Frontiers of Information and Platform Design in Operations Management

by

Qian Zhang (Ken)

Master of Economics, Pennsylvania State University

Submitted to the Graduate Faculty of

the Katz Graduate School of Business in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2024

UNIVERSITY OF PITTSBURGH

KATZ GRADUATE SCHOOL OF BUSINESS

This dissertation was presented

by

Qian Zhang (Ken)

It was defended on

July 23, 2024

and approved by

Dissertation Chair: Arian Aflaki, Assistant Professor, Business Analytics and Operations

Dissertation Co-chair: Jennifer Shang, Professor, Business Analytics and Operations

Esther Gal-Or, Glenn Stinson Chair Professor, Marketing and Managerial Economics

Ran Shorrer, Assistant Professor, Department of Economics, Penn State University

Dennis Galletta, Professor and Director of Doctoral Programs, Information Systems

Frontiers of Information and Platform Design in Operations Management

Qian Zhang (Ken), PhD

University of Pittsburgh, 2024

This dissertation revolves around the intricate domains of information and mechanism design in operations management, presented in two essays. The first essay delves into the concerns surrounding customer apprehensions about price discrimination and explores the optimal utilization of information to alleviate such concerns. Employing a correlated Bayesian persuasion framework, we uncover the conditions under which a binary inventory signal complements pricing strategies, concurrently enhancing firm revenue and customer welfare. In the second essay, we study the design of rating platforms in the presence of disconfirmation effects, i.e., when customers incorporate their prior belief of the product into their ratings. The study elucidates the pivotal role of reference effects in shaping rating convergence to true product quality.

In the first essay, we focus on the strategic information transmission between firms and strategic customers. Customers have heterogeneous valuations for the products and services, and firms use various price discrimination tactics where they charge different prices to different customers. This practice, i.e., customizing the price for individual customers, is known as personalized pricing (PP) and is implemented in various industries. We investigate whether pricing can informatively signal PP to customers and how firms should adjust pricing strategies in response to customer reactions. We also investigate whether disclosing inventory information can benefit firms and customers, ultimately advocating for increased transparency in PP practices. By modeling dynamic interactions between firms and a continuum of heterogeneous strategic customers over two periods, we unveil nuanced insights. We find that firms reduce the first-period price to persuade high-valuation customers to purchase in the first period, even when they do not intend to implement PP. This is because the mere presence of PP risk makes customers reluctant to reveal their identity. The firm then must "compensate" customers to persuade them to reveal their valuations. We show that the price alone cannot perfectly signal the firm's PP intention when the firm takes the

strategic customer behavior into account. We next consider whether customers can infer the implementation of PP from the inventory availability information. To study this, we focus on a class of binary signals where the firm marks the inventory low when it is below a threshold. Such inventory signals are commonly adopted by retailers such as IKEA and ZARA. We show that an inventory signal can improve the firm revenue only when customers believe the firm conducts PP with a sufficiently low probability. In this case, an inventory signal alleviates customer PP concerns and allows the firm to set higher prices. Additionally, we demonstrate that disclosing inventory availability information is a strategic complement to the prices when alleviating the customer PP concerns. Furthermore, an inventory signal, in addition to the firm, can benefit all customers. With the growing interest in PP regulation, requiring firms to disclose inventory availability information could be a viable policy to make PP more transparent and credibly reduce customer concerns.

In the second essay, we study the customers' social learning problem upon observing the product ratings. Customers and platforms increasingly rely on online ratings to assess the quality of products and services. However, customer ratings are susceptible to various biases. *Disconfirmation bias* is a specific form where customers incorporate the discrepancy between their prior expectations and post-purchase experiences into their ratings. We study the asymptotic behavior of ratings in the presence of disconfirmation bias in three rating systems: (i) complete system, where customers observe the entire rating history; (ii) aggregate system, where only the frequency of each rating option is available; and (iii) average ratings, where customers solely use the average of past ratings. Customers are Bayesian and update their quality beliefs upon observing the ratings. After experiencing the product, they rate it according to their heterogeneous ex-post utility and disconfirmation bias. In complete and aggregate systems, we show that customer beliefs converge to the intrinsic quality when disconfirmation bias is small. When this bias is large, there will be a discrepancy between converged beliefs and the intrinsic quality, although this discrepancy could be arbitrarily small. When the disconfirmation bias is intermediate, beliefs may diverge significantly from the intrinsic quality or not converge. However, we establish that the platform can guarantee correct learning by designing a sufficiently granular rating system, i.e., a system with more rating options. We confirm all these results in the system with average ratings, albeit with a bias-correcting rule. Finally, we characterize the learning speed in the aggregate system.

In summary, this thesis contributes to the literature on the interface of information, platform, and mechanism design in operations management, unraveling the intricacies of pricing strategies, information revelation mechanisms, and rating platforms.

Table of Contents

List of Tables

List of Figures

Preface

First of all, I would like to thank my advisor, Professor Arian Aflaki. His unwavering support, guidance, and encouragement during past years were priceless, and he is a great role model for my academic career in the future. I appreciate his support to conduct research in my interest. Besides the academic and research aspect, I would like to thank him personally. During this long Ph.D. journey as an international student, his help and suggestions on my life have meant a lot to me.

I would also like to thank my other thesis committee members. Throughout the years, my program mentor, Professor Jennifer Shang, has been a steadfast support for me, assisting me with my teaching, research, and personal life. She has affected and led me in ways that are beyond words. I will never forget her gentle guidance and edification. I want to express my deep gratitude to her for all of the help she has given me. I am very much grateful to Professor Ran Shorrer for kindly supervising me when I was at Penn State. His passion and proficiency in research have always inspired me. I am deeply thankful to Professor Esther Gal-Or for guiding me patiently, helping me smoothly transit my career from Economics to Operations Management, and enlightening me many research ideas. Her kind support and advice in my job search inspire me in many ways. I would also like to thank Professor Dennis Galletta for his insightful questions and suggestions. More importantly, as the director of Katz doctoral program, Dr. Galletta has shown consistent support and guidance to facilitate my career at Katz.

Similarly, I am also very thankful to all faculty members of business analytics and operations (BAO) at Katz school who have provided guidance and help during the past years as a Ph.D. student in BAO. Many other teachers and professors from the Department of Economics, Department of Statistics, and Swanson School of Engineering at the University of Pittsburgh have inspired me to grow in the academic field. I am very thankful for their instructions. I would like to thank Ms. Carrie Woods, Dr. Dennis Galletta, Ms. Rachael McAleer, and all members of the doctoral office for their kindliness, consideration, and help during this memorable journey. I am also very grateful to all Ph.D. candidates at BAO

(Krista, Mike, Joya, Jing, Tiffany, Titing, Gabriel, Kasra, Qiqi, and Denise), and other departments for being great colleagues and friends in the past years. The relevant discussions and personal support have created memories that will last a lifetime.

1.0 Introduction

Platforms fostering strategic interactions among multiple stakeholders, including firms, customers, and content contributors, have recently become ubiquitous. I have concentrated my scholarly pursuits on investigating the underlying mechanisms governing these platforms and deciphering the complex interactions and incentives among agents to facilitate system-wide efficiency and fairness. My interdisciplinary approach draws from the realms of operations management, information systems, game theory, information economics, revenue management, pricing strategies, and machine learning.

In the contemporary business environment, effective information design enables businesses to fine-tune their communications and offerings to their customers' needs and preferences. This, in turn, enhances customer satisfaction and loyalty. Similarly, platform design, particularly in the context of social learning and rating systems, influences how customers perceive and evaluate products and services. By mitigating customer rating biases, businesses can ensure that customer feedback is more accurate and reliable, leading to better decision-making and improved product offerings.

Due to the critical importance of these areas, information design and platform design require further research. Their interdisciplinary nature, spanning design, technology, business, and social sciences, complicates the development of a unified research framework. The rapid pace of technological advancement also means that research must often catch up to practical applications. Nevertheless, integrating thoughtful information and platform design in business practices leads to improved efficiency, agility, and competitive advantage, ultimately contributing to sustained success and growth in a dynamic marketplace.

Chapter 2 of the dissertation focuses on firm's optimal information provision problem in Personalized Pricing (PP), where the firm customizes prices for individual customers. Despite the evident benefits of PP for firms and potential benefits for customers, customers remain concerned about PP. Our research explores whether pricing can informatively signal PP to customers and how firms should adjust pricing strategies in response to customer reactions. We also investigate whether disclosing inventory information can benefit both firms and customers, ultimately advocating for increased transparency in PP practices. To address these questions, we build a dynamic Bayesian persuasion game of repeated interactions between a firm and a continuum of heterogeneous customers. We assume the firm and customers are a priori uncertain about the feasibility of PP. We refer to a firm that can personalize the prices as P-type, and a firm that must charge uniform prices to all customers as U-type. A P-type firm has the additional incentive to learn customers' valuations, and then can use this information to personalize the prices . Customers cannot directly observe the firm's pricing strategy or inventory. Hence, they may be unable to identify whether a high price reflects inventory scarcity or the firm's intention to learn customer valuations for PP. Consequently, strategic customers have the incentive to forgo the purchase and hide their private information. As a result, such customer concerns may hurt both firm's revenue and customer's utility. Therefore, we consider firm's optimal information provision strategy to alleviate customer's concern about PP.

Chapter 3 of the dissertation studies the rating systems in the presence of disconfirmation bias, where customers reflect the discrepancy between their prior expectations and post-purchase experience in their ratings. This bias has been the subject of empirical studies, where researchers examine its magnitude and direction. In our social learning context, we explore the dynamics of customer learning when the customer disconfirmation bias influences ratings. We show whether customer beliefs about intrinsic quality converges depends on customer heterogeneity, disconfirmation bias, and the rating system's granularity. We investigate this convergence in three rating systems: a system where customers observe individual past customer ratings (complete system), only observe the frequency of each rating option (aggregate system), and only observe the average of past ratings. In each of these rating systems, we study the asymptotic behavior of the ratings and the effect of the disconfirmation bias on this behavior. We also examine the implications of the granularity of the rating system, i.e., the number of rating options available to customers. Finally, we investigate the speed of convergence in the aggregate system.

This thesis considers the information transmission problems from both firm's and customer's perspectives. It bridges the gap between theoretical models and practical applications, offering guidelines for firms on how to enhance customer trust and market efficiency through informed decision-making. The findings underscore the importance of transparency, strategic information disclosure, and platform design.

2.0 Is Your Price Personalized? Alleviating Customer Concerns with Inventory Availability Information

Personalized pricing (PP), i.e., customizing prices for individual customers, can benefit firms and some customers. However, customer concerns about being targeted by such practices have raised debates on PP tactics. In a Bayesian persuasion framework, we study whether and when price can signal such PP implementation to customers. We also investigate whether disclosing inventory availability information can alleviate customer concerns and benefit the stakeholders, including the firm and customers. We consider a dynamic personalized pricing and information provisioning game between a firm and a market of heterogeneous customers. The firm may set the price to learn the customer valuations in the first period and exploit this information in the second period. Customers are uncertain about inventory availability and PP implementation. However, they update their beliefs upon receiving new information. In addition to myopic customers who make decisions based on their immediate utility, we study strategic customers who consider their future utility. We show that price is insufficient to signal PP, hurting the firm and customers. We find conditions under which a binary inventory signal complements price and simultaneously increases the firm's revenue and benefits customers. We establish the robustness of our insights in a series of extensions.

2.1 Introduction

Customers have heterogeneous valuations for products and services. Therefore, a single price charged to all customers often fails to capture the full revenue potential. While a high price increases the revenue for each unit sold, it excludes the customers with lower valuations and reduces demand. To overcome this challenge, firms use various price discrimination tactics. For example, they may assign different prices to distinct customer segments such as students and senior citizens US Council of Economic Advisers (2016) or intertemporally alter prices to skim the market Besanko and Winston (1990) and Elmaghraby and Keskinocak (2003). Firms should ideally charge customers their maximum willingness to pay to maximize revenue. This practice, i.e., customizing the price for individual customers, is known as *personalized pricing (PP)* and is implemented in various industries.

Despite the evident benefits of PP for firms, it imposes unique implementation challenges. For example, PP is successful only when the firm can efficiently learn the customer valuations. In addition to IT capabilities, this requires experimentation with prices which some firms do not desire Wallheimer (2018), Ban and Keskin (2021). Also, while the extant literature has shown that customers may benefit from PP Dubé and Misra (2019), they remain concerned about these practices. Specifically, PP can hurt the customer trust Garbarino and Lee (2003) and be perceived as unfair Haws and Bearden (2006). In 2000, when Amazon experimented with using customer purchase history to personalize the prices for DVD movies, it was not received well. Later, Amazon apologized and issued refunds to the customers Streitfeld (2000).

From the regulatory standpoint, while PP is generally not illegal, it may conflict with other acts such as consumer protection laws. As Organisation for Economic Co-operation and Development (OECD) reports, PP is problematic when it uses "techniques that are deceptive or misleading, lack transparency and are implemented without user choice or consent OECD (2020)." Firms intending to implement PP must disclose that in their data privacy policy. While most customers would never read through these policies Deloitte (2017), disclosing PP implementation risks customer setbacks. PP may also conflict with anti-discrimination laws, which protect customers against discrimination based on attributes such as gender, race, age, and nationality. Although firms may not directly target customers based on these attributes, customer valuations may be correlated with them Wallheimer (2018). For example, when Princeton Review charged different prices for its SAT tutoring course based on ZIP codes, it targeted Asians twice as likely to receive a high price as Non-Asians Angwin et al. (2016).

Due to these implementation and regulatory challenges, many firms remain wary of implementing PP Wallheimer (2018). In a consumer market study, European Commission (2018) did not find significant evidence of PP for identical products: they only observed it in 6% of the situations with a median price difference of less than 1.6%. US Council of

Economic Advisers (2016) reports that the extensive use of PP remains relatively limited among US companies.

