF(v, \theta) = K \end{cases}$$

□

Notice that both $\theta D(v, p^*(v)) p^*(v)$ and $u(v, \theta)$ increase in v and θ . Thus, the BMR prefers to stock products with high v and high θ when it does not face online competition.

3.8.1.2 Proofs for the main model with online-offline competition

Proof of Proposition 3.1. We show that the strategies in Proposition 3.1 are an equilibrium in two steps.

(Step 1) We take as given the retailers' equilibrium prices $p(v, \theta) = p^*(v)$ and derive the consumers' optimal shopping strategy. By assumption, consumers' beliefs are passive and correct equilibrium path, so $p^e(v, \theta) = p^*(v)$.¹²

(Step 1.1) First, observe that only visiting the BMR is suboptimal because visiting the OR is costless.

(Step 1.2) Second, observe that if a consumer plans to visit both the BMR and the OR, then visiting the BMR first is optimal. The reason is that if $\mathcal{B} \cap \mathcal{O} \neq \emptyset$, then the consumer can physically inspect these products in the BMR and save the expected hassle cost of returning online purchases, $\int_{\mathcal{B} \cap \mathcal{O}} (1 - \theta) r dF(v, \theta) > 0$.

¹²Passive beliefs is assumed in both rational expectations equilibrium and sequential search works. In sequential search, this implies that if a consumer is surprised by one retailer's deviation, the consumer's belief on the other retailer's prices does not change.

We calculate the expected utility of visiting the BMR first and visiting the OR second. In this case, the consumer obtains expected utility $\int_{\mathcal{B}} u(v, \theta, p^e(v, \theta)) dF(v, \theta) = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta)$ from the BMR and pays travel cost s . When the consumer visits the online store, he optimally searches for the products that give a non-negative expected utility (i.e., products that satisfy $u(v, \theta, p^e(v, \theta)) - (1 - \theta)r = u(v, \theta) - (1 - \theta)r \geq 0$). The consumer purchases all of these products and obtains $\int_{\mathcal{O} \cap \mathcal{B}^c} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta)$ from the online store. Thus, the total expected utility from this shopping trip is

$$\int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - s + \int_{\mathcal{O} \cap \mathcal{B}^c} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta).$$

(Step 1.3) Third, we calculate the expected utility of visiting only the OR. In this case, the consumer optimally searches online for the products that give a non-negative expected utility (i.e., products that satisfy $u(v, \theta, p^e(v, \theta)) - (1 - \theta)r = u(v, \theta) - (1 - \theta)r \geq 0$). The consumer purchases all of these products and obtains expected utility

$$\int_{\mathcal{O}} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta).$$

(Step 1.4) Fourth, we compare the two undominated strategies in (Step 1.2) and (Step 1.3). Direct calculation shows that if $s < \hat{s} = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta)$, then the consumer visits the BMR first and the OR second. Otherwise, the consumer visits only the online store.

(Step 2) Now, we take as given consumers' optimal shopping strategies and show that the optimal approach is to set price $p(v, \theta) = p^*(v)$.

(Step 2.1) First, given the consumers' shopping strategy, the BMR obtains profit $\int_{\mathcal{B}} \theta D(v, p) p dF(v, \theta) \times G(\hat{s})$. From Step 1, \hat{s} depends only on $p^e(v, \theta)$ and is constant in $p(v, \theta)$. Furthermore, $\theta D(v, p) p$ is maximized at $p(v, \theta) = p^*(v)$. Therefore, the BMR optimally sets price $p(v, \theta) = p^*(v)$ for product (v, θ) .

(Step 2.2) Second, given the consumers' shopping strategy, the OR obtains profit $\int_{\mathcal{O} \cap \{(v, \theta) | u(v, \theta) - (1 - \theta)r \geq 0\}} D(v, p) p - (1 - \theta)R dF(v, \theta) \times [1 - G(\hat{s})] + \int_{\mathcal{O} \cap \mathcal{B}^c \cap \{(v, \theta) | u(v, \theta) - (1 - \theta)r \geq 0\}} D(v, p) p - (1 - \theta)R dF(v, \theta) \times G(\hat{s})$, where the first term is the profit obtained from consumers with $s \geq \hat{s}$ and the second term is the profit obtained from consumers with $s < \hat{s}$. By the same logic as before, the OR optimally sets price $p(v, \theta) = p^*(v)$ for each product (v, θ) , as long as selling the product does not result in a net loss due to returns (i.e., $D(v, p^*(v)) p^*(v) - (1 - \theta)R \geq 0$). Due to the assumption $r \geq \underline{r}$, we have that $u(v, \theta) - (1 - \theta)r \geq 0$ implies $D(v, p^*(v)) p^*(v) - (1 - \theta)R \geq 0$. Thus, as long as consumers

are willing to buy, the OR can make a non-negative profit at price $p^*(v)$. Therefore, setting price $p^*(v)$ for all products in \mathcal{O} is optimal.

We have shown that the pricing strategy and consumer search strategy forms a best response cycle. Hence, it is a perfect Bayesian equilibrium. It should be noted that this proof did *not* ruled out the existence of other equilibria. See the Appendix for the discussion on equilibrium selection. \square

Proof of Proposition 3.2. We first derive the OR's product assortment and then derive the BMR's product assortment.

(Step 1) We show that the OR optimally chooses $\mathcal{O} = \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}$. First, consumers only search for products $\{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}$ online. Hence, $\mathcal{O} \subset \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}$. Given the assortment of the BMR, the optimization of OR can be written as

$$\begin{aligned} \max_{\mathcal{O} \subset \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}} & \int_{\mathcal{O}} \theta D(v, p^*(v)) p^*(v) - (1 - \theta)R dF(v, \theta) \times [1 - G(\hat{s}_1)] \\ & + \int_{\mathcal{O} \cap \mathcal{B}^c} \theta D(v, p^*(v)) p^*(v) - (1 - \theta)R dF(v, \theta) \times G(\hat{s}_1) \\ \text{s.t. } \hat{s}_1 & = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta) \end{aligned}$$

Using proof by contradiction, suppose that $\mathcal{O} \subsetneq \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}$ and $\{(v, \theta) \notin \mathcal{O} \mid u(v, \theta) - (1 - \theta)r \geq 0\}$ has a positive measure. Then, by our assumption $r \in [\max_v \frac{\int_{p^*(v)}^{\infty} D(v, p) dp}{D(v, p^*(v)) \times p^*(v)} R, p^*(\underline{v})]$, stocking all the products $\{(v, \theta) \notin \mathcal{O} \mid u(v, \theta) - (1 - \theta)r \geq 0\}$ strictly increases profits from sales and attracts more consumers to visit the OR first (i.e., \hat{s} strictly decreases). Thus, it is optimal to set $\mathcal{O} = \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}$.

(Step 2) Now, we derive the BMR's optimal product assortment.

(Step 2.1) We write out the BMR's profit maximization problem:

$$\begin{aligned} \max_{\mathcal{B} \subset [\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} & \int_{\mathcal{B}} \theta D(v, p^*(v)) p^*(v) dF(v, \theta) \times G(\hat{s}) \\ \text{s.t. } \hat{s} & = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta) \\ & \int_{\mathcal{B}} dF(v, \theta) \leq K \end{aligned}$$

Here, $u(v, \theta) = \theta \int_{p^*(v)}^{\infty} D(v, p^*(v)) dv$. Let $q(v, \theta) \in \{0, 1\}$ denote the BMR's stocking

decision for product (v, θ) . Then, the optimization can be rewritten as

$$\begin{aligned} & \max_{q(v, \theta) \in \{0, 1\}} \int \theta D(v, p^*(v)) p^*(v) q(v, \theta) dF(v, \theta) \times G(\hat{s}) \\ & \text{s.t. } \hat{s} = \int u(v, \theta) q(v, \theta) dF(v, \theta) - \int_{\mathcal{O}} [u(v, \theta) - (1 - \theta)r] q(v, \theta) dF(v, \theta) \\ & \int q(v, \theta) dF(v, \theta) \leq K \end{aligned}$$

This converts the problem from optimizing over sets into optimizing over functions.

(Step 2.2) We show that a solution to the BMR's profit maximization problem exists. For this existence proof, we define the notation $\pi_{q(v, \theta)}$ to be the BMR's profit when it chooses a stocking function $q(v, \theta)$. Letting π^* denote the supreme of the optimization problem, we can see that $\pi^* \leq \int \theta D(v, p^*(v)) p^*(v) dF(v, \theta) < \infty$. By definition of the supreme, there exists a sequence of functions $\{q_n(v, \theta)\}_{n=1}^{\infty}$, such that $\pi_{q_n(v, \theta)} \rightarrow \pi^*$. There must be a subsequence $\{q_{n_j}(v, \theta)\}_{j=1}^{\infty}$ such that the stocking function converges pointwise almost everywhere (i.e., for almost all points except a set of measure zero); otherwise, $\pi_{q_n(v, \theta)} \rightarrow \pi^*$ will be violated. Denote the limit of this sub-sequence as $q^*(v, \theta)$. By continuity, we have $\pi_{q_{n_j}(v, \theta)} \rightarrow \pi_{q^*(v, \theta)}$, so $\pi_{q^*(v, \theta)} = \pi^*$. The limit stocking function $q^*(v, \theta)$ satisfies both of the constraints because equalities and weak inequalities are preserved by limits. Hence, the $q^*(v, \theta)$ is an optimal solution.

This proves the existence of an optimum, and allows us to focus on the first order necessary conditions.

(Step 2.3) At the optimum, the constraint $\int q(v, \theta) dF(v, \theta) \leq K$ must hold with equality. Otherwise, the BMR can stock any remaining products until the store size capacity is reached. Doing so will attract more consumers (i.e., \hat{s} increase) and allow the BMR to make more profit off the new products.

(Step 2.4) We can write a "Lagrangian" for the optimization problem:

$$\begin{aligned} \mathcal{L} = & \int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} \theta D(v, p^*(v)) p^*(v) G(\hat{s}) q(v, \theta) dF(v, \theta) + \lambda_1 [K - \int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} q(v, \theta) dF(v, \theta)] \\ & + \lambda_2 [\int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} u(v, \theta) q(v, \theta) dF(v, \theta) - \int_{\mathcal{O}} [u(v, \theta) - (1 - \theta)r] q(v, \theta) dF(v, \theta) - \hat{s}] \end{aligned}$$

Collecting all terms with $q(v, \theta)$ and separating them out from other terms, we get:

$$\begin{aligned} \mathcal{L} &= \int_{\mathcal{O}} [\theta D(v, p^*(v)) p^*(v) G(\hat{s}) - \lambda_1 + \lambda_2(1 - \theta)r] q(v, \theta) dF(v, \theta) \\ &\quad + \int_{\mathcal{O}^c} [\theta D(v, p^*(v)) p^*(v) G(\hat{s}) - \lambda_1 + \lambda_2 u(v, \theta)] q(v, \theta) dF(v, \theta) + \lambda_1 K - \lambda_2 \hat{s} \end{aligned}$$

(Step 2.5) As in Rhodes et al. (2020), because the Lagrangian is linear in $q(v, \theta)$, the optimum product assortment should have $q(v, \theta) = 1$ if the coefficient is positive and should have $q(v, \theta) = 0$ if the coefficient is negative. Thus, BMR stocks a product (i.e., $q(v, \theta) = 1$) iff:

$$\begin{cases} \theta D(v, p^*(v)) p^*(v) G(\hat{s}) + \lambda_2(1 - \theta)r \geq \lambda_1, & (v, \theta) \in \mathcal{O} \\ \theta D(v, p^*(v)) p^*(v) G(\hat{s}) + \lambda_2 u(v, \theta) \geq \lambda_1, & (v, \theta) \notin \mathcal{O} \end{cases}$$

Plugging in the optimal $q(v, \theta)$, $\{\hat{s}, \lambda_1, \lambda_2\}$ can be solved from the system of equations:

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \hat{s}} = 0 \\ \hat{s} = \int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} u(v, \theta) q(v, \theta) dF(v, \theta) - \int_{\mathcal{O}} [u(v, \theta) - (1 - \theta)r] q(v, \theta) dF(v, \theta) \\ \int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} q(v, \theta) dF(v, \theta) = K \end{cases}$$

□

Proof of Proposition 3.3. By Proposition 3.2, the OR's assortment is $\mathcal{O} = \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}$. Since $u(v, \theta)$ increases in θ and $(1 - \theta)r$ decreases in θ , the OR stocks a product if θ is high enough. This proves the first part of Proposition 3.3.

By Proposition 3.2, the BMR's assortment is given by

$$\begin{cases} \theta D(v, p^*(v)) p^*(v) G(\hat{s}) + \lambda_2(1 - \theta)r \geq \lambda_1, & (v, \theta) \in \mathcal{O} \\ \theta D(v, p^*(v)) p^*(v) G(\hat{s}) + \lambda_2 u(v, \theta) \geq \lambda_1, & (v, \theta) \notin \mathcal{O} \end{cases}$$

For products not sold online, the direct (first term on the left hand side) and indirect (second term on the left hand side) effects both increase in v and θ . So, the BMR offers products with high v and high θ in this region. For products sold online, the direct and indirect effects both still increase in v . So, the BMR offer these products only when value is high enough. But, in this case, the direct effect increases in θ and the

indirect effect decreases in θ . So, depending on which effect dominates, the BMR can prefer high or low θ in this region. □

Proof of Proposition 3.4. To show prove Proposition 3.4, we consider a small neighborhood around $(v, \theta) = (0, 1)$ in the product variety space. By Proposition 3.2, the OR stocks products in this neighborhood, because $\mathcal{O} = \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\} = \{(v, \theta) \mid \theta \geq \frac{r}{r-u(v, \theta)}\} \ni (0, 1)$.

However, the BMR always avoids this region as long as $K < 1$. It does so because, from Proposition 3.2, the marginal benefit of stocking the product is close to 0 near $(v, \theta) = (0, 1)$, while the marginal opportunity cost λ_1 is positive as long as store size capacity is binding. □

3.8.1.3 Equilibrium selection for the last two stages

We define an equivalence condition for two equilibria. In words, we say that an equilibrium is *essentially unique* if, in any other equilibria, the retailers' pricing strategies only differ on a set of products that has measure zero. This idea is formalized below.

Definition 3.1. Given the assortment decisions in period 1, \mathcal{B}, \mathcal{O} , consider a perfect Bayesian equilibrium of the period 2 with retailer prices $p_{BMR}^1(v, \theta)$ and $p_{OR}^1(v, \theta)$. This equilibrium is *essentially unique* if for any other perfect Bayesian equilibrium with prices $p_{BMR}^2(v, \theta)$ and $p_{OR}^2(v, \theta)$, we have

$$\int_{\mathcal{B}} \mathbb{1}_{p_{BMR}^1(v, \theta) \neq p_{BMR}^2(v, \theta)} dF(v, \theta) + \int_{\mathcal{O}} \mathbb{1}_{p_{OR}^1(v, \theta) \neq p_{OR}^2(v, \theta)} dF(v, \theta) = 0.$$

Now we perturb the last two stages of the main model to introduce a small cost of browsing online, and show that the equilibrium presented in Proposition 3.1 is essentially unique.

Proposition 3.10. *When consumers have an arbitrarily small but positive marginal cost of searching online, $\epsilon > 0$, the perfect Bayesian equilibrium described in Proposition 3.1 is essentially unique.*

Proof. For this proof, let us introduce two notations. Let $p_{BMR}(v, \theta)$ be the price the BMR sets for (v, θ) , and let $p_{OR}(v, \theta)$ be the price the OR sets for (v, θ) . The

proof repeatedly exploits the hold-up problem to show that downwards price pressure disappears with costly consumer search. The main idea is similar to the Diamond paradox: With a small cost of searching online, the OR needs to undercut the BMR's price until $p_{OR} + \epsilon \leq p_{BMR}$ to compete for consumers; But this cannot be an equilibrium because, the moment a consumer sink the search cost ϵ and observes p_{OR} , the OR has incentive to increase price, making the consumer search suboptimal.

First, we use proof by contradiction to show that the two firms charge the same price $p^*(v)$ in the region $\mathcal{O} \cap \mathcal{B}$, so consumers do not cherry-pick for the lowest price in equilibrium. (Case 1) That a retailer i never charges price $p_i(v, \theta) > p^*(v)$ is trivial because lowering the price to $p^*(v)$ increases profits, given consumers' beliefs and shopping strategies. (Case 2) Suppose that $p_{BMR}(v, \theta) < p_{OR}(v, \theta) + \epsilon \leq p^*(v)$ for a positive measure of product $(v, \theta) \in \mathcal{O} \cap \mathcal{B}$. Since consumers hold beliefs $p_{BMR}^e(v, \theta) < p_{OR}^e(v, \theta) + \epsilon$, searching online for (v, θ) is not worthwhile if the consumer plans to visit the BMR. Given the optimal shopping strategy induced by this belief, the BMR will have incentive to deviate to charge a slightly higher price (while still maintaining $p_{BMR}(v, \theta) < p_{OR}^e(v, \theta) + \epsilon = p_{OR}(v, \theta) + \epsilon$) and make more profit. (Case 3) Suppose that $p_{OR}(v, \theta) + \epsilon < p_{BMR}(v, \theta) \leq p^*(v)$ for a positive measure of product $(v, \theta) \in \mathcal{O} \cap \mathcal{B}$. Since consumers hold beliefs $p_{OR}^e(v, \theta) + \epsilon < p_{BMR}^e(v, \theta)$, those who visit the BMR will inspect (v, θ) in the store and purchase online if the product fits. However, given this shopping strategy, once the consumers with $s < \hat{s}$ searches online (and the search cost ϵ is sunk), it becomes optimal for the OR to deviate and increase price to $p_{BMR}(v, \theta)$. (Case 4) Suppose $p_{BMR}(v, \theta) = p_{OR}(v, \theta) + \epsilon \leq p^*(v)$ holds for a positive measure of products. If consumers break indifference by searching and purchasing online, the BMR would want to deviate and slightly undercut price so that $p_{BMR}(v, \theta) < p_{OR}(v, \theta) + \epsilon$, which brings us back to Case 2. If consumers break indifference by purchasing at the BMR, then the OR would want to either undercut prices and go back to Case 3, or charge the high travel cost consumers $p^*(v)$. (Case 5) Therefore, we are left with the only possibility of $p_{BMR}(v, \theta) \leq p^*(v) \leq p_{OR}(v, \theta) + \epsilon$ almost everywhere (i.e., products that do not satisfy this has measure zero). If $p_{BMR}^e(v, \theta) = p_{OR}(v, \theta)^e + \epsilon$ for a positive measure of products, and consumers break indifference by searching online, then we are back to Case 4, and the BMR would want to slightly undercut price. If $p_{BMR}^e(v, \theta) < p_{OR}^e(v, \theta) + \epsilon$ or $p_{BMR}^e(v, \theta) = p_{OR}^e(v, \theta) + \epsilon$ and consumers break indifference by purchasing at the BMR, then consumers with low (high) travel cost buy $(v, \theta) \in \mathcal{O} \cap \mathcal{B}$ at the BMR (OR). Given this optimal shopping strategy induced by consumer beliefs, both retailers optimally

charge $p^*(v)$. In this situation, the belief $p_{BMR}^e(v, \theta) < p_{OR}^e(v, \theta) + \epsilon$ is correct on path, and therefore is the only possible equilibrium price for $(v, \theta) \in \mathcal{O} \cap \mathcal{B}$.

Second, we show that, in equilibrium, consumers do not search online for products with $u(v, \theta) - (1 - \theta)r < 0$. Using proof by contradiction, suppose there is an equilibrium where the OR sets price $p(v, \theta) < p^*(v)$ such that $u(v, \theta, p(v, \theta)) - (1 - \theta)r > \epsilon$ for a positive measure of products in $\{(v, \theta) \notin \mathcal{B} \mid u(v, \theta) - (1 - \theta)r < 0\}$. Then, $u(v, \theta, p^e(v, \theta)) - (1 - \theta)r > \epsilon$, and high travel cost consumers will want to search for these products and buy online. Given that consumers search for these products, because $p(v, \theta) < p^*(v)$, the OR would want deviate to increase the price until $u(v, \theta, p(v, \theta)) - (1 - \theta)r = 0$. Hence, this cannot be an equilibrium. Due to the hold-up problem of a costly search, consumers do not search online for products with $u(v, \theta) - (1 - \theta)r < 0$.