Firms also increasingly implement data-driven pricing in response to market conditions and inventories (Elmaghraby and Keskinocak 2003, Chen et al. 2021a, Keskin et al. 2022). For example, they may offer non-personalized promotions to all or a fraction of customers to boost demand and increase revenue. The prices also reflect the supply-demand mismatch, where poor-performing products are sold at lower prices. For example, in the airline industry, the prices are constantly updated based on the availability of the seats and demand for the routes (Gallego and Van Ryzin 1994). In retail, prices may vary based on inventories or pre-scheduled promotional events such as Holiday sales. While customers generally accept these reactive pricing strategies, they are averse to being targeted solely for their willingness-to-pay (Reinartz et al. 2018). With many price changes and non-transparent algorithmic pricing, it would be hard for customers to distinguish the price drivers. In a survey collected from European countries, more than a third of participants indicated concerns about a profile being made on them for online pricing, and 28% indicated they might end up paying more (European Commission 2018). In a experiment, European Commission (2018) found that 40% of the customers who received a high price correctly identified whether prices were personalized. When the participants were told their history was used, this percentage increased to 50%. Notably, over half of the customers remained unaware of the PP implementation. Thus, it is unclear whether price alone can serve as an instrument to signal PP.

Nonetheless, customers have become sophisticated in their shopping: they may forgo purchasing from firms they do not trust (Taylor 2004) and take action to anonymize their shopping or game the system (Conitzer et al. 2012, European Commission 2018, Wallheimer 2018). For example, they may use VPNs, delete their cookies, and refuse to log in before seeing the price if they suspect their personal information will be used for PP.

2.1.1 Contribution and Methodology

While PP and its value for the firm have received growing attention in the operations literature (Ban and Keskin 2021, Chen et al. 2021c, Elmachtoub et al. 2021), whether customers can detect and react to PP when firms reflect the market condition and product availability into prices is unknown. In this work, we study this problem and address whether the price can informatively signal PP to customers. We also study how firms should adjust their pricing strategies when customers strategically react to PP practices. Finally, we investigate whether disclosing product availability information to customers can benefit firms and customers.

To address these questions, we build a dynamic Bayesian game of repeated interactions between a firm and a continuum of heterogeneous customers over two periods. Due to the uncertainty in the firm's high-level strategy and potential regulatory restrictions, the firm and customers are a priori uncertain about the feasibility of PP before the first period. We refer to a firm that can personalize the prices as *P-type*, and a firm that must charge uniform prices to all customers as *U-type*. The firm and customers are also uncertain about the inventory availability. At the beginning of the first period, the firm learns its type and inventory and sets the price according to a pre-committed price mapping. The firm does not know the customer valuations a priori; however, it updates its belief about valuations over time using customers' purchase history. A P-type firm then can use this information to personalize the prices in the second period. Customers cannot directly observe the firm's pricing strategy or inventory. Hence, they may be unable to identify whether a high price reflects inventory scarcity or the firm's intention to learn customer valuations for PP. However, after observing the first-period price, they update their beliefs about these uncertainties. We initially consider myopic customers who do not consider the possibility of future PP and purchase if they receive a price lower than their valuations. Later, we extend the model to include strategic customers who consider this possibility when making purchase decisions. Moreover, we allow the firm to disclose information about the product availability by announcing whether inventory is below or above a threshold.

Finally, we consider three extensions to establish the robustness of our insights:

(i) A case where the firm discloses finer inventory information by announcing the precise inventory level.

(ii) A game with more than two periods and nonstationary customer valuations.

(iii) An environment where customer uncertainty regarding inventory availability stems from demand instead of inventories. In this case, we also allow the firm to decide the inventory.

2.1.2 Findings and Managerial Implications

We find that firms must reduce the first-period price to persuade high-valuation customers to purchase in the first period even when they do not implement PP. This is because the mere presence of PP risk makes customers reluctant to reveal their valuations. The firm then must compensate these customers by charging a lower price. Also, we find that firms implement PP more often (for a broader range of inventory realizations) with myopic customers than with strategic customers, leading to higher expected prices in both periods. Among many reasons, strategic customer behavior can explain why PP is not widely adopted in practice. We demonstrate that price alone cannot perfectly signal the firm's PP intention, even to the most strategic customers. Specifically, we identify two regions: (i) When the firm is P-type with high probability, the game has a partial-pooling equilibrium where customers cannot perfectly identify the firm type; however, they can update their beliefs upon observing the first-period price. (ii) When the probability of the firm being P-type is low, the game has a pooling equilibrium where the price does not convey any information about the firm type.

Having found that price does not always informatively signal the firm's type, we investigate the role of disclosing availability information to customers. The prior literature has primarily focused on the role of availability information in creating urgency for customers to purchase early (Yin et al. 2009, Cui et al. 2019, Calvo et al. 2020). In addition to this aspect, we consider whether customers can infer the PP implementation from the availability information. Generally speaking, if customers believe the price is high due to inventory scarcity, they are less likely to associate it with PP. To study this, we consider a class of binary signals where the firm marks the product low on inventory when inventory is below a threshold. Such inventory signals are commonly adopted by retailers such as IKEA and ZARA. We show that an inventory signal can improve the firm revenue only when customers believe the firm is P-type with a high probability. In this case, an inventory signal alleviates customer PP concerns and allows the firm to set prices closer to what is optimal with myopic customers. When the firm's likelihood of being P-type is low, an informative inventory signal can backfire and exacerbate customer concerns and hurt the firm revenue. In this situation, we also prove that fully disclosing the inventory—another commonly used inventory signal—does not outperform the binary signal. Interestingly, we find that an inventory signal benefits the firm only when price alone (in the absence of an inventory signal) can reveal "some" information regarding the firm type. Therefore, disclosing product availability information is a strategic complement to price when alleviating customer PP concerns.

We show that an inventory signal, in addition to the firm, can benefit all customers. High-valuation customers may enjoy lower prices. Low-valuation customers benefit from higher chances of obtaining a unit when the availability signal reduces the prices. In this case, they can afford the item in the first period, and in the second period, they are pooled with high-valuation customers, which increases their chances of obtaining a unit. Furthermore, we identify conditions under which the availability signal a priori benefits all the stakeholders, including the firm and customers. With the growing interest in PP regulation, requiring firms to disclose availability information could be a viable policy to make PP more transparent.

We illustrate that our insights remain when customer uncertainty stems from demand instead of inventory. Especially, we show that price continues not to be a perfect signal for PP. Furthermore, the firm can benefit from disclosing product popularity/availability information. Interestingly, when the firm optimizes inventory, the value of the availability signal increases. This is because the firm stocks less when it can inform customers about product scarcity.

Finally, we investigate a setting where customer valuations change over multiple—possibly more than two—periods. In this setting, we uncover the exploration-exploitation trade-off commonly present in dynamic games with incomplete information. We partition the selling horizon into three intervals: 1) In the first interval, which we label the *dormancy stage*, both firm types set the same price for customers, making all players unable to learn new information. 2) In the *exploration stage*, the P-type firm sets prices to learn customer types, and the U-type firm adjusts its prices to signal uniform pricing. 3) In the *exploitation stage*, the P-type firm implements personalized pricing. We also establish that price cannot perfectly separate the two firm types.

2.2 Literature Review

Price discrimination has received growing attention in the Operations, Marketing, and Economics literature. A body of literature studies intertemporal price discrimination in the presence of strategic customers (Su 2007, Aviv and Pazgal 2008, Cachon and Swinney 2009, Li et al. 2014, Kremer et al. 2017, Aviv et al. 2019, Aflaki et al. 2020). In this literature, price varies over time and customers time their purchases to maximize their utilities. Since the opportunity cost of waiting is higher for high-valuation customers than those with low valuations, the firm can gradually reduce the price to skim the market.

While in these papers customers "self-select" the prices based on their valuations and patience, a stream of work focuses on PP tactics where a firm customizes prices for different customers based on their distinct features (Cohen et al. 2020, Ban and Keskin 2021, Chen et al. 2022). These papers study the statistical properties of algorithms that balance the learning-and-earning trade-off in PP using customer features. Since PP requires granular feature-based learning, it results in a high-dimensional problem requiring special solution treatments. Aydin and Ziya (2009) study dynamic PP of limited inventories when the firm does not have perfect knowledge of the customer valuations; however, it receives a signal correlated with customer valuations. They find that a mere positive correlation between customer valuations and the signal is insufficient to set a high price. Elmachtoub et al. (2021) study the value of PP when the firm has limited knowledge of the distribution of customer valuations. Similar to these papers, we consider PP when the firm does not a priori know the customer valuations and balances the learning-and-earning trade-off. However, we focus on two-sided learning where the firm learns about the customers, and customers learn about the firm and may strategically respond to that.

We are also related to the literature on behavior-based pricing and marketing, where the firm discriminates against customers based on their purchase history. In a multi-period setting, this literature studies various aspects of behavior-based pricing, including the conditions under which price discrimination benefits firms facing strategic customers (Acquisti and Varian 2005), valuation enhancing offers based on customer purchase history (Pazgal and Soberman 2008), and interplay of price discrimination and product differentiation under competition (Jing 2017, Amaldoss and He 2019). Fudenberg and Villas-Boas (2006) surveys this literature. Most of this literature assumes customers have perfect knowledge of the firm's pricing strategy. We investigate this assumption by asking whether price can serve as an instrument to signal the firm's pricing strategy and whether availability information can complement such a signal.

A few papers consider privacy-concerned customers who strategically react to the firm's customer profiling policy. Taylor (2004) studies a situation where a firm learns the customer valuations by selling a product. The firm can then sell the customer information to another firm that personalizes the prices using this information. He finds that strategic customer behavior significantly impacts the firms' payoff and customer surplus. Conitzer et al. (2012) consider the repeated interactions of a monopolistic firm with strategic customers over two periods. The firm uses the customer purchase history from the first period to personalize the prices in the second period. Customers can exert effort to anonymize their purchases. They find that increasing the cost of anonymity might benefit the customers as long as it is sufficiently low. However, it often harms the customers when the firm can set this cost. In our model, customers are uncertain about the product availability and whether the firm implements PP. This two-dimensional uncertainty prevents customers from identifying whether their profile will be used for PP. Furthermore, our work adds the operational aspect of disclosing availability information to customers.

Allon et al. (2012) and Drakopoulos et al. (2021) study the correlation between price and inventory signals. In the presence of inventory availability uncertainty, Allon et al. (2012) show that a separating equilibrium exists in the game where customers perfectly learn the product availability by observing prices. However, the firm may set a suboptimal price to signal availability. Drakopoulos et al. (2021) study a pricing and inventory provisioning game where the firm sets uniform prices for all customers and sends personalized availability signals. They find customized information can significantly benefit the firm and has properties similar to PP. Yu et al. (2015) show that rationing capacity in advance selling can signal product quality. However, the possibility of rationing capacity may harm the seller. Lingenbrink and Iyer (2019) study whether a single-server queue should share information about the queue's state to strategic customers who decide whether to join the queue. They show the optimality of binary signaling mechanisms and that the firm may benefit from strategically concealing information. In addition to product availability (correspondingly queue length), in our model, customers are uncertain whether the firm can or intends to implement PP. In this environment, we study the welfare implications of an availability signal.

A body of literature investigates the role of information provisioning in social learning when customers learn about the product characteristics from other customer purchase experiences through levers such as product reviews and recommendation systems (Besbes and Scarsini 2018, Che and Hörner 2018, Papanastasiou et al. 2018, Acemoglu et al. 2022, Bimpikis and Papanastasiou 2019, Garg and Johari 2019). In the context of inventory information provisioning, Yin et al. (2009) show that displaying one unit of inventory instead of displaying all at once can create a shortage risk and improve the firm profit. Allon and Bassamboo (2011) prove that sharing inventory information with a homogeneous population of customers does not change the intertemporal customer behavior. They identify conditions under which inventory information would not be cheap talk when customers are heterogeneous. Aydinliyim et al. (2017) study the optimal disclosure of inventory availability information in online retailing. They show that an optimal inventory threshold exists below which the firm should disclose the exact inventory level, and above that, it should only disclose availability. Cui and Shin (2018) show that for a retailer selling multiple products, sharing aggregate inventory information could outperform sharing more granular product availability information. Cui et al. (2019) empirically establish that customers learn from availability information and a decrease in inventory availability increases future sales. Chen et al. (2021b) analyze the timing of inventory disclosure policies and show that a threshold policy that discloses the inventory information when inventory drops below a given level outperforms always-disclose and never-disclose policies. Similar to this literature, we study the inventory disclosure policies, and answer whether and under what conditions such signals benefit firms and customers.

2.3 Model

We study a firm selling a product over two periods to a continuum of customers with unit mass. Let M denote the market of customers. Each customer purchases at most one unit of the product in period $t \in \{1, 2\}$, totaling a maximum of two purchases over the two selling periods. We note that the product sold in one period could be different from the other. For example, customers repeatedly purchase flight tickets from travel agencies. These tickets may differ as they could be for different times or flights. In the retail context, customers purchase apparel (e.g., T-shirts) with various designs in different seasons. In this environment, firms can still learn about customers from their purchase history.

Customers are heterogeneous in their valuations for the product. Specifically, Customer $i \in M$ values the product at $v_i \in \{v_L, v_H\}$ in each period, where $v_H > v_L \geq 0$. We refer to customers with $v_i = v_L$ and $v_i = v_H$ as L-Type and H-type, respectively. The customer type is unknown to the firm; however, it has prior knowledge of its distribution. Specifically, a fraction α of customers are H-type, i.e., $\mathbb{P}(v_i = v_H) = \alpha$. Parameter α is common knowledge between the firm and customers. Let M_L and M_H denote the population of L-type and H-type customers, respectively.

As is common in the related literature, we assume that the firm and customers are a priori uncertain about the inventory levels at the beginning of each period (Allon et al. 2012, Drakopoulos et al. 2021). However, they have a common belief about its distribution and density functions denoted, respectively, by $F(\cdot)$ and $f(\cdot)$ supported on $\mathcal{I}$. While inventory levels are uncertain before the start of the first selling period, we assume their realizations are the same in both periods, denoted by I.[1] We furthermore assume I has a uniform distribution on interval $[\alpha, 1]$. That is because a fraction α of customers are H-type, and the market has a unit mass.[2] At the beginning of the first period, the firm learns the inventory realization while customers remain uncertain. This creates information asymmetry between

[1]We make this assumption mainly for tractability to focus on the pricing and information provisioning game dynamics between the firm and customers without ad-hoc treatment of various sub-game realizations. This is a reasonable assumption as we assume consistent demand over the two periods and consider firm's high-level pricing strategy.

[2]We also considered a variation where inventory is a binary random variable. The results remain qualitatively the same.

the firm and customers.

Although we refer to the supply of products as "inventory" throughout the paper, it may represent other forms of capacity, such as the number of seats in a flight. As we will see, inventory availability (henceforth *availability*), i.e., the probability that customers can obtain a unit, is the primary driver of the results rather than the absolute value of the inventory. In §2.6.1, we formally confirm this observation by considering a variation of the model where uncertainty stems from demand instead of inventory. We analytically show that our insights continue to hold. We also numerically establish the robustness of our results when the firm optimizes inventory.

The firm can implement *uniform pricing (UP)* or *personalized pricing (PP)*. With UP, the firm charges the same price to all customers. With PP, however, different customers may receive different prices. Specifically, the firm may use the customers' first-period purchase history to learn their valuations and target them with customized prices in the second period. As discussed in the introduction, firms face uncertainty over the possibility of using PP due to their high-level pricing strategy, and regulatory and implementation challenges. Let $\omega \in \Omega = \{0, 1\}$ denote this possibility: when $\omega = 1$, the firm is allowed to personalize the prices, while when $\omega = 0$, it cannot implement PP. The firm and customers do not know the state of ω before the first selling period. However, they have a common belief about its realization. Suppose $\mathbb{P}(\omega = 1) = \tau \in (0, 1)$. We refer to the firm of type $\omega = 1$ as P-type (for personalized pricing) and the firm with $\omega = 0$ as U-type (for uniform pricing). Like inventory, at the beginning of the first period, the firm learns the state of ω while customers remain uncertain, creating information asymmetry. Here we distinguish between the "possibility" of PP and the actual "implementation" of that: While a P-type firm can implement PP, it does not necessarily benefit from it. In other words, a P-type firm may find it optimal to implement UP.

Upon observing price p_{it}, Customer i decides whether to purchase the product or opt for her outside option. Let $a_{it} \in \{0, 1\}$ be the indicator function for the customer purchase decision, where $a_{it} = 1$ if the customer purchases in period t. Moreover, let $p_t = \{(p_{it}) : i \in M\}$ and $a_t = \{(a_{it}) : i \in M\}$ be period t vector of prices and customer decisions, respectively. In the next section, we discuss the details of the sequence of events and player decisions.

We focus on a two-period problem with two customer segments to analyze the complex dynamics of a two-sided learning game with information asymmetry. However, when customer valuations are non-stationary or in a game with more than two customer segments, the game may continue beyond two periods. §2.6.2 considers an extension with non-stationary customer valuations and a general finite selling horizon. We show that our framework is not bound to a two-customer-segment two-period problem.

2.3.1 Sequence of Events, Information Sets, and Signaling Mechanism

In each period, the firm and customers learn about the game dynamics and take actions according to their information sets. Let $\mathcal{H}_t^f$ and $\mathcal{H}_{it}^c$, respectively, denote the information sets of the firm and Customer i up until the beginning of period t. The dynamics of the game unfold as follows:

Period -1: Before the game starts, all players are uncertain about ω and I. Thus, $\mathcal{H}_{-1}^f = \varnothing$. Similar to Drakopoulos et al. (2021), we assume the first-period price mapping is set prior to the realization of uncertainty, i.e., the firm commits to the first-period price mapping $p_1(\omega, I) : \Omega \times \mathcal{I} \to [v_L, v_H]$ before learning ω and I. This assumption captures that the pricing strategy is designed in advance using prior information about the inventory and possibility of PP; however, the actual price depends on the precise realizations of ω and I when selling the product.

Period 0: The states of ω and I realize, and the firm sets price p_1 according to price mapping $p_1(\omega, I)$. We first consider a base model, where the firm does not disclose any information to customers other than the price. Later, we allow the firm to signal availability. An *inventory signaling mechanism* is a mapping $\Sigma : \mathcal{I} \to \mathcal{S}$, which discloses information about the inventory to customers. For most of the analysis, we focus on a class of binary signaling mechanisms, where the firm marks the product low-on-inventory when the inventory is below a threshold I_c. For example, IKEA labels the products that are unlikely to run out during the day as "in stock" and the products that may run out of inventory as "low in stock." ZARA marks the items short on inventory as "last items in stock." We study if such a signaling mechanism can benefit the firm and customers when customers infer information about the firm's pricing strategy. When this binary inventory signal is ineffective in improving the firm revenue, we also consider a signaling mechanism where the firm fully discloses the inventory level. Similar to the prices, we assume the firm commits to signaling mechanism Σ before the realization of inventory (Drakopoulos et al. 2021). However, signal $s \in \mathcal{S}$ is sent after the inventory realizes in period 0. In this period, the firm and customer information sets are $\mathcal{H}_0^f = \{\omega, I\}$ and $\mathcal{H}_{i0}^c = \varnothing$.

Period 1: All customers enter the market and stay until the last period. Customer i observes price p_1 and inventory signal s (if any inventory information is disclosed). Based on this information, she decides whether to purchase the product, i.e., she determines $a_{i1}(p_{i1}, s) \in \{0, 1\}$. In period 1, the firm and customers update their information as $\mathcal{H}_1^f = \{\omega, I, s\}$ and $\mathcal{H}_{i1}^c = \{v_i, p_1, s\}$. We assume any unsold inventory at the end of period 1 is liquidated at price zero.

Period 2: The firm observes the customer purchase decisions and sets the second-period prices. Let $p_{i2}(\omega, I, \mathcal{H}_2^f)$ be the second-period price charged to Customer i when the firm is of type ω and carries inventory I, where $\mathcal{H}_2^f = \{\omega, I, s, a_1(p_1, s)\}$. We note that if $\omega = 0$, we must have $p_{i2}(\omega, I, \mathcal{H}_2) = p_{j2}(\omega, I, \mathcal{H}_2)$ for all $i, j \in M$. In other words, a U-type firm cannot implement PP. However, for $\omega = 1$, the firm can use the customer purchase history $\mathcal{H}_2^f$ to personalize the price. We let $p_2(\omega, I, \mathcal{H}_2^f) = \{p_{i2}(\omega, I, \mathcal{H}_2^f) : i \in M\}$ be the set of all prices offered to customers in period t. Any unsold inventory at the end of the second period is disposed of at price zero.

2.3.2 Belief Updating and Inventory Availability

Both the firm and customers update their prior beliefs upon receiving new information. The firm updates its belief about customer valuations after observing their purchase decisions. Intuitively, if the firm charges price v_L, it cannot infer any information about the customer type. However, if a customer purchases the product at a price greater than v_L, the firm would learn that the customer is H-type. Accordingly, the firm updates its belief about the customer types following the Bayes' rule.

Customers update their beliefs about the firm type upon observing the price and inventory signal. Since they are also uncertain about inventory, they may not precisely learn the firm type upon receiving the first-period information. For example, a firm charging price $p_1 = v_H$ may do so because the inventory is low, and can sell out even at the highest viable price. Alternatively, a P-type firm may charge a higher price than v_L to learn the customer valuations and personalize the prices in the second period.

Also, since the firm has limited inventory, customers' decisions to purchase do not guarantee they would receive a unit. Let $D_t(p_t) = \int_{i \in M} a_{it}(p_{it}, s)\, di$ denote the total demand in period t. If $I < D_t(p_t)$, then a fraction of customers would not receive the product. In the case of insufficient inventory, we assume the firm first prioritizes the customers who pay higher prices (Su and Zhang 2008, Cachon and Feldman 2015). For the customers who pay the same price, we assume the inventory is randomly allocated between the customers demanding the product (Dana Jr and Petruzzi 2001, Aflaki and Swinney 2021). Let $\widetilde{\xi}_{it} \in [0, 1]$ be the probability that Customer i would be allocated a unit in period t.

Note that the availability depends on the firm type and inventory as well as prices charged to all customers. If the firm sets uniform price v_L, all customers demand the product, and the probability that a customer receives the product would be $\widetilde{\xi}_{it} = I$ for all $i \in M$ and $t \in \{1, 2\}$. Also, if a firm sets a uniform price higher than v_L, then L-type customers cannot afford the product. In this case, H-type customers are guaranteed to receive a unit. Finally, if the firm implements PP, H-type customers receive the product with certainty as they are prioritized over the L-type customers. In contrast, L-type customers are randomly allocated a unit from the leftover products. Thus, we have $\widetilde{\xi}_{i2} = 1$ for $i \in M_H$ and $\widetilde{\xi}_{i2} = \frac{I-\alpha}{1-\alpha}$ for $i \in M_L$.

Since the inventory and firm type are unknown a priori, the customers and firm form beliefs about the availabilities. Let $\xi^c_{it}(\mathcal{H}^c_{ij})$ and $\xi^f_{it}(\mathcal{H}^f_{ij})$ be Customer i and the firm beliefs about the availability in period t at the beginning of period j. These beliefs are updated in each period following the Bayes' Rule based on the firm and customer information sets.

2.3.3 Customer Behavior and Utility

Customer $i \in M$ who obtains a unit of the product at price p_{it}, receives utility $v_i - p_{it}$. Hence, the customer ex-ante utility-to-go function at the beginning of period $t \in \{1, 2\}$ is given by

$$u_{it}(p_{i1}, p_{i2}, \mathcal{H}^c_{it}) = \mathbb{E}_{I,\omega}\left[\sum_{j=t}^{2} a_{ij}\left(p_{ij}, s\right) \xi^c_{ij}(\mathcal{H}^c_{ij})(v_i - p_{ij}) \middle| \mathcal{H}^c_{it}\right]. \tag{1}$$

If a customer decides to purchase the product at the posted price p_{ij} (i.e., $a_{ij}\left(p_{ij}, s\right) = 1$), and if she is assigned a unit (with probability $\xi^c_{ij}(\mathcal{H}^c_{ij})$), she receives utility $(v_i - p_{ij})$ from her purchase.

Recall that a P-type firm may use the customer purchase history in the first period to learn the valuations and charge a personalized price in the second period. L-type customers never pay a higher price than v_L and are not concerned that the firm would learn their valuations. However, an H-type customer may receive a higher price because of PP. These customers may benefit from not revealing their types to the firm. We consider two customer types: *Myopic customers* do not include their future utility when making first-period decisions. Therefore, they purchase a unit in period t if they receive a positive surplus. *Strategic customers* recognize the possibility of PP and may choose their outside option if doing so would increase their total utility. Without loss of generality, we assume the customer's outside option yields a zero surplus. For example, customers may buy from another seller who does not implement PP but always charges price v_H.[3]

In §2.6.2, we consider an extension where customer valuations may change over time, and the game extends beyond two periods. We demonstrate that customers choosing their outside option in a period may purchase from the firm in the later periods.

A customer choosing the outside option effectively *hides* her type. However, hiding type comes at the risk of not being able to obtain the product in the second period. This is because customers who pay a higher price are prioritized in being allocated a unit, and customers paying v_L enter a lottery to receive the product. Therefore, strategic customers balance the trade-off between (i) utility of the first-period purchase, (ii) the impact on future

[3]Customers may use other methods to anonymize their types. For example, they may refuse to log in, use VPN, or delete their cookies (OECD 2018). This, requires customers to exert costly effort (Conitzer et al. 2012), which reduces their utilities. Our insights continue to hold in such a model.

prices if the firm is P-type, and (iii) the probability of obtaining a unit in the second period. Consequently, they maximize their ex-ante surplus by dynamically making purchase decisions at the beginning of each period. In other words, Customer i solves

$$\max_{a_{it} \in \{0,1\}} u_{it}(p_{i1}, p_{i2}, \mathcal{H}_{it}^c) \text{ for } t \in \{1, 2\}. \tag{2}$$

2.3.4 Firm Behavior and Revenue

The firm considers its inventory, type, and customer behavior when setting the prices and designing the inventory signal. While a U-type firm cannot use the customer purchase history to personalize the prices, it may adjust them to separate itself from P-type. Similarly, a P-type firm may not find it optimal to implement PP if the gain from doing so is not justified by the revenue loss due to strategic customer behavior. Given these trade-offs, the firm revenue-to-go function at the beginning of period t is given by

$$R_t(\mathcal{H}_t^f, p_t) = \mathbb{E}_{I,\omega}\left[\sum_{j=t}^{2} \int_{i \in M} a_{ij}\left(p_{ij}, \Sigma(I)\right) \xi_{ij}^f(\mathcal{H}_j^f) p_{ij}\, di \middle| \mathcal{H}_t^f\right]. \tag{3}$$

The firm earns revenue p_{ij} if Customer i purchases in period j, i.e., $a_{ij}\left(p_{it}, \Sigma(I)\right) = 1$, and is allocated a unit from inventory, which occurs with probability $\xi_{ij}^f(\mathcal{H}_j^f)$. The firm dynamically maximizes its ex-ante revenue by setting prices and the inventory signal (if applicable) for $t \in \{1, 2\}$:

$$\begin{aligned} \max_{p_t, \Sigma(I)} \quad & R_t(\mathcal{H}_t^f, p_t) \\ \text{s.t.} \quad & (p_1, \Sigma) \in \arg\max_{p_1(\omega, I), \Sigma(I)} R_1(\mathcal{H}_{-1}^f, p_t) \\ & (1 - \omega)(p_{i2} - p_{l2}) = 0, \ \forall i, l \in M \ \& \ \omega \in \Omega \end{aligned} \tag{4}$$

The first constraint ensures that the first-period prices are selected from the price menu designed in period -1. The second constraint in Problem (4) guarantees that the firm sets uniform prices if it is U-type. We note that if the firm sends an inventory signal, it also impacts the customers' beliefs about the inventory and firm type, which influence customer decisions $a_{it}\left(p_{it}, s\right)$.