Finally, we show that the optimal choice is for a retailer i to charge $p_i(v, \theta) = p^*(v)$ for its exclusive products. Similar to the proof of Proposition 3.1, the number of consumers who visit each retailer depends only on consumers' anticipated price, but it does not depend on the product's actual price. Thus, the optimal approach is for a retailer i to charge $p_i(v, \theta) = p^*(v)$ for its exclusive products. \square

3.8.2 Proofs for the Application: Buy Online and Return in Store Agreements

Proof of Lemma 3.3. Similar to Proposition 3.1, we prove this result in two steps.

(Step 1) We take as given the retailers' equilibrium prices $p(v, \theta) = p^*(v)$ and derive consumers' optimal shopping strategy.

(Step 1.1) First, as in the proof of Proposition 3.1, observe that visiting only the BMR is suboptimal.

(Step 1.2) Second, suppose that a consumer visits the BMR first and visits the OR second. In this case, he obtains expected utility $\int_{\mathcal{B}} u(v, \theta) dF(v, \theta)$ from the BMR and pays travel cost s . When the consumer visits the online store, he optimally searches for the products that give a non-negative expected utility (i.e., products that satisfy $u(v, \theta) - (1 - \theta)r \geq 0$). The consumer purchases each of these products and obtains $\int_{\mathcal{O} \cap \mathcal{B}^c} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta)$ from the online store. Thus, the total expected utility

from this shopping trip is

$$\int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - s + \int_{\mathcal{O} \cap \mathcal{B}^c} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta).$$

(Step 1.3) Third, suppose that a consumer visits the OR first and visits the BMR second. In this case, returning products is costless. So the consumer first optimally searches for all products sold online. The consumer purchases all of these products and obtains $\int_{\mathcal{O}} u(v, \theta) dF(v, \theta)$ from the online store. Then, the consumer pays travel cost s to visit the BMR to return unwanted goods and to purchase more. The expected utility of purchasing additional goods is $\int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta)$. The total expected utility from this shopping trip is

$$\int_{\mathcal{O}} u(v, \theta) dF(v, \theta) + \int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) - s.$$

(Step 1.4) Fourth, suppose that a consumer visits only the OR. In this case, returning products is costly. So he optimally searches online for the products that give a non-negative expected utility, (i.e., products that satisfy $u(v, \theta) - (1 - \theta)r \geq 0$). The consumer purchases all of these products and obtains expected utility

$$\int_{\mathcal{O}} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta).$$

(Step 1.5) Finally, we compare the strategies discussed above. Notice that the shopping strategy in (Step 1.3) yields a higher payoff than the shopping strategy in (Step 1.2). The reason is that, with a return agreement, visiting the BMR store last saves return hassles. Thus, we only have to compare the shopping strategies in (Step 1.3) and (Step 1.4). A direct comparison shows that the consumer visits the BMR first and visits the OR second if

$$\int_{\mathcal{O}} u(v, \theta) dF(v, \theta) + \int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) - s \geq \int_{\mathcal{O}} [u(v, \theta) - r(1 - \theta)]^+ dF(v, \theta).$$

Rearranging terms, we get

$$\begin{aligned}
s &\leq \int_{\mathcal{O}} u(v, \theta) dF(v, \theta) + \int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O}} [u(v, \theta) - r(1 - \theta)]^+ dF(v, \theta) \\
&= \int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) + \int_{\mathcal{O} \cap \{(v, \theta) | u(v, \theta) > r(1 - \theta)\}} r(1 - \theta) dF(v, \theta) \\
&\quad + \int_{\mathcal{O} \cap \{(v, \theta) | u(v, \theta) \leq r(1 - \theta)\}} u(v, \theta) dF(v, \theta) \\
&= \int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) = \hat{s}_{BORS}.
\end{aligned}$$

Otherwise, the consumer only visits the online store.

(Step 2) To complete this proof, we take as given consumers' optimal shopping strategies and show that it is optimal to price $p(v, \theta) = p^*(v)$. This process is analogous to the proof of Proposition 3.1 The idea is that consumers decide on their shopping strategy based on the anticipated price and not on the actual price. Hence, the retailers have incentives to charge $p(v, \theta) = p^*(v)$. Thus, consumers anticipate $p_i(v, \theta) = p^*(v)$ because they understand the hold-up problem of costly search. \square

Proof of Lemma 3.4. We first derive the OR's product assortment and then derive the BMR's product assortment.

(Step 1) We show that the OR optimally chooses $\mathcal{O} = \{(v, \theta) | \theta D(v, p^*(v))p^*(v) \geq (1 - \theta)\tilde{R}\}$.

(Step 1.1) First, consider a product in the region $\{(v, \theta) | u(v, \theta) \geq (1 - \theta)r\}$. All consumers want to buy the product, regardless of how they plan to return unwanted products. Because consumers with $s \leq \hat{s}_{BORS}$ return to the BMR, and other consumers return directly, the OR's expected loss from returning is $(1 - \theta)[G(\hat{s}_{BORS})\tilde{R} + (1 - G(\hat{s}_{BORS}))R]$. Because $r \geq \frac{\int_p^\infty D(v, p) dv}{D(v, p^*(v)) \times p^*(v)} R \geq \frac{\int_p^\infty D(v, p) dv}{D(v, p^*(v)) \times p^*(v)} \tilde{R}$, we have that the expected profit of selling the product $\theta D(v, p^*(v))p^*(v) - (1 - \theta)[G(\hat{s})\tilde{R} + (1 - G(\hat{s}))R] \geq 0$. Thus, the OR should stock all products in $\{(v, \theta) | u(v, \theta) \geq (1 - \theta)r\}$.

(Step 1.2) Second, consider a product in the region $\{(v, \theta) | u(v, \theta) < (1 - \theta)r\}$. In this region, only consumers who plan to return to the BMR will search and buy these products. Thus, the OR's expected loss from returning is $(1 - \theta)\tilde{R}$. Because all consumers visit the OR first, there is no need to lose profit to attract consumers. So the

OR stocks a product if it gives non-negative expected profit. That is,

$$\theta D(v, p^*(v))p^*(v) - (1 - \theta)\tilde{R} \geq 0.$$

(Step 1.3) Combining the previous two steps, we have that the OR stocks a product if $u(v, \theta) \geq (1 - \theta)r$, or if $u(v, \theta) < (1 - \theta)r$ and $\theta D(v, p^*(v))p^*(v) - (1 - \theta)\tilde{R} \geq 0$. Because $r \geq \frac{\int_p^\infty D(v, p) dv}{D(v, p^*(v)) \times p^*(v)} R \geq \frac{\int_p^\infty D(v, p) dv}{D(v, p^*(v)) \times p^*(v)} \tilde{R}$, we have that

$$\mathcal{O} = \{(v, \theta) \mid \theta D(v, p^*(v))p^*(v) \geq (1 - \theta)\tilde{R}\}.$$

(Step 2) Now, we derive the BMR's optimal product assortment.

(Step 2.1) First, we write out the BMR's profit maximization problem:

$$\begin{aligned} \max_{\mathcal{B} \subset [\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} & \int_{\mathcal{B} \cap \mathcal{O}^c} \theta D(v, p^*(v))p^*(v) dF(v, \theta) G(\hat{s}_{BORS}) \\ \text{s.t. } & \hat{s}_{BORS} = \int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) \\ & \int_{\mathcal{B}} dF(v, \theta) \leq K \end{aligned}$$

(Step 2.2) Second, we argue that the BMR does not stock products available online because consumers always visit the OR first. Stocking products that are available online only takes up store capacity, does not change profit, and does not change how many consumers visit the BMR.

(Step 2.3) Third, if $K \geq \int_{\mathcal{O}^c} dF(v, \theta)$, then the store size is not binding. So the optimal solution is simple: The BMR stocks all products not available online (i.e., $\mathcal{B} = \mathcal{O}^c$) because this strategy maximizes the profit from each product and also maximizes the number of consumers who visit the BMR (maximizes \hat{s}_{BORS}). Otherwise, the capacity constraint is binding. The argument is analogous to the proof of Proposition 3.2.

The following steps consider the case with $K < \int_{\mathcal{O}^c} dF(v, \theta)$.

(Step 2.4) Fourth, let $q(v, \theta) \in \{0, 1\}$ denote the BMR's stocking decision for product

(v, θ) . Then, the optimization can be rewritten as

$$\begin{aligned} & \max_{q(v, \theta) \in \{0, 1\}} \int_{\mathcal{O}^c} \theta D(v, p^*(v)) p^*(v) q(v, \theta) dF(v, \theta) G(\hat{s}_{BORS}) \\ & \text{s.t. } \hat{s}_{BORS} = \int_{\mathcal{O}^c} u(v, \theta) q(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) \\ & \int q(v, \theta) dF(v, \theta) = K \end{aligned}$$

(Step 2.5) Fifth, we show that a solution to the BMR's profit maximization problem exists. This existence proof is analogous to that in the proof of Proposition 3.2.

(Step 2.6) Sixth, we can write a "Lagrangian" for the optimization problem:

$$\begin{aligned} \mathcal{L} = & \int_{\mathcal{O}^c} \theta D(v, p^*(v)) p^*(v) G(\hat{s}_{BORS}) q(v, \theta) dF(v, \theta) + \lambda_1 [K - \int_{\mathcal{O}^c} q(v, \theta) dF(v, \theta)] \\ & + \lambda_2 [\int_{\mathcal{O}^c} u(v, \theta) q(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) - \hat{s}_{BORS}] \end{aligned}$$

Collecting all terms with $q(v, \theta)$ and separating them out from the other terms, we get:

$$\begin{aligned} \mathcal{L} = & \int_{\mathcal{O}^c} [\theta D(v, p^*(v)) p^*(v) G(\hat{s}_{BORS}) - \lambda_1 + \lambda_2 u(v, \theta)] q(v, \theta) dF(v, \theta) \\ & + \lambda_1 K + \lambda_2 [\int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) - \hat{s}_{BORS}] \end{aligned}$$

(Step 2.7) Seventh, as in Rhodes et al. (2020), because the Lagrangian is linear in $q(v, \theta)$, the optimum product assortment should have $q(v, \theta) = 1$ if the coefficient is positive and $q(v, \theta) = 0$ if the coefficient is negative. Thus, BMR stocks a product (i.e., $q(v, \theta) = 1$) iff:

$$\theta D(v, p^*(v)) p^*(v) G(\hat{s}_{BORS}) + \lambda_2 u(v, \theta) \geq \lambda_1, \text{ and } (v, \theta) \notin \mathcal{O}$$

Plugging in the optimal $q(v, \theta)$, we can solve $\{\hat{s}_{BORS}, \lambda_1, \lambda_2\}$ from the system of equations

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \hat{s}_{BORS}} = 0 \\ \hat{s}_{BORS} = \int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} u(v, \theta) q(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) \\ \int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} q(v, \theta) dF(v, \theta) = K \end{cases}$$

□

Proof of Proposition 3.5. By Lemma 3.4, $\mathcal{B} \cap \mathcal{O} = \emptyset$, so the two retailers fully differen-

tiate when there is a BORS agreement.

Comparing Lemma 3.4 with Proposition 3.2, the products $\{(v, \theta) \mid u(v, \theta) < (1 - \theta)r, \theta D(v, p^*(v))p^*(v) \geq (1 - \theta)\tilde{R}\}$ were previously unavailable online, but now with a BORS agreement in place, they are now offered. These products have lower θ than the products $\{(v, \theta) \mid u(v, \theta) \geq (1 - \theta)r\}$. The BMR is constrained to stock products in \mathcal{O}^c . Because \mathcal{O} expanded, the BMR offers more niche products. \square

For the next two proofs we introduce some new notations. Let \mathcal{O}_1 and \mathcal{B}_1 be the two retailers' product assortment in the absence of a BORS agreement, and let \mathcal{O}_2 and \mathcal{B}_2 be the two retailers' product assortment in the presence of a BORS agreement.

Proof of Proposition 3.6. From Lemma 3.3, we see that consumers always visit the OR first when there is a returns agreement. From Proposition 3.5, we have $\mathcal{O}_1 \subset \mathcal{O}_2$. Therefore, BORS leads to consumers purchasing $\mathcal{O}_1 \cap \mathcal{B}_1$ and $\mathcal{O}_2 \cap \mathcal{O}_1^c$ online. Since more purchases are made online, consumers also make more returns.

Now, we analyze how the BMR's store traffic changes. Without a BORS agreement, consumers visit the BMR if $s < \hat{s}$. Simple algebra shows that

$$\begin{aligned} \hat{s} &= \int_{\mathcal{B}_1} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O}_1 \cap \mathcal{B}_1} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta) \\ &\leq u(\bar{v}, \bar{\theta}) \cdot K \end{aligned}$$

With a BORS agreement, consumers visit the BMR if $s < \hat{s}_{BORS}$. Further, observe that

$$\begin{aligned} \hat{s}_{BORS} &= \int_{\mathcal{B}_2 \cap \mathcal{O}_2^c} u(v, \theta) dF(v, \theta) + \int_{\mathcal{O}_2} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) \\ &> \int_{\mathcal{O}_2} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) \end{aligned}$$

Therefore, if $K < \frac{1}{u(\bar{v}, \bar{\theta})} \int_{\mathcal{O}_2} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta)$, then $\hat{s}_{BORS} > \hat{s}$. Thus, more consumers visit the BMR when there is a BORS agreement and K is small enough. \square

Proof of Proposition 3.7. From Proposition 3.2 and Lemma 3.3, the OR always benefits from the BORS agreement.

Without a returns agreement, the BMR's profit is bounded above by

$$\begin{aligned}\pi_B &= \int_{\mathcal{B}_1} \theta D(v, p^*(v)) p^*(v) dF(v, \theta) \times G(\hat{s}_{BORS}) \\ &< \bar{\theta} D(\bar{v}, p^*(v)) p^*(v) KG[u(\bar{v}, \bar{\theta})K]\end{aligned}$$

With a returns agreement, the BMR's profit is bounded below by

$$\begin{aligned}\pi_B &= \int_{\mathcal{B}_2} \theta D(v, p^*(v)) p^*(v) dF(v, \theta) \times G(\hat{s}) \\ &> \underline{\theta} D(\underline{v}, p^*(v)) p^*(v) KG\left[\int_{\mathcal{O}_2} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta)\right]\end{aligned}$$

Thus, if

$$K \leq \frac{1}{u(\bar{v}, \bar{\theta})} G^{-1}\left\{\frac{\underline{\theta} D(\underline{v}, p^*(v))}{\bar{\theta} D(\bar{v}, p^*(v))} G\left[\int_{\mathcal{O}_2} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta)\right]\right\},$$

then the BMR also benefits from the BORS agreement. Thus, the two retailers can reach a BORS agreement without any transfers.

Notice that $\int_{\mathcal{O}_2} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta)$ is increasing in r and decreasing in \tilde{R} (because \mathcal{O}_2 is dependent on \tilde{R}). Thus, if r is large enough and \tilde{R} is small enough, such that

$$\int_{\mathcal{O}_2} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) \geq G^{-1}\left\{\frac{\bar{\theta} D(\bar{v}, p^*(v))}{\underline{\theta} D(\underline{v}, p^*(v))} G[u(\bar{v}, \bar{\theta})K]\right\},$$

then the BMR also benefits from the BORS agreement. If this inequality is violated, then a large r and small \tilde{R} only guarantee that the OR benefits a lot from the returns agreement. In this case, it can be profitable to compensate the BMR to reach an agreement. □

3.8.3 Proofs for the Extension

3.8.3.1 Proofs for the extension 1: Consumers without access to the OR

In this subsection, we first analyze the case without a BORS agreement in Lemma 3.5 and Proposition 3.11. Then, we consider the case where a BORS agreement is reached in Lemma 3.6 and Proposition 3.12. Finally, we state the proof of Proposition 3.8.

Suppose there is a fraction of $\gamma \in [0, 1]$ consumers that are loyal to the BMR.

Lemma 3.5. *In the absence of a BORS agreement, each retailer sets price $p^*(v)$ for product (v, θ) . There exists \hat{s} as defined in Lemma 3.1 and $\hat{s}_l = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta)$, such that:*

1. *Non-loyal consumers behave as in Proposition 3.1: If $s \leq \hat{s}$, they first visit the BMR to buy products $(v, \theta) \in \mathcal{B}$ and then search the OR to buy $\{(v, \theta) \in \mathcal{O} \cap \mathcal{B}^c \mid u(v, \theta) - (1 - \theta)r \geq 0\}$; If $s > \hat{s}$, they only search and buy products $\{(v, \theta) \in \mathcal{O} \mid u(v, \theta) - (1 - \theta)r \geq 0\}$ online.*
2. *Loyal consumers visit the BMR to buy $(v, \theta) \in \mathcal{B}$ if $s \leq \hat{s}_l = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta)$, and they take their outside option if $s > \hat{s}_l$.*

Proof. The proofs for pricing and for the non-loyal consumers' shopping strategy are analogous to Proposition 3.1. Non-loyal consumers can choose to visit the BMR or can take their outside option. The expected utility of visiting the BMR is

$$\int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - s.$$

Because taking the outside option yields 0 utility, loyal consumers visit the BMR if $s \leq \hat{s}_l = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta)$. □

Proposition 3.11. *In the absence of a BORS agreement,*

1. *OR's assortment is the same as in Proposition 3.1: $\mathcal{O} = \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}$.*
2. *BMR stocks a product iff*

$$\begin{cases} \theta D(v, p^*(v)) p^*(v) [(1 - \gamma)G(\hat{s}) + \gamma G(\hat{s}_l)] + \lambda_2(1 - \theta)r + \lambda_3 u(v, \theta) \geq \lambda_1, & (v, \theta) \in \mathcal{O} \\ \theta D(v, p^*(v)) p^*(v) [(1 - \gamma)G(\hat{s}) + \gamma G(\hat{s}_l)] + \lambda_2 u(v, \theta) + \lambda_3 u(v, \theta) \geq \lambda_1, & (v, \theta) \notin \mathcal{O} \end{cases},$$

where $\{\hat{s}, \hat{s}_l, \lambda_1, \lambda_2, \lambda_3\}$ solves

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \hat{s}} = \frac{\partial \mathcal{L}}{\partial \hat{s}_l} = 0 \\ \hat{s} = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [u(v, \theta) - (1 - \theta)r] dF(v, \theta) \\ \hat{s}_l = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) \\ \int_{\mathcal{B}} dF(v, \theta) = K \end{cases}.$$

Proof. (Step 1) First, the proof for $\mathcal{O} = \{(v, \theta) \mid u(v, \theta) - (1 - \theta)r \geq 0\}$ is identical to Proposition 3.1.

(Step 2) Second, we derive the BMR's optimal product assortment.