Table 2.1: Caption: Notation Summary (Chapter 2)

Variables	
v_i	Customer $i \in M$ values the product at $v_i \in \{v_L, v_H\}$ in each period, where $v_H > v_L \geq 0$
α	A fraction α of customers are H-type, i.e., $\mathbb{P}\left(v_i = v_H\right) = \alpha$
I	Inventory realization for both periods
ω	$\omega \in \Omega = \{0, 1\}$ denote this possibility: when $\omega = 1$, the firm is allowed to personalize the prices, while when $\omega = 0$, it cannot implement PP
τ	$\mathbb{P}(\omega = 1) = \tau \in (0, 1)$ as the probability that the firm is able to conduct PP
p_{it}	The price for Customer i at period t
a_{it}	$a_{it} \in \{0, 1\}$ be the indicator function for the customer purchase decision, where $a_{it} = 1$ if the customer purchases in period t
$\mathcal{H}_t^f$	Information set of the firm at period t
$\mathcal{H}_{it}^c$	Information set of the Customer i at period t
s	Inventory signal sent by the firm, either s_a (Above) or s_b (Below)
Σ	An inventory signaling mechanism, which is a mapping from the inventory realization to an inventory signal
$D_t(p_t)$	The total demand in period t
$\widetilde{\xi}_{it}$	The probability that Customer i would be allocated a unit in period t
$\xi_{it}^c(\mathcal{H}_{ij}^c)$	Customer i's belief about the availability in period t at the beginning of period j
$\xi_{it}^f(\mathcal{H}_{ij}^f)$	Firm's belief about the availability in period t at the beginning of period j
u_{it}	Customer i's expected utility-to-go at the beginning of period t
R_t	Firm's expected revenue-to-go at the beginning of period t
p_t^m	Equilibrium price at period t when customers are myopic
I^m	Unique equilibrium threshold, under which P-type firm sets prices $p_1^m > v_L$ if $I < I^m$ when customers are myopic
p_t^n	Equilibrium price at period t when customers are strategic
γ	Size for the interval, i.e., $\overline{I} - \underline{I} \equiv \gamma(\tau, \alpha)$, which is a function of τ and α

Variables	(Continued)
I_c	A threshold that marks the inventory as low if $I \leq I_c$ and high if $I > I_c$
$\overline{p}_b$	Optimal pricing function when $s = s_b$, which is a function of I_c
$\overline{p}_a$	Optimal pricing function when $s = s_a$, which is a function of I_c
I_c^*	the equilibrium inventory signal cut-off
p_{it}^*	Equilibrium pricing function for Customer i at period t
CS^n	Expected total customer surplus without inventory signal
CS^s	Expected total customer surplus with inventory signal
Q_L^n	Expected number of units received by L-type customer without inventory signal
Q_L^s	Expected number of units received by L-type customer with inventory signal
s^d	Extension: Demand signal sent by the firm, either s_a^d (Above) or s_b^d (Below)
Σ^d	Extension: A demand signaling mechanism, which is a mapping from the demand realization to an demand signal
α_c	Extension: A threshold that marks the demand as low if $\alpha \leq \alpha_c$ and high if $\alpha > \alpha_c$
$\tilde{v}_{it}$	Extension: Customer i's valuation at period t
λ	Extension: H-type customers would have valuation v_H with probability λ
t^m	A threshold that partitions the selling horizon into two stages. In the exploration stage, for $t \leq t^m$, the firm explores the customer valuations by charging v_H to all customers. In the exploitation stage, for $t > t^m$, the firm personalizes the prices and exploits its knowledge about customers

2.3.5 Solution Concept

We base our analysis on the Bayesian persuasion framework (Kamenica and Gentzkow 2011, Anunrojwong et al. 2023) and study the sender-preferred subgame perfect Bayesian equilibrium (SPBE) of the game. An SPBE consists of prices p_t, an inventory signaling mechanism Σ (if applicable), and customer purchase decisions $a_{it}(p_{it}, s)$ such that for $t \in \{1, 2\}$ and $i \in M$ they satisfy the following conditions:

1) Firms and customers update their beliefs according to the Bayes' rule upon receiving new

information.
2) Myopic customers purchase the product in each period they receive a positive surplus. Strategic customers maximize their ex-ante utility by dynamically solving Problem (2).
3) The firm maximizes its revenue-to-go function by solving Problem (4).
4) The equilibrium is sender-preferred, i.e., the best outcome for the firm among all the possible equilibria.

2.4 Equilibrium Analysis

In this section, we analyze the SBPE of the game between the firm and customers. We initially consider a model with myopic customers to illustrate the PP dynamics without customer concerns. Then, we analyze a model with strategic customers. When customers are strategic, first, we study the equilibrium in a base model where the firm only sets the price and does not send any inventory signal. In the base model, we drop argument s from the notations that depend on inventory signal. For example, we denote the customer purchase decision by $a_{it}(p_{it})$ instead of $a_{it}(p_{it}, s)$. Finally, we consider a scenario where the firm discloses inventory information to customers. Throughout, we use notation $\mathbb{1}(\cdot)$ for the characteristic function.

2.4.1 Myopic Customers

Following the Solution concept in §2.3.5, myopic customers purchase the product in each period they receive a positive surplus without considering the impact of their decisions on future prices. The following proposition characterizes the SBPE of the game with myopic customers.

Proposition 2.1. *A unique SBPE to the game with myopic customers exists. Let $p_1^m(\omega, I)$ and $p_2^m(\omega, I, \mathcal{H}_2^f)$ be the equilibrium first and second-period prices with myopic customers. Then,*
(i) A U-type firm sets prices

$$p_1^m(0,I) = p_2^m(0,I,\mathcal{H}_2^f) = \begin{cases} v_H & \text{if } I \leq \frac{\alpha v_H}{v_L} \\ v_L & \text{otherwise.} \end{cases} \tag{5}$$

(ii) A unique threshold $I^m \in [\frac{\alpha v_H}{v_L}, 1]$ *exists such that a P-type firm sets prices*

$$p_1^m(1,I) = \begin{cases} v_H & \text{if } I \leq I^m \\ v_L & \text{otherwise,} \end{cases} \quad \text{and}$$

$$p_{i2}^m(1,I,\mathcal{H}_2^f) = \begin{cases} \mathbb{1}(v_i = v_H)v_H + \mathbb{1}(v_i = v_L)v_L & \text{if } I \leq I^m \\ v_L & \text{otherwise.} \end{cases}$$

Proof. All proofs appear in Appendix A.1.

With myopic customers, both firm types charge either price v_L or v_H in both periods. For any price $p \in (v_L, v_H)$, all H-type customers purchase the product, while L-type customers are excluded. In this situation, the firm is incentivized to increase the price to improve its revenue without losing demand. Furthermore, for sufficiently a low inventory ($I \leq \frac{\alpha v_H}{v_L}$), a U-type firm is better off to set price v_H and sell only to the H-type customers. However, if the inventory is high ($I > \frac{\alpha v_H}{v_L}$), the revenue loss from the exclusion of the L-type customers dominates the revenue gain from charging a higher price. In this case, the firm sets price v_L in both periods.

Part (ii) of the proposition shows that, similar to a U-type firm, a P-type firm would set price v_H only when inventory is below a threshold ($I \leq I^m$). Again, for sufficiently a high inventory ($I > I^m$), the firm is better off selling to all customers in the first period by setting price v_L. However, this prevents the firm from learning the customer valuations. Moreover, we have $I^m > \frac{\alpha v_H}{v_L}$. This is because setting first-period price v_H in addition to offering a higher revenue for each unit sold (which also benefits a U-type firm) allows the firm to learn the customer valuations and personalize the prices in the second period. This additional value motivates a P-type firm to set first-period price v_H for a larger range of inventory realizations than a U-type firm. This implies that $\mathbb{E}_I[p_1^m(1,I)] > \mathbb{E}_I[p_1^m(0,I)]$, i.e., a P-type firm charges a higher "expected" price in the first period than a U-type firm. This correlation between the first-period price and firm type allows customers to infer information about the firm type and inventory availability by observing the first-period price. This information does not change the behavior of myopic customers as they do not consider their future utilities. In the next section, we consider strategic customers who use this information when making purchase decisions.

2.4.2 Strategic Customers

Following §2.3.5, strategic customers solve Problem (2). In this game, different firm types may have the incentive to differentiate from or mimic the other firm. In §2.4.2.1, we consider a model without any inventory signal. Later, in §2.4.2.2, we allow the firm to signal availability to customers.

2.4.2.1 Base Model: Equilibrium Without an Inventory Signal

In this base model, the firm sets the prices, and customers update their beliefs and make decisions based on these prices. First, we characterize the customer best responses in the following lemma:

Lemma 2.1. *In any SBPE, we have $a_{i2}(p_{i2}) = \mathbb{1}(v_i \geq p_{i2})$. Furthermore, there exists a unique threshold $\overline{p} \in (v_L, v_H)$ such that H-type customers purchase in the first period if and only if $p \leq \overline{p}$. Thus, we have*

$$a_{i1}(p_{i1}) = \begin{cases} \mathbb{1}(p_{i1} \leq v_L) & \text{if } v_i = v_L \\ \mathbb{1}(p_{i1} \leq \overline{p}) & \text{if } v_i = v_H. \end{cases}$$

Furthermore, $\overline{p}$ is non-increasing in τ.

Intuitively, all customers who receive a non-negative surplus would purchase the product in the second period. Also, L-type customers would buy in the first period if they receive a price not greater than their valuations, i.e., price v_L. These customers are not concerned about revealing their valuations as even a P-type firm implementing PP would charge second-period price v_L to these customers. However, H-type customers have the incentive to hide their types to receive the product at a lower price in the second period. Lemma 2.1 shows that these customers would purchase if the price is sufficiently low. If the customers were myopic, all H-type customers would purchase the product at price v_H, and we would have $\overline{p} = v_H$. However, strategic customers should be persuaded to reveal their types. Hence, the firm sets price $\overline{p} < v_H$.

The existence of unique threshold $\overline{p}$ is not a priori trivial. On the one hand, a higher first-period price reduces the customer surplus from purchasing the product in the first period. This motivates customers to choose their outside option. On the other hand, a higher first-period price may signal an inventory shortage, which increases the availability risk. A higher availability risk favors revealing H-type to the firm by purchasing at a high price in the first period. Lemma 2.1 shows that customers purchase in the first period if they receive sufficiently a low price $\overline{p}$. Additionally, as the lemma shows, $\overline{p}$ is non-increasing in τ. Recall that τ is the likelihood of the firm being P-type. Consequently, as τ increases, customer concerns regarding the use of their purchase history for PP increase. This prevents the firm from setting a high first-period price.

Threshold $\overline{p} > v_L$ critically depends on the availability risk. In fact, a P-type firm cannot persuade H-type customers to reveal their types without this risk (Taylor 2004). In our problem, however, the availability risk discourages customers from hiding their types for two reasons: First, hiding type would pool them with L-type customers and decreases their chances of obtaining a unit in the second period. Second, a high price may reflect inventory scarcity, reducing its association with PP.

Having the customer behavior on hand, next, we establish the existence and uniqueness of the SBPE for the game with strategic customers and characterize the equilibrium.

Proposition 2.2. *A unique SBPE to the game exists. Define* $\underline{I} = \min(\frac{\alpha\overline{p}}{v_L}, 1)$ *and* $\overline{I} = \min\left(\frac{\alpha\overline{p}}{v_L} + \frac{\alpha(v_H - v_L)}{v_L}, 1\right)$*, where* $\overline{p}$ *is given in Lemma 2.1. Let* $p_1^n(\omega, I)$ *and* $p_2^n(\omega, I, \mathcal{H}_2^f)$ *be the equilibrium first and second-period prices. Then,*

(i) A U-type firm sets prices

$$p_1^n(0, I) = \begin{cases} \overline{p} & \text{if } I \leq \underline{I} \\ v_L & \text{otherwise,} \end{cases} \quad \text{and} \quad p_2^n(0, I, \mathcal{H}_2^f) = \begin{cases} v_H & \text{if } I \leq \dfrac{\alpha v_H}{v_L} \\ v_L & \text{otherwise.} \end{cases}$$

(ii) A P-type firm sets prices

$$p_1^n(1, I) = \begin{cases} \overline{p} & \textit{if } I \leq \overline{I} \\ v_L & \textit{otherwise,} \end{cases} \quad \textit{and} \quad p_{i2}^n(1, I, \mathcal{H}_2^f) = \begin{cases} a_{i1}(\overline{p})v_H + (1 - a_{i1}(\overline{p}))v_L & \textit{if } I \leq \overline{I} \\ v_L & \textit{otherwise,} \end{cases}$$

and customers follow the strategy characterized in Lemma 2.1.

Similar to the case with myopic customers, a P-type firm implements PP only if the inventory is sufficiently low. However, the inventory threshold below which the P-type firm implements PP is lower with strategic customers than myopic customers, i.e., $\overline{I} \leq I^m$. This is because customer strategic behavior limits a P-type firm's ability to personalize the prices without raising customer concerns. Hence, with strategic customers, the firm can implement PP only when the inventory is sufficiently low such that customers attribute a high price to an inventory shortage. This concern does not exist with myopic customers, which allows implementing PP for a higher range of inventory.