(Step 2.1) We write out the BMR's profit maximization problem:

$$\begin{aligned} \max_{\mathcal{B} \subset [\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} & \int_{\mathcal{B}} \theta D(v, p^*(v)) p^*(v) dF(v, \theta) \times [(1 - \gamma)G(\hat{s}) + \gamma G(\hat{s}_l)] \\ \text{s.t. } & \hat{s} = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [u(v, \theta) - (1 - \theta)r]^+ dF(v, \theta) \\ & \int_{\mathcal{B}} dF(v, \theta) \leq K \\ & \hat{s}_l = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) \end{aligned}$$

Let $q(v, \theta) \in \{0, 1\}$ denote the BMR's stocking decision for product (v, θ) . Then, the optimization can be rewritten as

$$\begin{aligned} \max_{q(v, \theta) \in \{0, 1\}} & \int \theta D(v, p^*(v)) p^*(v) q(v, \theta) dF(v, \theta) \times [(1 - \gamma)G(\hat{s}) + \gamma G(\hat{s}_l)] \\ \text{s.t. } & \hat{s} = \int u(v, \theta) q(v, \theta) dF(v, \theta) - \int_{\mathcal{O}} [u(v, \theta) - (1 - \theta)r] q(v, \theta) dF(v, \theta) \\ & \int q(v, \theta) dF(v, \theta) \leq K \\ & \hat{s}_l = \int u(v, \theta) q(v, \theta) dF(v, \theta) \end{aligned}$$

(Step 2.2) There exists a solution to BMR's profit maximization problem. Also, the constraint $\int q(v, \theta) dF(v, \theta) \leq K$ must hold with equality at the optimum. The arguments are analogous to those in Proposition 3.1.

(Step 2.3) We can write a "Lagrangian" for the optimization problem:

$$\begin{aligned} \mathcal{L} = & \int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} \theta D(v, p^*(v)) p^*(v) [(1 - \gamma)G(\hat{s}) + \gamma G(\hat{s}_l)] q(v, \theta) dF(v, \theta) \\ & + \lambda_1 [K - \int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} q(v, \theta) dF(v, \theta)] \\ & + \lambda_2 [\int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} u(v, \theta) q(v, \theta) dF(v, \theta) - \int_{\mathcal{O}} [u(v, \theta) - (1 - \theta)r] q(v, \theta) dF(v, \theta) - \hat{s}] \\ & + \lambda_3 [\int u(v, \theta) q(v, \theta) dF(v, \theta) - \hat{s}_l] \end{aligned}$$

Collecting all terms with $q(v, \theta)$ and separating them from the other terms, we get

$$\begin{aligned}
\mathcal{L} = & \int_{\mathcal{O}} \left\{ \begin{array}{l} \theta D(v, p^*(v)) p^*(v) [(1 - \gamma)G(\hat{s}) + \gamma G(\hat{s}_l)] \\ -\lambda_1 + \lambda_2(1 - \theta)r + \lambda_3 u(v, \theta) \end{array} \right\} q(v, \theta) dF(v, \theta) \\
& + \int_{\mathcal{O}^c} \left\{ \begin{array}{l} \theta D(v, p^*(v)) p^*(v) [(1 - \gamma)G(\hat{s}) + \gamma G(\hat{s}_l)] \\ -\lambda_1 + \lambda_2 u(v, \theta) + \lambda_3 u(v, \theta) \end{array} \right\} q(v, \theta) dF(v, \theta) \\
& + \lambda_1 K - \lambda_2 \hat{s} - \lambda_3 \hat{s}_l
\end{aligned}$$

(Step 2.4) As in Rhodes et al. (2020), because the Lagrangian is linear in $q(v, \theta)$, the optimum product assortment should have $q(v, \theta) = 1$ if the coefficient is positive and $q(v, \theta) = 0$ if it is negative. Thus, BMR stocks a product (i.e., $q(v, \theta) = 1$) iff

$$\begin{cases} \theta D(v, p^*(v)) p^*(v) [(1 - \gamma)G(\hat{s}) + \gamma G(\hat{s}_l)] + \lambda_2(1 - \theta)r + \lambda_3 u(v, \theta) \geq \lambda_1, & (v, \theta) \in \mathcal{O} \\ \theta D(v, p^*(v)) p^*(v) [(1 - \gamma)G(\hat{s}) + \gamma G(\hat{s}_l)] + \lambda_2 u(v, \theta) + \lambda_3 u(v, \theta) \geq \lambda_1, & (v, \theta) \notin \mathcal{O} \end{cases}$$

After plugging in the optimal $q(v, \theta)$, we can determine $\{\hat{s}_1, \hat{s}_l, \lambda_1, \lambda_2, \lambda_3\}$ from the following system of equations:

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \hat{s}} = \frac{\partial \mathcal{L}}{\partial \hat{s}_l} = 0 \\ \hat{s} = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [u(v, \theta) - (1 - \theta)r] dF(v, \theta) \\ \hat{s}_l = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) \\ \int_{\mathcal{B}} dF(v, \theta) = K \end{cases}$$

□

Next, we consider the case where a BORS agreement is reached.

Lemma 3.6. *When the BORS agreement is reached, each retailer sets price $p^*(v)$ for product (v, θ) .*

1. *Non-loyal consumers behave as in Lemma 3.3: If $s \leq \tilde{s}_{BORS}$, they first search and buy all products online and then visit the BMR to return unwanted goods and to purchase matched products in $\mathcal{B} \setminus \mathcal{O}$; if $s > \tilde{s}_{BORS}$, they only search and buy products $\{(v, \theta) \in \mathcal{O} \mid u(v, \theta) - (1 - \theta)r \geq 0\}$ online.*
2. *Loyal consumers behave as in Lemma 3.5: They visit the BMR if $s \leq \hat{s}_l$ and take their outside option if $s > \hat{s}_l$.*

Proof. The proof is analogous to Lemma 3.3. For each consumer segment, we compute the expected utility of each possible shopping strategy. Consumers observe their private travel cost and choose the strategy that yields the highest expected payoff. \square

Proposition 3.12. *With a BORS agreement,*

1. *OR's assortment is the same as in Lemma 3.4: $\mathcal{O} = \{(v, \theta) \mid \theta D(v, p^*(v))p^*(v) - (1 - \theta)\tilde{R} \geq 0\}$.*
2. *BMR stocks a product iff*

$$\begin{cases} \theta D(v, p^*(v))p^*(v)\gamma G(\hat{s}_l) + \lambda_3 u(v, \theta) \geq \lambda_1, & (v, \theta) \in \mathcal{O} \\ \theta D(v, p^*(v))p^*(v)[(1 - \gamma)G(\tilde{s}_{BORS}) + \gamma G(\hat{s}_l)] + \lambda_2 u(v, \theta) + \lambda_3 u(v, \theta) \geq \lambda_1, & (v, \theta) \notin \mathcal{O} \end{cases},$$

where $\{\tilde{s}_{BORS}, \hat{s}_l, \lambda_1, \lambda_2, \lambda_3\}$ solves

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \tilde{s}_{BORS}} = \frac{\partial \mathcal{L}}{\partial \hat{s}_l} = 0 \\ \tilde{s}_{BORS} = \int_{\mathcal{O}^c} u(v, \theta) q(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) \\ \hat{s}_l = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) \\ \int_{\mathcal{B}} dF(v, \theta) = K \end{cases}$$

Proof. (Step 1) First, the proof for $\mathcal{O} = \{(v, \theta) \mid \theta D(v, p^*(v))p^*(v) - (1 - \theta)\tilde{R} \geq 0\}$ is identical to Lemma 3.3.

(Step 2) Second, we derive the BMR's optimal product assortment.

(Step 2.1) We write out the BMR's profit maximization problem:

$$\begin{aligned} & \max_{\mathcal{B} \subset [v, \bar{v}] \times [\theta, \bar{\theta}]} \int_{\mathcal{B} \cap \mathcal{O}} \theta D(v, p^*(v))p^*(v) dF(v, \theta) \times (1 - \gamma)G(\tilde{s}_{BORS}) \\ & \quad + \int_{\mathcal{B}} \theta D(v, p^*(v))p^*(v) dF(v, \theta)\gamma G(\hat{s}_l) \\ & \text{s.t. } \tilde{s}_{BORS} = \int_{\mathcal{B} \cap \mathcal{O}^c} u(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta). \\ & \quad \int_{\mathcal{B}} dF(v, \theta) \leq K \\ & \quad \hat{s}_l = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) \end{aligned}$$

Let $q(v, \theta) \in \{0, 1\}$ denote the BMR's stocking decision for product (v, θ) . Then, the optimization can be rewritten as

$$\begin{aligned}
& \max_{q(v, \theta) \in \{0, 1\}} \int_{\mathcal{O}} \theta D(v, p^*(v)) p^*(v) q(v, \theta) dF(v, \theta) \times (1 - \gamma) G(\tilde{s}_{BORS}) \\
& \quad + \int \theta D(v, p^*(v)) p^*(v) q(v, \theta) dF(v, \theta) \gamma G(\hat{s}_l) \\
& \text{s.t. } \tilde{s}_{BORS} = \int_{\mathcal{O}^c} u(v, \theta) q(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta). \\
& \quad \int q(v, \theta) dF(v, \theta) \leq K \\
& \quad \hat{s}_l = \int u(v, \theta) q(v, \theta) dF(v, \theta)
\end{aligned}$$

(Step 2.2) There exists a solution to BMR's profit maximization problem. Also, the constraint $\int q(v, \theta) dF(v, \theta) \leq K$ must hold with equality at the optimum if $K < 1$. The arguments are analogous to those in Proposition 3.1.