Moreover, we have $\underline{I} \leq \frac{\alpha v_H}{v_L}$. In other words, similar to the P-type firm, the inventory threshold for charging a high price for a U-type firm is also lower with strategic customers than with myopic customers. Therefore, from Propositions 2.1 & 2.2, we find that a U-type firm always sets a higher price with myopic customers than with strategic customers. Although a U-type firm cannot implement PP, it sets a lower price for strategic customers to "differentiate" itself from a P-type firm. In response, a P-type firm decreases its price to "mimic" U-type and reduce customer concerns regarding future PP. A lower first-period price compensates H-type customers for potential loss in their second-period surplus, further encouraging them to purchase in the first period.

Interestingly, both firm types set the same equilibrium price $\overline{p}$ in the first period if they intend to sell only to the H-type customers. As a result, customers cannot perfectly distinguish between the two firm types upon observing first-period price $\overline{p}$. This implies that both firms cannot charge a price higher than v_L without raising customer concerns regarding future PP. However, since $\overline{I} \geq \underline{I}$, a higher price is more likely to be associated with a P-type than a U-type firm. Because of this, customers may be able to "partially" distinguish between the two firm types. Specifically, if $I \in (\underline{I}, \overline{I}]$, the U-type firm sets v_L while the P-type firm sets $\overline{p}$.

Furthermore, for $I > \overline{I}$, although both firm types set the same first-period price v_L, it would not raise customer concerns because v_L is the lowest possible price and prevents the firm from learning the customer valuations. However, for $I \leq \underline{I}$, both firm types set a similar price in the first period while only the P-type firm implements PP. In this environment, customers can only probabilistically distinguish between the two firms, leading to a "partial-pooling" equilibrium. Consequently, the existence of non-empty interval $(\underline{I}, \overline{I}]$ is the only source of information about the firm type. If such an interval does not exist, customers cannot update their beliefs about the firm type. In fact, a larger size for this interval, i.e., $\overline{I} - \underline{I} \equiv \gamma(\tau, \alpha)$, increases the strength of the price signal in determining the firm type. The following corollary is a direct consequence of this discussion and Proposition 2.2.

Corollary 1. *With strategic customers,*
(i) If $\gamma(\tau, \alpha) > 0$, the game has a partial-pooling equilibrium, where the customers update their beliefs about the firm type; however, they cannot perfectly distinguish them.
(ii) If $\gamma(\tau, \alpha) = 0$, the game has a pooling equilibrium where the first-period price does not convey any information about the firm type.

While $\gamma(\tau, \alpha)$ is independent of the inventory realization, it critically depends on parameters τ and α. We study the behavior of $\gamma(\tau, \alpha)$ in these parameters in the next proposition.

Proposition 2.3. *In the game with strategic customers,*
(i) $\gamma(\tau, \alpha)$ is non-decreasing in τ.
(ii) $\gamma(\tau, \alpha)$ initially increases, then decreases, and finally remains constant in α.
(iii) There exists thresholds $\overline{\tau}(\alpha)$ and $\overline{\alpha}$ such that $\gamma(\tau, \alpha) = 0$ for any $\tau \leq \overline{\tau}(\alpha)$ or $\alpha \geq \overline{\alpha}$.

Part (i) shows that the price signal strengthens as τ increases. This is because when τ is small, the customers are not too concerned about the threat of being targeted by PP. This enables both firm types to increase the first-period price, which pools them together and reduces the strength of the first-period price as a signal to separate the firms. The customer PP concerns increase as τ increases beyond a threshold. This makes the U-type firm reduce its price for intermediate inventory realizations to separate itself from the P-type firm and persuade customers to purchase in the first period. The P-type firm, however, continues to charge a high price to extract customer surplus in the second period. Thus, the first-period price becomes a stronger signal of the firm type, and $\gamma(\tau, \alpha)$ increases in τ. For sufficiently a high value of $\tau < 1$, there always exists an inventory interval with a constant length in τ where the two firms charge different prices. Hence, $\gamma(\tau, \alpha)$ remains constant at a positive value as τ increases beyond a threshold.

Part (ii) illustrates that the behavior of $\gamma(\tau, \alpha)$ is not monotonic in α. An increase in α increases the value of learning customer valuations for a P-type firm because it can personalize the prices for a larger population of H-type customers. It also increases the U-type firm's desire to set a high price in the first period. For sufficiently a small α, both firm types serve all customers by setting price v_L, as L-type customers comprise a large market population and excluding them is costly. This is specifically true for a U-type firm as the firm cannot utilize the customer information in the second period. Therefore, as α initially increases, the P-type firm's incentive to set a high price increases faster than the L-type firm's incentive. This asymmetry in the firm incentives leads to the initial increasing behavior of $\gamma(\tau, \alpha)$ as a function of α. As α increases beyond a threshold, the P-type firm charges $\overline{p}$ for all inventory values, while the U-type firm continues to set price v_L when the inventory is sufficiently large. An increase in α beyond this threshold encourages the U-type firm to set $\overline{p}$ for a larger inventory range and become more similar to the P-type firm. As the two firm types set the same prices for more inventory realizations, the price becomes a weaker signal to differentiate the firms. Consequently, after this threshold, $\gamma(\tau, \alpha)$ decreases in α until it eventually becomes zero. In this case, i.e., for $\alpha \geq \overline{\alpha}$, the game accepts only a pooling equilibrium.

2.4.2.2 Equilibrium With an Inventory Signal

As discussed earlier, uncertainty over two dimensions, i.e., firm type and inventory, prevents customers from perfectly identifying if a high price is a result of inventory scarcity or the firm's intention to learn the customer valuations for PP.[4] In this section, we allow the firm to partially resolve this uncertainty by disclosing inventory availability information.[5]

[4] Inventory scarcity is relative to demand. In §2.6.1, we consider a scenario where demand uncertainty is the root of availability risk.

[5] Firms also disclose whether they implement PP in their data privacy policies. However, these disclosures are often lost in long documents not read by customers. For example, a survey by Deloitte (2017) showed that 91% of customers accept legal terms and conditions without reading them for installing applications and using online services. This number is likely higher in retail. Thus, we study inventory availability

We consider a class of binary signaling mechanisms similar to the examples of IKEA and ZARA discussed in §2.3.1. Specifically, we assume that the firm sets a threshold I_c and marks the inventory as low if $I \leq I_c$ and high if $I > I_c$ (we use subscript c for cut-off). Such an inventory signal has the form

$$\Sigma^c(I) = \begin{cases} s_b & \text{if } I \leq I_c \\ s_a & \text{otherwise,} \end{cases} \tag{6}$$

where s_b is sent when inventory is "below" threshold I_c and s_a is sent when inventory is "above" this threshold. We refer to s_b and s_a as *LI* (for Low-Inventory) and *HI* (for High-Inventory) signals, respectively. Since the structure of the signal is designed before the inventory realization, it is a priori stochastic. However, the actual signal is sent according to function $\Sigma^c(I)$ after the realization of uncertainty.

In this binary mechanism, customers immediately infer the HI signal upon not observing LI. Hence, this signaling mechanism is equivalent to a mechanism where the firm sends signal s_b when inventory is low and does not send any signal when it is high. Furthermore, signaling mechanism (6) does not convey any information when $I_c = 1$. Therefore, the model with this signaling mechanism is a generalization of the base model without the inventory signal studied in §2.4.2.1.

First, we characterize the customer behavior in the following lemma.

Lemma 2.2. *In any SBPE, Customer i who receives price p_{i2} and signal $s \in \{s_a, s_b\}$ makes decision $a_{i2}(p_{i2}, s) = \mathbb{1}(v_i \geq p_{i2})$ in the second period. Additionally, there exist unique thresholds $\overline{p}_a(I_c) \in (v_L, v_H)$ and $\overline{p}_b(I_c) \in (v_L, v_H]$ such that H-type customers purchase the product in the first period if and only if $p_1 \leq \overline{p}_b(I_c)$ when they receive signal s_b and $p_1 \leq \overline{p}_a(I_c)$ when they receive signal s_a. Hence,*

$$a_{i1}(p_1, s) = \begin{cases} \mathbb{1}(p_1 \leq v_L) & \text{if } v_i = v_L \\ \mathbb{1}(p_1 \leq \overline{p}_b(I_c)) & \text{if } v_i = v_H \text{ and } s = s_b \\ \mathbb{1}(p_1 \leq \overline{p}_a(I_c)) & \text{if } v_i = v_H \text{ and } s = s_a. \end{cases}$$

information often utilized by sellers as an operational proxy to signal the pricing strategy.

Furthermore, for any signaling mechanism $\Sigma^c(I)$*, there exist thresholds* $\underline{I}_b(I_c)$*,* $\overline{I}_b(I_c)$*,* $\underline{I}_a(I_c)$*, and* $\overline{I}_a(I_c)$ *as functions of* I_c *such that*

$$\overline{p}_b(I_c) = \begin{cases} v_H - \frac{\tau}{2} \cdot \frac{I_c-\alpha}{1-\alpha} \cdot (v_H - v_L) & \text{if } I_c \in [\alpha, \underline{I}_b(I_c)] \\ v_H - \frac{\tau}{2} \cdot \frac{(I_c-\alpha)^2}{1-\alpha} \cdot \frac{(v_H-v_L)}{(1-\tau)\underline{I}_b(I_c)+\tau I_c-\alpha} & \text{if } I_c \in (\underline{I}_b(I_c), \overline{I}_b(I_c)] \\ \overline{p} & \text{if } I_c \in (\overline{I}_b(I_c), 1], \end{cases}$$

where $\overline{p}$ *is given in Lemma 2.1. Also,*

$$\overline{p}_a(I_c) = \begin{cases} v_H - \frac{\tau(\overline{I}_a(I_c)-I_c)}{(1-\tau)\underline{I}_a(I_c)+\tau\overline{I}_a(I_c)-I_c} \cdot \frac{I_c+\overline{I}_a(I_c)-2\alpha}{2(1-\alpha)}(v_H - v_L) & \text{if } I_c \in [\alpha, \underline{I}_a(I_c)] \\ v_H - \frac{I_c+\overline{I}_a(I_c)-2\alpha}{2(1-\alpha)}(v_H - v_L) & \text{if } I_c \in (\underline{I}_a(I_c), \overline{I}_a(I_c)] \\ v_H - \frac{2\overline{I}_a(I_c)-2\alpha}{2(1-\alpha)}(v_H - v_L) & \text{if } I_c \in (\overline{I}_a(I_c), 1]. \end{cases}$$

Finally, $\overline{p}_b(I_c)$ *and* $\overline{p}_a(I_c)$ *are non-increasing in* I_c*, and* $\overline{p}_a(I_c) \leq \overline{p} \leq \overline{p}_b(I_c)$*.*

Similar to the base model, H-type customers purchase the product in the first period only if the price is lower than a threshold. However, this threshold depends on the inventory signal they receive. As illustrated in the lemma, $\overline{p}_a \leq \overline{p} \leq \overline{p}_b$. Intuitively, an HI signal increases the likelihood of the price being high due to PP. Hence, the firm must reduce the price to compensate H-type customers for this possibility. In contrast, an LI signal correlates with a higher likelihood of a high price in response to an inventory shortage. Moreover, when inventory is limited, H-type customers, by hiding their types, enter a lottery with L-type customers to obtain a unit, which significantly reduces their chances. Because of these reasons, the firm can increase its first-period price without the risk of losing H-type customers.

Lemma 2.2 also shows that $\overline{p}_a(I_c)$ and $\overline{p}_a(I_c)$ are both non-increasing in I_c. This is because, as I_c increases, an LI signal becomes a weaker indicator of inventory scarcity. Also, an HI signal becomes a stronger indicator of product availability. Consequently, independent of the signal sent to the customers, H-type customers associate a high price with a higher likelihood of PP implementation. Therefore, they require a lower price to purchase in the first period as I_c increases. Figure 2.1 illustrates this behavior.

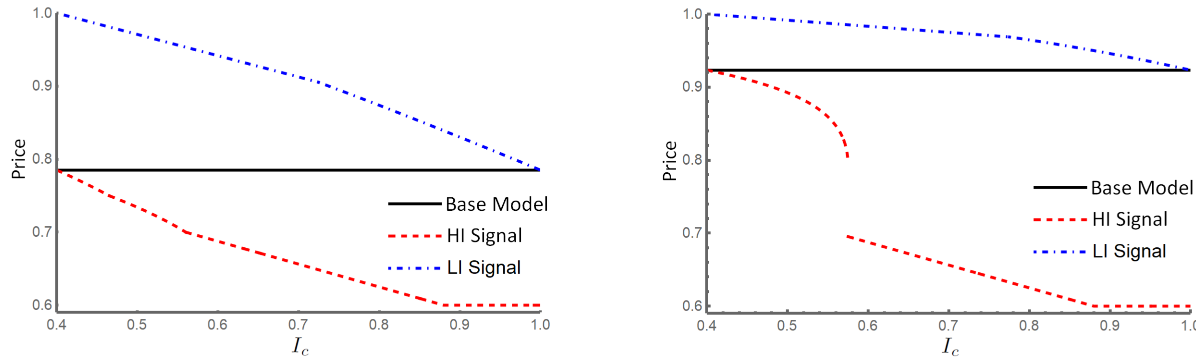

(a) Prices when $\tau = 0.7$.

(b) Prices when $\tau = 0.2$.

Figure 2.1: Prices $\overline{p}$, $\overline{p}_b(I_c)$, and $\overline{p}_a(I_c)$ as functions of signal cut-off I_c, when $v_H = 1$, $v_L = 0.5$, and $\alpha = 0.4$.

Let $p_1^*(\omega, I)$ and $p_2^*(\omega, I, \mathcal{H}_2^f)$ be the first and second-period equilibrium prices, respectively. Also, let I_c^* be the equilibrium inventory signal cut-off and $\Sigma^*(I)$ be the corresponding signaling mechanism. Moreover, we use superscript * for the values of the functions defined in Lemma 2.2 calculated at point I_c^*. For example, we define $\overline{I}_b^* = \overline{I}_b(I_c^*)$ and $p_b^* = \overline{p}_a(I_c^*)$.