(Step 2.3) We can write a Lagrangian for the optimization problem:

$$\begin{aligned}
\mathcal{L} &= \int_{\mathcal{O}} \theta D(v, p^*(v)) p^*(v) q(v, \theta) dF(v, \theta) \times (1 - \gamma) G(\tilde{s}_{BORS}) \\
& \quad + \int \theta D(v, p^*(v)) p^*(v) q(v, \theta) dF(v, \theta) \gamma G(\hat{s}_l) \\
& \quad + \lambda_1 [K - \int_{[\underline{v}, \bar{v}] \times [\underline{\theta}, \bar{\theta}]} q(v, \theta) dF(v, \theta)] \\
& \quad + \lambda_2 [\int_{\mathcal{O}^c} u(v, \theta) q(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) - \tilde{s}_{BORS}] \\
& \quad + \lambda_3 [\int u(v, \theta) q(v, \theta) dF(v, \theta) - \hat{s}_l]
\end{aligned}$$

Collecting all terms with $q(v, \theta)$ and separating them from other terms, we get

$$\begin{aligned}
\mathcal{L} &= \int_{\mathcal{O}} \{\theta D(v, p^*(v)) p^*(v) \gamma G(\hat{s}_l) + \lambda_3 u(v, \theta) - \lambda_1\} q(v, \theta) dF(v, \theta) \\
& \quad + \int_{\mathcal{O}^c} \left\{ \theta D(v, p^*(v)) p^*(v) [(1 - \gamma) G(\tilde{s}_{BORS}) + \gamma G(\hat{s}_l)] \right. \\
& \quad \quad \left. - \lambda_1 + \lambda_2 u(v, \theta) + \lambda_3 u(v, \theta) \right\} q(v, \theta) dF(v, \theta). \\
& \quad + \lambda_1 K - \lambda_2 \tilde{s}_{BORS} - \lambda_3 \hat{s}_l
\end{aligned}$$

(Step 2.4) As in Rhodes et al. (2020), because the Lagrangian is linear in $q(v, \theta)$, the optimum product assortment should have $q(v, \theta) = 1$ if the coefficient is positive and

$q(v, \theta) = 0$ if it is negative. Thus, BMR stocks a product (i.e., $q(v, \theta) = 1$) iff

$$\begin{cases} \theta D(v, p^*(v)) p^*(v) \gamma G(\hat{s}_l) + \lambda_3 u(v, \theta) \geq \lambda_1, & (v, \theta) \in \mathcal{O} \\ \theta D(v, p^*(v)) p^*(v) [(1 - \gamma) G(\tilde{s}_{BORS}) + \gamma G(\hat{s}_l)] + \lambda_2 u(v, \theta) + \lambda_3 u(v, \theta) \geq \lambda_1, & (v, \theta) \notin \mathcal{O} \end{cases}$$

Plugging in the optimal $q(v, \theta)$, $\{\tilde{s}_{BORS}, \hat{s}_l, \lambda_1, \lambda_2, \lambda_3\}$ can be solved from the system of equations

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \tilde{s}_{BORS}} = \frac{\partial \mathcal{L}}{\partial \hat{s}_l} = 0 \\ \tilde{s}_{BORS} = \int_{\mathcal{O}^c} u(v, \theta) q(v, \theta) dF(v, \theta) + \int_{\mathcal{O}} \min\{u(v, \theta), r(1 - \theta)\} dF(v, \theta) \\ \hat{s}_l = \int_{\mathcal{B}} u(v, \theta) dF(v, \theta) \\ \int_{\mathcal{B}} dF(v, \theta) = K \end{cases}$$

□

Now, we are ready to prove Proposition 3.8.

Proof of Proposition 3.8. Comparing Lemma 3.4 with Proposition 3.12, in the region $(v, \theta) \in \mathcal{O}$, the marginal benefit of stocking a product becomes positive when $\gamma > 0$. By continuity, the marginal benefit of stocking a product in \mathcal{O} exceeds that in \mathcal{O}^c with a sufficiently large γ . Therefore, a significant loyal segment of the BMR generates product line overlap in the presence of BORS.

The long-run equilibrium product assortments with(out) BORS is given in Proposition 3.11 (12). When the two retailers reach a BORS agreement, the BMR's marginal benefit of selling overlapping products (i.e., $(v, \theta) \in \mathcal{O}$) decreases by $\theta D(v, p^*(v))(1 - \gamma)G(\hat{s}) + \lambda_2(1 - \theta)r$. So, with BORS, the BMR offers less products in \mathcal{O} and more products in \mathcal{O}^c . Since \mathcal{O} expands after offering BORS, \mathcal{O}^c contains more niche products, and the BMR offers more products with lower fit probability, θ . This decreases the consumer surplus of BMR's loyal consumers. As the fraction of the BMR's loyal consumers shrinks, $1 - \gamma$ increases, and $\theta D(v, p^*(v))(1 - \gamma)G(\hat{s}) + \lambda_2(1 - \theta)r$ increases. This greater decline in the marginal benefit of selling overlapping products leads to more niche product assortment of the BMR and lower consumer surplus for its loyal segment.

□

3.8.4 Proofs for the extension 2: Price competition

In the equilibrium with price competition, consumer behavior is similar to Proposition 3.1. We have a slightly different cut-off consumer given by

$$\tilde{s} = \int_{\mathcal{B}} \mathbb{E}[u(v, \theta, p_{BMR})] dF(v, \theta) - \int_{\mathcal{O} \cap \mathcal{B}} [\mathbb{E}[u(v, \theta, p_{\min})] - (1 - \theta)r]^+ dF(v, \theta),$$

where expectations are taken over random prices. Consumers with $s > \tilde{s}$ optimally choose to only shop online, while consumers with $s \leq \tilde{s}$ shops in-store first and online second. We define $\bar{p}(v, \theta) = \min\{p^*(v), \hat{p}(v, \theta)\}$, where $\hat{p}(v, \theta)$ is the solution to $u(v, \theta, \hat{p}) - (1 - \theta)r = 0$ and $p^*(v) = \arg \max_p \theta D(v, p)$ is the monopoly price. Denote $\underline{p}(v, \theta)$ as the smallest solution to $\theta D(v, \underline{p}(v, \theta)) - (1 - \theta)r = 0$. The next lemma characterizes the support of the random prices.

Lemma 3.7. *There is a unique cut-off \tilde{v} such that $\theta D(v, \bar{p}(v, \theta)) - (1 - \theta)r = 0$. The OR never sells a product $(v, \theta) \in \mathcal{B}$ if $v \leq \tilde{v}$. Thus, on equilibrium path, we can restrict attention to $v > \tilde{v}$ for overlapping products $(v, \theta) \in \mathcal{B} \cap \mathcal{O}$.¹³ In equilibrium, overlapping products have prices $p_{OR}(v, \theta), p_{BMR}(v, \theta) \in [\underline{p}(v, \theta), \bar{p}(v, \theta)]$.*

Proof. Observe that the price distribution of the two retailers must have the same support. This is proven by Varian (1980), and the logic is that if price distributions have different supports, then the retailer with a larger support upper-bound would want to undercut prices and the other retailer would want to raise prices. The rest of this proof follows in two steps. We first show $p(v, \theta) \leq \bar{p}(v, \theta)$, and then show $p(v, \theta) \geq \underline{p}(v, \theta)$.

First, we prove $p(v, \theta) \leq \bar{p}(v, \theta)$. No retailer would set $p(v, \theta) > \bar{p}(v, \theta) = p^*(v)$, because lowering prices to the monopoly price $p^*(v)$ attracts more consumers and leads to higher profit per consumer. On the other hand, suppose the OR sets $p(v, \theta) > \bar{p}(v, \theta) = \hat{p}(v, \theta)$. Then $u(v, \theta, p(v, \theta)) - (1 - \theta)r < 0$, and the consumers with $s > \tilde{s}$ will not buy the product from the OR because return costs are too high. Thus, the OR loses its captive consumers ($s > \tilde{s}$) at this high price, and both retailers compete for the consumers with $s \leq \tilde{s}$. This is a Bertrand type competition that drives prices down to $\hat{p}(v, \theta)$.

¹³Off equilibrium path, if there is a product with $v \leq \tilde{v}$ that are sold by both retailers (i.e., $(v, \theta) \in \mathcal{B} \cap \mathcal{O}$ and $v \leq \tilde{v}$), then $p_{OR}(v, \theta) = p_{BMR}(v, \theta) = p^*(v)$ is the unique equilibrium price. Intuitively, the OR cannot compete on prices, because of a net returns loss from its captive consumers. So, the BMR does not face a downwards pressure on pricing.

Second, we prove $p(v, \theta) \geq \underline{p}(v, \theta)$. If the OR sets price $\bar{p}(v, \theta)$, it obtains profit $[1 - G(\tilde{s})][\theta D(v, \bar{p}(v, \theta))\bar{p}(v, \theta) - (1 - R)\theta]$ from its captive consumers ($s > \tilde{s}$). Since $p_{BMR} \leq \bar{p}(v, \theta)$, the OR loses all non-captive consumers. If the OR sets price $\underline{p}(v, \theta)$, it obtains profit $[1 - G(\tilde{s})][\theta D(v, \underline{p}(v, \theta))\underline{p}(v, \theta) - (1 - R)\theta]$ from its captive consumers, and obtains $G(\tilde{s})\theta D(v, \underline{p}(v, \theta))\underline{p}(v, \theta)$ from non-captive consumers. Here, notice that non-captive consumers visit the BMR and then search online, so product fit is observable and there will be no returns. In the mixed strategy equilibrium, all prices in the support must give the same profit. Therefore, $\theta D(v, \underline{p}(v, \theta))\underline{p}(v, \theta) = [1 - G(\tilde{s})]\theta D(v, \bar{p}(v, \theta))\bar{p}(v, \theta)$.

Finally, when the OR randomize over prices of a product (v, θ) , its expected profit is simply given by $[1 - G(\tilde{s})][\theta D(v, \bar{p}(v, \theta))\bar{p}(v, \theta) - (1 - R)\theta]$. Therefore, it is only profitable for the OR to sell a product $(v, \theta) \in \mathcal{B}$ if $\theta D(v, \bar{p}(v, \theta))\bar{p}(v, \theta) - (1 - R)\theta > 0$, or equivalently, $v > \tilde{v}$. □

Proof of Proposition 3.9. By Lemma 3.7, when the two retailers compete on prices, the equilibrium prices have common support $[\underline{p}(v, \theta), \bar{p}(v, \theta)]$. As in other mixed strategy equilibria, each party randomizes to make their opponent indifferent between all actions played in equilibrium. In the following, we first derive $H_{BMR}(p)$ and then derive $H_{OR}(p)$.