We assert the following proposition:

Proposition 2.4. *An SBPE to the signaling game between the firm and customers exists. Furthermore,*

(i) A U-type firm sets,

$$p_1^*(0, I) = \begin{cases} p_b^* & \text{if } I \leq \min(\underline{I}_b^*, I_c^*) \\ p_a^* & \text{if } I_c^* < I \leq \underline{I}_a^* \\ v_L & \text{otherwise,} \end{cases} \quad \text{and} \quad p_2^*(0, I, \mathcal{H}_2^f) = \begin{cases} v_H & \text{if } I \leq \dfrac{\alpha v_H}{v_L} \\ v_L & \text{otherwise.} \end{cases}$$

(ii) A P-type firm sets,

$$p_1^*(1, I) = \begin{cases} p_b^* & \text{if } I \leq \min(\overline{I}_b^*, I_c^*) \\ p_a^* & \text{if } I_c^* < I \leq \overline{I}_a^* \\ v_L & \text{otherwise.} \end{cases}$$

Additionally,

$$p^*_{i2}(1, I, \mathcal{H}^f_2) = \begin{cases} a_{i1}(p^*_b, s_b)v_H + (1 - a_{i1}(p^*_b, s_b))\, v_L & \text{if } I \leq \min(\overline{I}^*_b, I^*_c) \\ a_{i1}(p^*_a, s_a)v_H + (1 - a_{i1}(p^*_a, s_a))\, v_L & \text{if } I^*_c < I \leq \overline{I}^*_a \\ v_L & \text{otherwise,} \end{cases}$$

where customers follow the strategy characterized in Lemma 2.2.
(iii) We have $\overline{I}^*_b \geq \underline{I}^*_b$ and $\overline{I}^*_a \geq \underline{I}^*_a$.

Proposition 2.4 shows that while the first-period price "values" are the same for both firm types (i.e., they both set either v_L, $\overline{p}^*_a$, or $\overline{p}^*_b$), the "inventory intervals" under which these prices are set may be different. Specifically, the proposition establishes that $\overline{I}^*_b \geq \underline{I}^*_b$ and $\overline{I}^*_a \geq \underline{I}^*_a$. Thus, for any inventory signal, a P-type firm sets a high price for a larger range of inventory realizations than a U-type firm. This observation combined with Lemma 2.2 implies that a P-type firm, in equilibrium, sets a higher price than a U-type firm. Hence, customers can use a combination of price and inventory signals to update their beliefs about the likelihood of future PP.

While Proposition 2.4 establishes the existence of an SBPE, it does not guarantee that an inventory signal benefits the firm: when $I^*_c = 1$, the inventory signal does not communicate any information, and the firm achieves the same revenue as the base model. A firm may choose $I^*_c = 1$ because an informative inventory signal may lead to a lower first-period price, which can harm both firm types. The following proposition finds the conditions under which a binary inventory signal improves the firm revenue.

Proposition 2.5. *A unique $I^*_c < 1$ exists if and only if $\tau > \overline{\tau}(\alpha)$, where $\overline{\tau}(\alpha)$ is defined in Proposition 2.3. Furthermore, $\overline{\tau}(\alpha)$ is increasing in α.*

Condition $I^*_c < 1$ represents the case where an inventory signal benefits the firm; otherwise, the firm could set $I^*_c = 1$ to achieve the same revenue as the base model. Thus, this proposition shows that an inventory signal improves the firm revenue only when $\tau > \overline{\tau}(\alpha)$. When τ is large, the firm's likelihood of being P-type is high. Therefore, firms intending to price higher than v_L must compensate customers by offering a sufficiently low price to persuade them to purchase in the first period. An inventory signal, in this case, helps customers differentiate between the U-type and P-type firms, which enables them to set a higher price than the base model. To illustrate, note that by setting a sufficiently high signal cut-off $I_c^* < 1$, the firm sends an LI signal with a high probability which enables it to set price $p_b^* > \overline{p}$ and improve its revenue. Also, an HI signal is sent when inventory is large, in which case firms set price v_L even in the base model. Consequently, an HI signal does not significantly reduce the firm expected revenue. In sum, an LI signal improves the revenue, and an HI signal does not substantially harm the firm. As such, inventory signal benefits the firm when τ is large.

In contrast, when τ is small, customers' PP concerns alleviate. In this case, both firms would behave the same, and an inventory signal does not help separate them. From Proposition 2.3, recall that $\overline{\tau}(\alpha)$ is the threshold such that for $\tau \leq \overline{\tau}(\alpha)$ price alone cannot even partially separate the firms, i.e., $\gamma(\tau, \alpha) = 0$. Hence, customers do not learn about the firm type by observing the first-period price. Moreover, in this case, Proposition 2.5 shows that $I_c^* = 1$. In other words, an inventory signal does not convey any information to the customers either. Therefore, this proposition implies that an inventory signal benefits firms only when customers can "complement" their inventory information with prices.

Furthermore, the proposition shows that $\overline{\tau}(\alpha)$ is increasing in α. In other words, as the fraction of H-type customers increases, the inventory signal benefits the firm for a smaller range of high τ values. This is because an increase in α increases a U-type firm's incentive to set a high price in the first period and exclude the L-type customers. Simultaneously, it increases the PP value for a P-type firm. In both cases, i.e., a high uniform or personalized price, inventory becomes a less determinant factor of the firm's pricing strategy than customer valuations. Consequently, an inventory signal does not convey much information to customers regarding the firm type and pricing strategy when α is high.

Figure 2.2 plots the regions of τ and α where an inventory signal improves the firm revenue. We observe that when α is sufficiently low ($\alpha \lesssim 0.5$), the inventory signal improves the firm revenue for all values of τ. This follows because $\overline{\tau}(\alpha)$ is an increasing function of α and condition $\tau > \overline{\tau}(\alpha)$ translates to condition $\alpha < \overline{\alpha}$ for some threshold $\overline{\alpha}$. Also, when α is sufficiently large ($\alpha \gtrsim 0.67$), both firm types only sell to the H-type customers. In this case, H-type customers do not have much incentive to hide their types, as doing so would significantly reduce their chances of obtaining a unit. Therefore, the firm does not need to

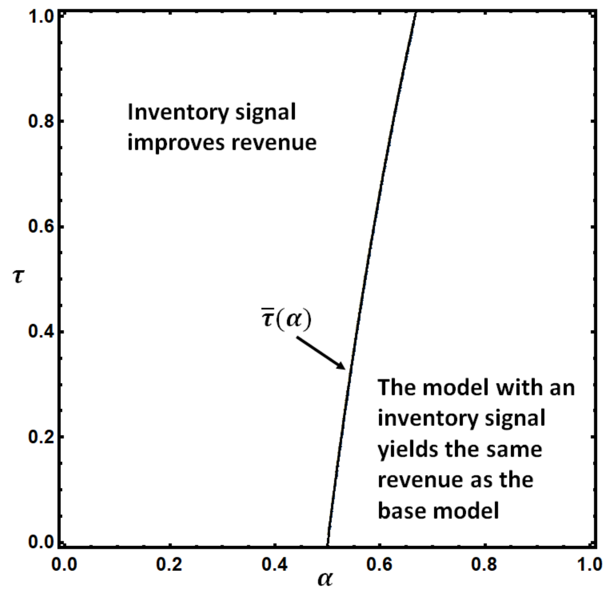

Figure 2.2: Regions of τ and α where an inventory signal improves the firm revenue. In this figure, we have $v_H = 1$ and $v_L = 0.5$.

send an informative inventory signal to persuade customers to reveal their types. Thus, for the entire range of τ, the firm with an inventory signal achieves the same revenue as the base model. However, for intermediate values of α ($0.5 \lesssim \alpha \lesssim 0.67$), a binary inventory signal only benefits the firm when $\tau > \overline{\tau}(\alpha)$, as shown in Proposition 2.5.

In § A.2, we consider a variation of the inventory signaling mechanism where the firm fully discloses the inventory level to customers. We show that when the binary inventory signal does not improve the firm revenue, a full-disclosure signal also does not convey any information to the customers and does not benefit the firm.

In sum, a binary inventory signal benefits the firm when τ is sufficiently large or when α is sufficiently low. These are the cases when customers use the price to update their beliefs about the firm type. As such, an inventory signal is a complement rather than a substitute to price when informing customers about PP. Next, we study the customer surplus.

2.5 Customer Surplus

In this section, we study how an inventory signal impacts customers. L-type customers only buy the product at price v_L. Hence, they would always receive a zero surplus. The total customer surplus is then given by

$$CS^n = \int_{i \in M_H} u_{i0}(p_1^n(\omega, I), v_H, \mathcal{H}_{i0}^c)\, di = \alpha \times u_{i0}(p_1^n(\omega, I), v_H, \mathcal{H}_{i0}^c) \text{ for } i \in M_H,$$

in the base model, and

$$CS^s = \int_{i \in M_H} u_{i0}(p_1^*(\omega, I), v_H, \mathcal{H}_{i0}^c)\, di = \alpha \times u_{i0}(p_1^*(\omega, I), v_H, \mathcal{H}_{i0}^c) \text{ for } i \in M_H,$$

in the model with an inventory signal. Recall that $u_{i0}(\cdot)$ is the ex-ante utility Customer i receives over both periods before period 1. Therefore, integrating $u_{i0}(\cdot)$ over $i \in M_H$ yields the customer surplus.

Since L-type customers receive a zeros surplus, the CS does not capture the impact on these customers whether they purchase or not. To study the impact on these customers, we also consider the expected number of units an L-type customer receives during the two periods. For Customer $i \in M_L$, this expected number is given by

$$Q_L^n = \mathbb{E}_{I,\omega}[a_{i1}(p_{i1}^n)\xi_{i1}^c(\mathcal{H}_{i0}^c) + a_{i2}(p_{i2}^n)\xi_{i2}^c(\mathcal{H}_{i0}^c)],$$

in the base model, and

$$Q_L^s = \mathbb{E}_{I,\omega}[a_{i1}(p_{i1}^*, s)\xi_{i1}^c(\mathcal{H}_{i0}^c) + a_{i2}(p_{i2}^*, s)\xi_{i2}^c(\mathcal{H}_{i0}^c)],$$

in the model with an inventory signal. Note that both values of $\xi_{i1}^c(\mathcal{H}_{i0}^c)$ and $\xi_{i2}^c(\mathcal{H}_{i0}^c)$, i.e., the availabilities in the first and second periods, are calculated based on the customer information before the first period, i.e., $\mathcal{H}_{i0}^c$. Hence, Q_L calculates the expected number of units an L-type customer receives before observing the first-period price and inventory signal.

Three interplaying forces determine the impact of an inventory signal on customers: (i) the first-period price directly influences the surplus of the H-type customers. Also, L-type customers can only purchase the product if the price is set at v_L. (ii) The first-period price impacts the likelihood of customers receiving the product. If the firm charges price v_L in the first period, all customers can purchase the product. However, H-type customers compete with L-type customers to randomly receive a unit. In contrast, if the firm charges a first-period price higher than v_L, all H-type customers attempting to purchase the product would obtain a unit. At the same time, the remaining items are randomly assigned to the L-type customers. (iii) A price higher than v_L enables the firm to learn the valuations of those who purchased in the first period. Hence, a P-type firm can fully extract the customer surplus in the second period. This also reduces the L-type customers' chances of receiving a unit as they receive less priority than the higher-paying H-type customers. Albeit, PP increases the affordability of a unit for L-type customers in the second period as they would receive price v_L. The following proposition characterizes the impact of the signaling game on customers.

Proposition 2.6. *Let $\overline{\alpha} = \overline{\tau}^{-1}(\tau)$. There exists a unique threshold $\underline{\alpha} \leq \overline{\alpha}$ such that*
(i) If $\alpha < \underline{\alpha}$, then $CS^s > CS^n$ and $Q_L^s > Q_L^n$.
(ii) If $\underline{\alpha} < \alpha < \overline{\alpha}$, then $CS^s < CS^n$ and $Q_L^s < Q_L^n$.
(iii) If $\overline{\alpha} \leq \alpha$, then $CS^s = CS^n$ and $Q_L^s = Q_L^n$.

Figure 2.3 plots the regions identified in the proposition. Interestingly, besides the firm, an inventory signal can benefit all customers through a higher surplus for the H-type customers and a higher expected number of products for the L-type customers. This happens specifically when α is sufficiently low, i.e., $\alpha < \underline{\alpha}$. In this case, PP does not offer a significant value to the firm because only a small population of customers are H-type, and knowing their valuations does not significantly add to the second-period revenue. Hence, when $\alpha < \underline{\alpha}$, a P-type firm implements PP only if the inventory is sufficiently low. To benefit from this, the firm sets a low I_c^* so that an LI signal conveys significant information about the availability. For example, when $I_c^* = 0$, an LI signal perfectly informs customers about a stockout. Simultaneously, an HI signal would not contain much information about the inventory. In the example with $I_c^* = 0$, an HI signal does not communicate any information to the customers. However, a low I_c^* increases the likelihood of sending an HI signal. In other words, although a low I_c^* weakens the HI signal, it occurs with a higher probability. This leads to a lower price than the base model, as shown in Lemma 2.2. A lower price then benefits the H-type

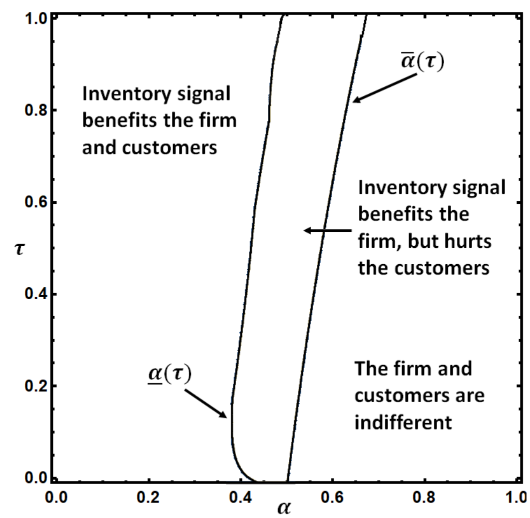

Figure 2.3: Regions of τ and α where an inventory signal improves the firm revenue and customer surplus. In this figure, we set $v_H = 1$ and $v_L = 0.5$.

customers and customer surplus.