Denote the BMR's price distribution as $H_{BMR}(p)$. To make the OR indifferent between all prices within $[\underline{p}(v, \theta), \bar{p}(v, \theta)]$, we need

$$\begin{aligned} & [1 - G(\tilde{s})][\theta D(v, p_1)p_1 - (1 - \theta)R] + G(\tilde{s})\theta D(v, p_1)p_1[1 - H_{BMR}(p_1)] \\ & = [1 - G(\tilde{s})][\theta D(v, p_2)p_2 - (1 - \theta)R] + G(\tilde{s})\theta D(v, p_2)p_2[1 - H_{BMR}(p_2)] \end{aligned}$$

for all $p_1, p_2 \in [\underline{p}(v, \theta), \bar{p}(v, \theta)]$. Here, the first term on each side represents the profit obtained from the captive consumers with $s \geq \tilde{s}$, and the second term represents the profit obtained from non-captive consumers ($s < \tilde{s}$) in the event that $p_{OR} < p_{BMR}$.¹⁴ Rearranging, we have

$$\theta D(v, p_1)p_1[1 - H_{BMR}(p_1)] - \theta D(v, p_2)p_2[1 - H_{BMR}(p_2)] = \frac{1 - G(\tilde{s})}{G(\tilde{s})} [\theta D(v, p_2)p_2 - \theta D(v, p_1)p_1].$$

¹⁴We assume that the non-captive consumers (who are currently in the BMR and checking for prices online) break indifference by purchasing in-store.

Dividing both sides by $p_1 - p_2$ and taking $p_2 \rightarrow p_1$, we get

$$\theta[D(v, p_1)p_1]'[1 - H_{BMR}(p_1)] - \theta D(v, p_1)p_1 h_{BMR}(p_1) = \frac{1 - G(\tilde{s})}{G(\tilde{s})}\theta[D(v, p_1)p_1]',$$

where h_{BMR} represents the probability density function of p_{BMR} . Rearranging, we have

$$h_{BMR}(p_1) + \frac{[D(v, p_1)p_1]'}{D(v, p_1)p_1}H_{BMR}(p_1) = \frac{[D(v, p_1)p_1]'}{G(\tilde{s})D(v, p_1)p_1}.$$

This is an ordinary differential equation with solution given by

$$H_{BMR}(p) = \frac{1}{D(v, p)p} \left[\frac{D(v, p)p}{G(\tilde{s})} + const \right],$$

for all $p \in [\underline{p}(v, \theta), \bar{p}(v, \theta)]$. Plugging in $H_{BMR}(\underline{p}(v, \theta)) = 0$, we get $const = -\frac{D(v, \underline{p}(v, \theta))\underline{p}(v, \theta)}{G(\tilde{s})}$.

Hence,

$$H_{BMR}(p) = \begin{cases} 1, & p \geq \bar{p}(v, \theta) \\ \frac{D(v, p)p - D(v, \underline{p}(v, \theta))\underline{p}(v, \theta)}{G(\tilde{s})D(v, p)p}, & \underline{p}(v, \theta) \leq p < \bar{p}(v, \theta) \\ 0, & p \leq \underline{p}(v, \theta) \end{cases}$$

Denote the OR's price distribution as $H_{OR}(p)$. To make the BMR indifferent between all prices within $[\underline{p}(v, \theta), \bar{p}(v, \theta)]$, we need

$$G(\tilde{s})\theta D(v, p_1)p_1[1 - H_{OR}(p_1)] = G(\tilde{s})\theta D(v, p_2)p_2[1 - H_{OR}(p_2)].$$

for all $p_1, p_2 \in [\underline{p}(v, \theta), \bar{p}(v, \theta)]$. Rearranging, we have

$$D(v, p_2)p_2H_{OR}(p_2) - D(v, p_1)p_1H_{OR}(p_1) = D(v, p_2)p_2 - D(v, p_1)p_1.$$

Dividing both sides by $p_2 - p_1$ and taking $p_1 \rightarrow p_2$, we get

$$[D(v, p_2)p_2]'H_{OR}(p_2) + D(v, p_2)p_2h_{OR}(p_2) = [D(v, p_2)p_2]',$$

where h_{OR} represents the probability density function of p_{OR} . Rearranging, we have

$$h_{OR}(p_2) + \frac{[D(v, p_1)p_1]'}{D(v, p_1)p_1}H_{OR}(p_2) = \frac{[D(v, p_1)p_1]'}{D(v, p_1)p_1}.$$

This is an ordinary differential equation with solution given by

$$H_{OR}(p) = \frac{1}{D(v,p)p} [D(v,p)p + const],$$

for all $p \in [\underline{p}(v, \theta), \bar{p}(v, \theta)]$. Plugging in $H_{OR}(\underline{p}(v, \theta)) = 0$, we get $const = -D(v, \underline{p}(v, \theta))\underline{p}(v, \theta)$. Hence,

$$H_{OR}(p) = \begin{cases} 1, & p \geq \bar{p}(v, \theta) \\ \frac{D(v,p)p - D(v, \underline{p}(v, \theta))\underline{p}(v, \theta)}{D(v,p)p}, & \underline{p}(v, \theta) \leq p < \bar{p}(v, \theta) \\ 0, & p \leq \underline{p}(v, \theta) \end{cases}$$

□

3.8.5 An example with closed-form solutions

3.8.5.1 Without a BORS agreement

Consider the case where products available are $\theta \sim U[0, 1]$, and $v > 0$ is a fixed number. Consumers' demand function (or purchase probability) is given by $D(v, p) = 8v(1 - p)$, and the travel cost to BMR $s \stackrel{iid}{\sim} U[0, 1]$. The marginal hassle cost of returning $r \geq v$. To avoid corner solutions, suppose that the BMR's capacity is moderate; that is, $K \in [\frac{5r - \sqrt{9r^2 + 8rv}}{4(r+v)}, \frac{\min\{r, v\}}{r+v}]$.

By Proposition 3.1, products are priced at $p^*(v) = \frac{1}{2}$. So we can compute consumers' expected consumption surplus $u(v, \theta) = v\theta$ and the retailers' expected profit $\theta D(v, p^*(v))p^*(v) = 2v\theta$. By Proposition 3.2, the OR stocks products with a high fit probability, $\theta \in [\frac{r}{r+v}, 1]$. From among the products that are sold online, the BMR optimally stocks those with low fit probability, $\theta \in [\frac{r}{r+v}, \hat{\theta}]$. For products that are not available online, the BMR stocks the ones with high fit probability, $\theta \in [\hat{\theta} - K, \frac{r}{r+v}]$. So the products $\theta \in [\hat{\theta} - K, \hat{\theta}]$ are sold by the BMR. Here, $\hat{\theta} = \frac{5K + \sqrt{\frac{12r + K^2r + K^2v}{r+v}}}{6} \in (\frac{r}{r+v}, 1)$. We can see that the BMR's store size capacity binds in equilibrium, and it sells products with moderate fit probability.

3.8.5.2 With a BORS agreement

Following the same setup as above, consider the case where a BORS agreement is reached. By Lemma 3.4, the OR stocks products $\theta \in [\frac{\hat{R}}{\hat{R}+v}, 1]$. The BMR avoids stocking products that are available online. When the BMR is small ($K < \frac{\hat{R}}{\hat{R}+v}$), it stocks

exclusive products with high fit probability, $\theta \in [\frac{\hat{R}}{\hat{R}+v} - K, \frac{\hat{R}}{\hat{R}+v}]$. When the BMR is large ($K \geq \frac{\hat{R}}{\hat{R}+v}$), it stocks all products that are not available online, $\theta \in [0, \frac{\hat{R}}{\hat{R}+v}]$.

Comparing with Example 1, we have that $\frac{\hat{R}}{\hat{R}+v} - K < \hat{\theta} - K$, so a larger variety of products is available. The OR expands to sell products that have a lower match probability, $\theta \in [\frac{\hat{R}}{\hat{R}+v}, \frac{r}{r+v})$. The BMR is forced in this case to sell more niche products. In aggregate, the equilibrium product variety increases.

3.8.6 Numerical analysis

For the purpose of plotting Figures 3.2-3.5, we make the following assumptions: The set of all products (v, θ) is jointly uniform on $[0, 1] \times [0, 1]$, consumers' travel cost $s \stackrel{iid}{\sim} U[0, 1]$, the demand function for a product that fits is $D(v, p) = 8v(1 - p)$, and consumers' marginal cost of returning $r = \frac{1}{2}$. These assumptions imply that $p^* = \frac{1}{2}$, $\theta p^* D(v, p^*) = 2\theta v$, and $u(v, \theta) = \theta \int_{p^*}^{\infty} D(v, p) dp = \theta v$.

Here, we describe how to compute the equilibrium of our main model. The computation for other cases are analogous. For simplicity of notations, let us define two sets:

$$S_1 = \{v, \theta \in [0, 1] \mid 2\theta v \hat{s} + \lambda_2(1 - \theta)r \geq \lambda_1, \theta v - (1 - \theta)r \geq 0\}$$

$$S_2 = \{v, \theta \in [0, 1] \mid 2\theta v \hat{s} + \lambda_2\theta v \geq \lambda_1, \theta v - (1 - \theta)r < 0\}$$

Observe that the two sets have multiple non-linear boundaries.

By Proposition 3.2, without a BORS agreement, the OR stocks a product if

$$\theta v - (1 - \theta)r \geq 0.$$

The BMR's stocking decision is given by

$$\begin{cases} 2\theta v \hat{s} + \lambda_2(1 - \theta)r \geq \lambda_1, & \theta v - (1 - \theta)r \geq 0 \\ 2\theta v \hat{s} + \lambda_2\theta v \geq \lambda_1, & \theta v - (1 - \theta)r < 0 \end{cases}$$

where $\{\hat{s}, \lambda_1, \lambda_2\}$ is given by

$$\begin{cases} 2 \int_{S_1} \theta v d\theta dv + 2 \int_{S_2} \theta v d\theta dv & = \lambda_2 \\ (1 - \theta)r \int_{S_1} d\theta dv + \int_{S_2} \theta v d\theta dv & = \hat{s} \\ \int_{S_1} d\theta dv + \int_{S_2} d\theta dv & = K \end{cases}$$

To obtain the equilibrium values of $\{\hat{s}, \lambda_1, \lambda_2\}$, we need to solve a system of three equations (with multiple non-linear integrals) and three unknowns. With our assumptions on the functional forms, there are four different integrals in the system of equations, $\int_{S_i} \theta v d\theta dv$ and $\int_{S_i} d\theta dv$, $i = 1, 2$. Each integral has a complicated functional form because of its nonlinear domain with three parameters $\{\hat{s}, \lambda_1, \lambda_2\}$. Below, we provide two methods to calculate the integrals and solve the system of equations.

[Method 1] This method uses Mathematica to analytically evaluate the integrals and only works with very simple functional forms, such as the ones we assumed. We can use the Boole command in Mathematica to define two indicator functions that take values 1 if and only if $(v, \theta) \in S_i$, $i = 1, 2$. Notice that $\int_{S_i} \theta v d\theta dv = \int_{[0,1]^2} \theta v \times \mathbb{1}_{(v,\theta) \in S_i} d\theta dv$ and $\int_{S_i} d\theta dv = \int_{[0,1]^2} \mathbb{1}_{(v,\theta) \in S_i} d\theta dv$. Hence, we can use the Integrate command in Mathematica to analytically evaluate the four integrals and use the FindRoot command to obtain the solutions $\{\hat{s}, \lambda_1, \lambda_2\}$.

[Method 2] This method approximates the integral equations and works with any functional form. First, we discretize the three-dimensional space of $\{\hat{s}, \lambda_1, \lambda_2\}$. For example, we can take 100 cuts in each dimension; this generates a lattice of 10^6 points. On each point, the sets S_1 and S_2 are functions of only v and θ . So the four integrals are real numbers on the discretized space and are easily computable. Thus, we can calculate the precise values of each equation on the lattice of 10^6 points. Then, we use interpolation methods to connect the discrete values of each equation using a smooth polynomial. The interpolation gives an approximation of the integral equations. We can use standard root-solving techniques (e.g., Newton's method) to obtain the solutions to the approximated equations.

Using either method, we find that $\{\hat{s}, \lambda_1, \lambda_2\} \approx \{0.0483348, 0.0808304, 0.24453\}$ when $K = \frac{1}{4}$, and $\{\hat{s}, \lambda_1, \lambda_2\} \approx \{0.0832457, 0.0873155, 0.411026\}$ when $K = \frac{1}{2}$.

3.9 Bibliography

- ALI, F. AND J. YOUNG (2021): "US ecommerce grows 32.4% in 2020," *Digital Commerce* 360.
- ARMSTRONG, M. (2017): "Ordered consumer search," *Journal of the European Economic Association*, 15, 989–1024.

- ARMSTRONG, M. AND J. VICKERS (2021): “Patterns of competitive interaction,” *Econometrica: Journal of the Econometric Society*.
- ARROW, K. J., T. HARRIS, AND J. MARSCHAK (1951): “Optimal inventory policy,” *Econometrica: Journal of the Econometric Society*, 250–272.
- BELL, D. R., J. CHOI, AND L. LODISH (2012): “What Matters Most in Internet Retailing,” *Sloan Management Review*.
- BRONNENBERG, B. J. (2015): “The provision of convenience and variety by the market,” *The RAND Journal of Economics*, 46, 480–498.
- (2020): “Innovation and Distribution: An Equilibrium Model of Manufacturing and Retailing,” *working paper*.
- CAVALLO, A. (2017): “Are online and offline prices similar? Evidence from large multi-channel retailers,” *American Economic Review*, 107, 283–303.
- CHINTAGUNTA, P. K. (1998): “Inertia and variety seeking in a model of brand-purchase timing,” *Marketing Science*, 17, 253–270.
- DALY, J. (2021): “Small specialty retail stores in the US,” *IBIS World*, INDUSTRY REPORT 45399.
- DENNIS, S. (2018): “Physical Retail Isn’t Dead. Boring Retail Is,” *Forbes*.
- DIAMOND, P. (1987): “Consumer differences and prices in a search model,” *The Quarterly Journal of Economics*, 102, 429–436.
- DIXIT, A. K. AND J. E. STIGLITZ (1977): “Monopolistic competition and optimum product diversity,” *The American economic review*, 67, 297–308.
- DUKES, A. AND L. LIU (2016): “Online shopping intermediaries: The strategic design of search environments,” *Management Science*, 62, 1064–1077.
- DUKES, A. J., T. GEYLANI, AND K. SRINIVASAN (2009): “Strategic assortment reduction by a dominant retailer,” *Marketing Science*, 28, 309–319.
- ERTEKIN, N., M. E. KETZENBERG, AND G. R. HEIM (2020): “Assessing impacts of store and salesperson dimensions of retail service quality on consumer returns,” *Production and Operations Management*, 29, 1232–1255.

- GALLINO, S. AND A. MORENO (2014): “Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information,” *Management Science*, 60, 1434–1451.
- GALLINO, S., A. MORENO, AND I. STAMATOPOULOS (2017): “Channel integration, sales dispersion, and inventory management,” *Management Science*, 63, 2813–2831.
- GAO, F. AND X. SU (2017): “Omnichannel retail operations with buy-online-and-pick-up-in-store,” *Management Science*, 63, 2478–2492.
- GODES, D. AND D. MAYZLIN (2004): “Using online conversations to study word-of-mouth communication,” *Marketing science*, 23, 545–560.
- GOLDFARB, A. AND C. TUCKER (2011): “Online display advertising: Targeting and obtrusiveness,” *Marketing Science*, 30, 389–404.
- HOLMES, T. J. (2001): “Bar codes lead to frequent deliveries and superstores,” *RAND Journal of Economics*, 708–725.
- HUANG, X. AND D. ZHANG (2020): “Service product design and consumer refund policies,” *Marketing Science*, 39, 366–381.
- JANAKIRAMAN, N. AND L. ORDÓÑEZ (2012): “Effect of effort and deadlines on consumer product returns,” *Journal of Consumer Psychology*, 22, 260–271.
- JING, B. (2018): “Showrooming and webrooming: Information externalities between online and offline sellers,” *Marketing Science*, 37, 469–483.
- JOHANSSON, A. (2020): “How to Remain Competitive in a Saturated Online Retail Market,” *Entrepreneur*.
- KÖK, A. G., M. L. FISHER, AND R. VAIDYANATHAN (2015): *Assortment planning: Review of literature and industry practice*, Springer.
- KÖK, A. G. AND Y. XU (2011): “Optimal and competitive assortments with endogenous pricing under hierarchical consumer choice models,” *Management Science*, 57, 1546–1563.
- KUKSOV, D. AND C. LIAO (2018): “When showrooming increases retailer profit,” *Journal of Marketing Research*, 55, 459–473.

- MAYZLIN, D. AND J. SHIN (2011): “Uninformative advertising as an invitation to search,” *Marketing Science*, 30, 666–685.
- MAYZLIN, D. AND H. YOGANARASIMHAN (2012): “Link to success: How blogs build an audience by promoting rivals,” *Management Science*, 58, 1651–1668.
- NAGESWARAN, L., S.-H. CHO, AND A. SCHELLER-WOLF (2020): “Consumer return policies in omnichannel operations,” *Management Science*, 66, 5558–5575.
- NAGESWARAN, L., E. H. HWANG, AND S.-H. CHO (2021): “Offline Returns for Online Retailers via Partnership,” *Available at SSRN*.
- OFEK, E., Z. KATONA, AND M. SARVARY (2011): ““Bricks and clicks”: The impact of product returns on the strategies of multichannel retailers,” *Marketing Science*, 30, 42–60.
- OKOLI, A. (2017): “7 Technologies to Watch: Bringing Offline Retail Experiences Online,” *Zoovu.com*.
- RABINOVICH, E., R. SINHA, AND T. LASETER (2011): “Unlimited shelf space in Internet supply chains: Treasure trove or wasteland?” *Journal of Operations Management*, 29, 305–317.
- RAO, S., E. RABINOVICH, AND D. RAJU (2014): “The role of physical distribution services as determinants of product returns in Internet retailing,” *Journal of Operations Management*, 32, 295–312.
- REAGAN, C. (2019): “That sweater you don’t like is a trillion-dollar problem for retailers. These companies want to fix it,” *CNBC*.
- REINARTZ, W., N. WIEGAND, AND M. IMSCHLOSS (2019): “The impact of digital transformation on the retailing value chain,” *International Journal of Research in Marketing*, 36, 350–366.
- REPKO, M. (2021): “A more than \$761 billion dilemma: Retailers’ returns jump as online sales grow,” *CNBC*.
- RHODES, A. (2015): “Multiproduct retailing,” *The Review of Economic Studies*, 82, 360–390.

- RHODES, A., M. WATANABE, AND J. ZHOU (2020): “Multiproduct Intermediaries,” *Journal of Political Economy*, forthcoming.
- SEETHARAMAN, P., A. AINSLIE, AND P. K. CHINTAGUNTA (1999): “Investigating household state dependence effects across categories,” *Journal of Marketing Research*, 36, 488–500.
- SELYUKH, A. (2018): “Online Shoppers Say They Rarely Return Purchases. Why?” *NPR.org*.
- SHIN, J. AND K. SUDHIR (2010): “A customer management dilemma: When is it profitable to reward one’s own customers?” *Marketing Science*, 29, 671–689.
- SHULMAN, J. D., A. T. COUGHLAN, AND R. C. SAVASKAN (2010): “Optimal reverse channel structure for consumer product returns,” *Marketing Science*, 29, 1071–1085.
- (2011): “Managing consumer returns in a competitive environment,” *Management Science*, 57, 347–362.
- SHULMAN, J. D., M. CUNHA JR, AND J. K. SAINT CLAIR (2015): “Consumer uncertainty and purchase decision reversals: Theory and evidence,” *Marketing Science*, 34, 590–605.
- SKROVAN, S. (2017): “What Matters Most in Internet Retailing,” *Retail Dive*.
- SPENCE, M. (1976): “Product differentiation and welfare,” *The American Economic Review*, 66, 407–414.
- THE ECONOMIST, . (2017): “Sorry, we’re closed - The decline of established American retailing threatens jobs,” *The Economist*.
- THOMAS, L. (2021): “Kohl’s says it added 2 million new customers in 2020, thanks to Amazon,” *CNBC*.
- VARIAN, H. R. (1980): “A model of sales,” *The American economic review*, 70, 651–659.
- YOUNG, J. (2020): “US ecommerce sales grow 14.9% in 2019,” *Digital Commerce 360*.