Furthermore, in this case, when inventory has an intermediate value, the firm sets first-period price v_L with an inventory signal whereas it charges $\overline{p}$ in the base model. In this region, L-type customers have a positive chance of receiving the product in the first period with an inventory signal, while they cannot afford it in the base model. Additionally, in this range, the first-period price $\overline{p}$ in the base model enables the firm to learn the customer valuations and implement PP in the second period. However, in the model with inventory signal, first-period price v_L prevents the firm from learning customer valuations. In this situation, H-type customers are prioritized in the base model, while they are pooled with L-type customers in the model with inventory signal. This increases the chances of L-type customers receiving the product in the second period. As such, L-type customers may have a higher chance of obtaining a unit in "both" periods as the firm sends an inventory signal. Note that this reduces the product availability for the H-type customers. However, Proposition 2.6 shows that the effect of a lower first-period price dominates the effect of lower availability and the customer surplus increases when α is small. In fact, we observe that H-type and L-type customers benefit from an inventory signal for the same range of

α. This is because both customer segments benefit from an inventory signal only when it reduces the expected first-period price.

As α increases beyond threshold $\underline{\alpha}$, the inventory signal enables the firm to learn the customer valuations at a higher price than the base model, which reduces the customer surplus. Also, as H-type customers purchase the product at a high price, they would be prioritized over the L-type customers in the second period. This reduces the chances of L-type customers obtaining a unit in the second period. Additionally, the firm's tendency to set a higher first-period price eliminates the L-type customer's ability to purchase in the first period. Hence, for $\alpha \in (\underline{\alpha}, \overline{\alpha})$, both customer types are worse-off as the firm sends an inventory signal. This is the region specified in Part (ii) of the proposition.

As α exceeds threshold $\overline{\alpha}$, the firm sets $I_c^* = 1$ as shown in Proposition 2.5, which leads to the same game dynamics with and without the inventory signal. Thus, both customer types would be treated equally in both models.

Finally, we note that for $\alpha < \overline{\alpha}$, we have $\tau > \overline{\tau}(\alpha)$ as given in Proposition 2.5. This is the region where the firm benefits from sending an informative inventory signal, i.e., a signal with $I_c^* < 1$. Since $\underline{\alpha} \leq \overline{\alpha}$, we find that an inventory signal benefits "all" the stakeholders, including the firm and customers when $\alpha < \underline{\alpha}$. In light of growing interest in PP regulations, requiring firms to disclose such inventory information can be a viable policy design, particularly for markets populated with low valuation customers (i.e., low α) and high possibility of PP (i.e., high τ) where inventory availability disclosure benefits all the stakeholders.

2.6 Extensions

In this section, we consider two extensions to establish the robustness of our results and modeling approach. For brevity, we only discuss the main insights in this section.[6]

[6] The details of the analyses, including the supporting results, are available upon request.

2.6.1 Demand Uncertainty and Optimal Inventory

Consider a case where the firm and customers are a priori uncertain about α, i.e., the fraction of H-type customers. Parameter α captures the popularity of the product and determines the demand for any price higher than v_L. Let $f_\alpha(\cdot)$ and $F_\alpha(\cdot)$ be the density and distribution functions of α. We assume that the firm learns the precise value of α at the beginning of the selling horizon while customers remain uncertain. This reflects that firms are often more informed about demand than their customers through market research. We continue to assume that $I_1 = I_2 = I$ to focus on the high-level implications of implementing PP. Furthermore, we assume α has a uniform distribution on interval $[0, 1]$. We consider the same sequence of events as in §2.3.1.

In this section, an availability signaling mechanism maps α to a measure of product popularity. For example, the firm can mark the product as "trendy" or "selling fast." The firm can also signal availability by sharing the fill rate or the service level information. Let $\Sigma^d : \alpha \to \mathcal{S}^d = \{s_b^d, s_a^d\}$ be this signaling mechanism where s_b^d and s_a^d are the signals sent when the product availability is high ($\alpha < \alpha_c$) and low ($\alpha \geq \alpha_c$), respectively.

In this model, we analytically replicate most of the results from the main model with modifications to the equilibrium values. Specifically, the customers' purchase decisions follow a threshold policy where they purchase the product if the price is below a threshold. Also, a P-type firm implements PP if the product is sufficiently popular, i.e., when α is sufficiently high. Moreover, we find that price alone cannot serve as a signal to entirely separate the two firm types.

When the firm can send a signal regarding product availability, we establish the existence of the equilibrium. We also show that the expected revenue function has at most two local maximizers in α_c such that one is below I and the other is I. When $\alpha_c = I$, the firm merely signals whether inventory is sufficient to meet the demand of all H-type customers. However, it does not inform customers regarding the degree of product availability. This strategy would particularly be optimal when τ and I are relatively large. In this case, the firm is P-type with high probability and has the incentive to personalize the price. Therefore, it benefits from not communicating too much information about the product availability.

However, when τ is small, the firm benefits from informing customers about shortage situations, which is achieved by setting a low value for α_c. Figure 2.4 plots the ex-ante expected revenue function as a function of the signal cut-off α_c and illustrates this discussion.

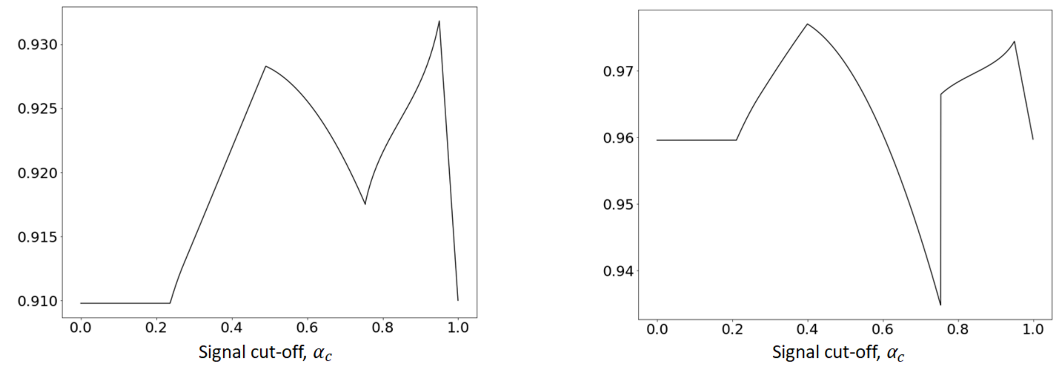

(a) $\tau = 0.7$.

(b) $\tau = 0.45$.

Figure 2.4: Ex-ante expected revenue as functions of signal cut-off α_c, when $v_H = 1$, $v_L = 0.1$, and $I = 0.95$.

Finally, we numerically investigate the optimal inventory level with and without an availability signal. We observe that the firm stocks less inventory when it uses an availability signal. This increases the shortage situations, which will be communicated to customers through the availability signal.

Interestingly, when the firm optimally sets inventory, the value of using an availability signal increases. This observation is not a priori intuitive: setting inventory enables the firm to reduce stocks and signal shortage via price even without any availability signal. This undermines the value of an availability signal. However, such a signal allows the firm to convey information about the inventory shortage more effectively. This re-purposes the role of prices from signaling availability to matching customer valuations. Figure 2.5 shows this behavior.

2.6.2 Customer Non-Stationary Valuations and a $T-$Period Model

In the main model, we considered two customer segments in a two-period model. In such a model, if a customer purchases the product at a price higher than v_L, she reveals her type

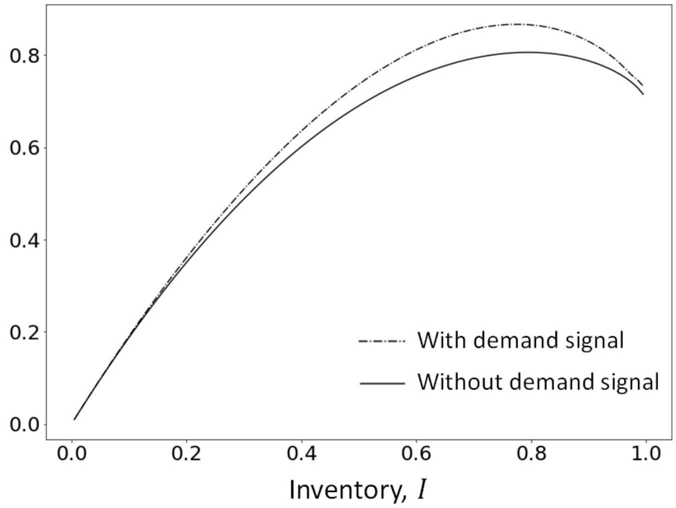

Figure 2.5: Optimal Ex-ante expected revenues with and without demand signaling as functions of inventory I, when $v_H = 1$, $v_L = 0.1$, and $\tau = 0.7$.

to the firm. In this environment, we found that some customers may hide their types in the first period. However, this result should not be interpreted as customers permanently hide their types until the last period when there are more than two periods. This is because, first, when the customer valuations are non-stationary, or there are more than two segments, purchasing a product does not fully reveal the customers' private valuations. Additionally, by hiding her type, a customer forgoes a positive utility, which might not be optimal in any period. This section shows how our model could be extended to more than two periods when customer valuations are non-stationary.

Consider $T > 0$ selling periods and a continuum of customers with unit mass. Similar to our main model, Customer i has valuation $\tilde{v}_{it} \in \{v_L, v_H\}$ in period t. We consider two customer types: L-type customers (denoted by set L) always possess valuation v_L, i.e., $\tilde{v}_{it} = v_L$ for all $t \in \{1, 2, ..., T\}$ and $i \in L$. However, H-type customers (denoted by set H) would have valuation v_H with probability λ and v_L with probability $1 - \lambda$ (Taylor 2004). In other words, $\mathbb{P}(\tilde{v}_{it} = v_H) = \lambda$ for $i \in H$. For example, a student traveler may rarely be willing to pay a high price for airfare. An H-type traveler, however, may be willing to pay a premium for a business travel. The same customer may not be willing to pay a high price for a leisure trip.

We assume that $\lambda v_H \geq v_L$. Otherwise, the firm always charges v_L to all customers. Also, let α denote the fraction of H-type customers. The firm and customer types as well as "valuations" are private, but their distributions are common knowledge. The sequence of events unfolds following §2.3.1. Additionally, customers learn their valuations at the beginning of each period, and the game continues beyond period 2. Also, since customer valuations change over time, the firm can only explore customer types rather than their valuations directly. In this setting, our main model with two periods and stationary customer valuations is a special case of this model with $T = 2$ and $\lambda = 1$.

The uncertainty regarding firm and customer types introduces an exploration-exploitation trade-off for both the firm and customers, where they initially take actions to learn about the other player's type and subsequently make decisions based on their acquired knowledge to maximize payoff. We establish that, for a P-type firm facing myopic customers, there exist threshold t^m that partitions the selling horizon into two stages. In the exploration stage, for $t \leq t^m$, the firm explores the customer valuations by charging v_H to all customers. In the exploitation stage, for $t > t^m$, the firm personalizes the prices and exploits its knowledge about customers.

For a P-type firm facing strategic customers, there exist thresholds $t_1^s \leq t_2^s$ and $\bar{p}_t$ for $t \in \{1, 2, ..., T-1\}$ such that the selling horizon is partitioned into three intervals. In the dormancy stage, for $t \leq t_1^s$, the firm charges v_L to all customers. In the exploration stage, for $t \in (t_1^s, t_2^s]$, the firm explores customer valuations by setting price $\bar{p}_t$ for the customers who have not purchased before and price v_H for previous purchasers. In the exploitation stage, for $t > t_2^s$, the firm sets price v_H for previous H-type purchasers and v_L for the other customers.

Comparing the cases with myopic and strategic customers, we observe that in addition to the exploration and exploitation stages, there is a stage when the firm charges the lowest possible price to all strategic customers. This is the stage when strategic customers intend to learn about the firm type. Also, any purchase at a high price enables a P-type firm to extract customer surplus for the rest of the selling horizon, which is costly for customers at the early stages. Thus, these customers would only purchase at a price higher than v_L once they are assured the firm is U-type and the utility lost due to hiding their types dominates the cost of revealing their types to a P-type firm.

The above observations suggest that our two-period model in the main part of the paper could be interpreted as a reparametrization of time into the exploration and exploitation stages after a period of dormancy.

2.7 Conclusions

We studied a pricing and inventory information provisioning game between a monopolistic firm and a market of heterogeneous customers. The firm has repeated interactions with customers and can use their purchase history to learn their valuations and personalize the price. Customers are uncertain about the product availability and the possibility of being targeted by PP. In this environment, we investigated a previously unanswered question: Can price serve as an instrument to signal the implementation of PP? We demonstrate that when high-valuation customers comprise a significant population of customers and when PP is anticipated, the price alone does not communicate any information about the PP implementation. We also find that customer concerns hurt the firm revenue even if the firm does not intend to implement PP. This is because the firm must reduce the price to alleviate these concerns. In this situation, we prove that an inventory availability signal can help customers identify the use of PP. This additional information can alleviate customer concerns and help improve the firm revenue. We show that an inventory signal only benefits the customers when the price is also an informative signal of PP. Thus, price and inventory signal are complements in informing customers about the PP implementation.

While most of the literature focuses on the role of inventory information in creating purchase urgency for customers, our paper uncovers a new role for the disclosure of such information: communicating the firm's pricing strategy and its intention to use the customer purchase history for PP. Our paper also has important policy implications. There has been growing interest in regulating PP due to customers' concerns and their potential surplus loss. Preventing firms from implementing PP might hurt the stakeholders. Our results suggest that an operational intervention such as requiring disclosure of inventory information under some circumstances can benefit all the stakeholders, including firms and customers.

Our work shows that communicating inventory availability information as a signal for firm's pricing strategy has important implications for firms, customers, and policymakers. In this work, we focused on the optimal binary inventory signals. We also studied full inventory disclosure signals when the binary signal does not improve the firm revenue. One can explore the optimal signal structure for the inventory disclosure policies in the space of all possible signaling mechanisms. Furthermore, future work can investigate the interplay of PP with signals other than inventory availability. For example, product reviews that correlate with quality can signal the implementation of PP. Also, future empirical work can verify our theoretical findings and the role of inventory in resolving uncertainty about the firm's pricing strategy.

3.0 Rating Systems under Customer Disconfirmation Bias: Asymptotic Behavior and Granularity

Customers and platforms increasingly rely on online ratings to assess the quality of products and services. However, customer ratings are susceptible to various biases. *Disconfirmation bias* is a specific form where customers incorporate the discrepancy between their prior expectations and post-purchase experiences into their ratings. We study the asymptotic behavior of ratings in the presence of disconfirmation bias in three rating systems: (i) complete system, where customers observe the entire rating history; (ii) aggregate system, where only the frequency of each rating option is available; and (iii) average ratings, where customers solely use the average of past ratings. Customers are Bayesian and update their quality beliefs upon observing the ratings. After experiencing the product, they rate it according to their heterogeneous ex-post utility and disconfirmation bias. In complete and aggregate systems, we show that customer beliefs converge to the intrinsic quality when disconfirmation bias is small. When this bias is large, there will be a discrepancy between converged beliefs and the intrinsic quality, although this discrepancy could be arbitrarily small. When the disconfirmation bias is intermediate, beliefs may diverge significantly from the intrinsic quality or not converge. However, we establish that the platform can guarantee correct learning by designing a sufficiently granular rating system, i.e., a system with more rating options. We confirm all these results in the system with average ratings, albeit with a bias-correcting rule. Finally, we characterize the learning speed in the aggregate system.

3.1 Introduction

Whether choosing restaurants, booking hotels, or shopping, customers increasingly rely on online reviews to gauge the quality of the products and services. In a survey, 89% of global customers stated they check online reviews when shopping (Trustpilot 2020) and 77% of customers look for websites with reviews and ratings (PowerReviews 2023). While reviews are often more comprehensive evaluations and can include written articulations, ratings are typically simpler expressions of customer experience using a numerical score. For example, many platforms such as Amazon use a 5-star system that allows customers to rate the products on a 1-5 scale. IMDb, an online database for digital entertainment, allows users to rate movie titles on a 1-10 scale while Steam, a video game distribution service, collects and reports positive-negative ratings. Critical questions in this environment are whether ratings can correctly reflect the underlying quality of the product or service and whether the design of the rating system affects this reflection.

Although customers have subjective preferences for product horizontal attributes such as color, the vertical attributes such as quality are objective. In other words, all customers would perceive a higher quality product as superior to a product with lower quality. If customers were to incorporate only these vertical characteristics in their ratings, we could naturally expect ratings to be an unbiased estimator of intrinsic quality. However, customers are not necessarily objective and may include their subjective preferences when rating their experience. For example, Besbes and Scarsini (2018) consider customers who experience quality differently and Acemoglu et al. (2022) study customers who rate based on their heterogeneous ex-post utilities. In this environment, it is not clear whether ratings can informatively uncover the intrinsic quality. Specifically, in these models, customers inevitably exhibit a selection bias: only those who derive sufficiently high utility purchase the product. These customers tend to have a higher ex-post utility, which can inflate their ratings. Similarly, a higher rating can motivate lower-valuation customers to purchase, which can lower the ratings.

Nevertheless, the selection bias is the result of rational purchasing choices of heterogeneous customers. However, customers also exhibit cognitive biases due to their prior expectations. Particularly, empirical work has established that customers with inferior experience than their prior expectations are more likely to rate the product negatively and those whose experience exceeds their expectations tend to rate more positively (Talwar et al. 2007, Ho et al. 2017). This well-established behavioral phenomenon is usually referred to as *disconfirmation* or *expectation* bias in various contexts. The term disconfirmation (as opposed to confirmation) in this context is used because customer ratings positively correlate with the sign of the discrepancy between their prior belief and actual experience.

3.1.1 Contribution and Methodology

In this work, we study the behavior of ratings in the presence of customer disconfirmation bias and its interplay with various rating systems. Specifically, customers who purchase the product rate based on the weighted average of two terms: (i) their ex-post utility which depends on the intrinsic quality, idiosyncratic heterogeneity, and price; and (ii) discrepancy between their prior quality belief and post-purchase experience. We refer to this weighted average as *the reference utility*. The weight of the average between the two terms determines the strength of the disconfirmation bias. In one end of the spectrum, customers purely rate their ex-post heterogeneity aligned with majority of the literature. In the other end, customers purely rate based on their disconfirmation bias which is endogenously formed upon observing the ratings. In the context of rating systems, this bias has been the subject of empirical studies, where researchers examine its magnitude and direction. To the best of our knowledge, we are the first to analytically study the evolution of ratings under the disconfirmation bias and the effect of various rating systems on this evolution.

We consider rating systems with an arbitrary finite set of rating options. Customers are Bayesian and form beliefs about intrinsic quality upon observing the ratings. We consider three rating systems: In the first system, customers observe the individual ratings of all past customers. We refer to this model as the *complete system*. In the second model, which we call the *aggregate system*, customers only observe the frequency of each rating option. Lastly, we study a system where customers solely analyze average ratings. In each of these rating systems, we study the asymptotic behavior of the ratings and the effect of the disconfirmation bias on this behavior. We also examine the implications of the granularity of the rating system, i.e., the number of rating options available to customers. Finally, we investigate the speed of convergence in the aggregate system both when the customer heterogeneity is low and high.

While we make assumptions to filter out the selection bias and focus on disconfirmation bias, our framework allows us to generalize in multiple directions. For example, unlike most literature, we do not require significant richness in customer heterogeneity. This enables us to derive distinct insights about the interplay of disconfirmation bias, customer heterogeneity, and granularity of the rating system. Also, we consider a relatively general rating system and distribution for the prior quality beliefs. A general distribution for quality beliefs enables us to even out the space of rating options and interpret the granularity of the rating system in the language of the number of rating options for customers.

3.1.2 Findings and Managerial Implications

We establish the following results:

1) When the customer disconfirmation bias is sufficiently small, the customer beliefs would eventually converge to the intrinsic quality in the complete and aggregate systems. In this environment, customers can incorporate these biases into their beliefs and uncover the intrinsic quality. Furthermore, in the system with average ratings, beliefs converge to a value that is strictly increasing in the intrinsic quality. In other words, average ratings enable customers to rank the quality of different products correctly. However, the converged beliefs have a systematic error due to limited information or cognitive ability. Since this is a systematic error, we find a bias-correcting rule that can fix it.

2) When the disconfirmation bias is large, all three models lead to biased quality beliefs. In this case, the disconfirmation bias is significant to the extent that customers can no longer distinguish between others' heterogeneous experiences and their cognitive biases. This results in a discrepancy between customer beliefs and the intrinsic quality. However, we show that this discrepancy is small for a large disconfirmation bias. Specifically, in the extreme case when customers solely reflect their disconfirmation bias rather than ex-post utility, the beliefs correctly converge to the intrinsic quality. In this setting, customers are so cognitively biased that they do not consider their ex-post heterogeneity, making it easier for future customers to uncover the effect of disconfirmation bias.

3) When the disconfirmation bias is intermediate, all three systems can lead to a large gap between customer beliefs and the intrinsic quality. In this case, customers' reference utility and rating choices convey the least information to future customers.

4) While disconfirmation bias can lead to incorrect learning, the platform can guarantee correct learning by designing a sufficiently granular rating system, i.e., by giving more rating options to customers. Notably, the granularity required to ensure correct convergence decreases in the degree of customer heterogeneity and is proportional to $1/(1-\alpha)$, where α is the weight of customer disconfirmation bias. In other words, the platform does not need to significantly increase granularity unless disconfirmation bias is very high. Interestingly, however, the case of high disconfirmation bias corresponds to when the gap between quality beliefs and intrinsic quality would be small (as discussed in Insight (2)). Hence, the platform may prefer to keep the number of rating options limited at the expense of a slight discrepancy.

5) We characterize the learning speed under the aggregate system. When heterogeneity is small, we show that there exist ratings that can separate two given quality values when they arise. We formulate a lower bound on the expected number of customers required for such a rating to occur as a function of the distribution of customer heterogeneity, disconfirmation bias, and the rating system's granularity. We show that this lower bound is monotonically increasing in the granularity of the rating system. However, it may be increasing or decreasing in the disconfirmation bias. Specifically, when customer beliefs are too far from the intrinsic quality, the disconfirmation bias can slow down learning. In contrast, when customer beliefs are close to the intrinsic quality, an increase in disconfirmation bias makes learning faster.

When customer heterogeneity is large, we show that customer learning is exponentially fast. We also bound the learning speed using Kullback-Leibler (KL) divergence between the probability of different ratings conditioned on various beliefs. Using these bounds, we observe that the granularity of the rating system does not necessarily translate to faster learning.

The remainder of the paper is organized as follows: In §3.2, we review the related literature. §3.3 introduces the model setup, including the customer rating behavior and rating system design. §3.4-3.6 study the asymptotic behavior of the ratings in the complete, aggregate, and average rating systems. The learning speed in the aggregate system is analyzed in §3.7. Finally, §3.8 concludes the paper.

3.2 Literature Review

Social learning, i.e., when decision-makers learn from others' actions, has received growing attention in the Economics and Operations literature. The seminal works of Banerjee (1992) and Bikhchandani et al. (1992) illustrate the possibility of an informational cascade and herd behavior wherein rational agents mimic the actions of their peers under social learning. This environment was later extended in various directions to incorporate factors such as heterogeneous preferences (Smith and Sørensen 2000), imperfect information (Çelen and Kariv 2004, Herrera and Hörner 2013), network effects (Acemoglu et al. 2011, Mossel et al. 2014, Lobel and Sadler 2015), and non-Bayesian learning (Jadbabaie et al. 2012).

When customers do not directly observe the purchase decisions and quality experiences of others, reviews and online ratings can serve as an instrument to convey information to future customers. A body of literature empirically studies online ratings. Using an experiment, Hu et al. (2006) show that online ratings may not converge to the intrinsic quality. Chevalier and Mayzlin (2006) find improved reviews for books can increase sales. Li and Hitt (2008) study the effect of self-selection bias on the evolution of ratings. Using the case of books on Amazon, they argue that marketing strategies should target customers who are more likely to write positive reviews to encourage them to purchase early.

Anderson and Sullivan (1993) show that quality satisfaction is not solely impacted by customer pre-purchase expectations; Instead, it is best captured by the disconfirmation between the perceived quality and expectations. Talwar et al. (2007) study the behavior of ratings using numerical and textual review analysis. They show that customer ratings partly reflect customers' discrepancies between intrinsic quality and their prior expectations. Ho et al. (2017) formalizes this in the context of disconfirmation bias. Using a hierarchical Bayesian model and data from an e-commerce website, they establish a positive disconfirmation bias in customer ratings. Motivated by these empirical findings, we analytically study the evolution of ratings and the impact of rating system design on this evolution.

Crapis et al. (2017) consider customers who are uncertain about the product quality and learn from a two-scale rating system where previous customers either like or dislike the product based on their ex-post utility. Customers are non-Bayesian and estimate the quality using a Maximum Likelihood approach. In this setting, they establish the convergence of quality beliefs to the intrinsic quality. They also find that various pricing policies that incorporate learning can increase the firm revenue. Besbes and Scarsini (2018) study sequential ratings of Bayesian customers who report their precise ex-post utility. They focus on two information availability schemes where customers either observe all past ratings or their sample mean. They find that customers learn the intrinsic quality asymptotically when observing the entire rating history; however, their estimate may be biased with the sample mean. Ifrach et al. (2019) focus on a binary rating system where purchasing customers would like the product if they receive a positive ex-post utility and dislike it otherwise. They prove that customer quality beliefs converge to the intrinsic quality as long customers continue purchasing the product. Since the firm's pricing influences customers' purchase decisions and ratings, they also consider the pricing problem and show that a single price is optimal if the set of possible prices is finite. The dynamic pricing with social learning has been further studied by Yu et al. (2016), Papanastasiou and Savva (2017), Shin et al. (2023), and Stenzel et al. (2020).

Among other mechanisms to control the evolution of social learning, Papanastasiou et al. (2014) consider a two-stage game where, in the first stage, the product is sold to customers who would then rate the product. The ratings inform future customers about the product quality in the second period. They show that strategic stockouts in the first stage can benefit the firm by boosting ratings when higher-valuation customers have a higher chance of obtaining the product in the first stage. Their model assumes a positive correlation between customer ex-ante expectations and their ratings. Maglaras et al. (2023) study the problem of ranking products when displayed to customers who incur a search cost. Using a fluid approximation, they establish that customers can learn the product quality and compare the performance of various ranking policies.

The closest to our work is Acemoglu et al. (2022) who study the behavior of ratings under a general rating system and customer selection bias. This bias stems from ex-ante and ex-post customer heterogeneity: Higher ratings can motivate customers with lower ex-ante beliefs to purchase the product. These customers, may end up with a low ex-post utility, motivating them to rate negatively. In the same vein, only those customers who derive sufficiently high utility may purchase the product. These customers tend to have higher ex-post utility, motivating them to rate positively. While their model considers the customers' selection bias, they do not incorporate the disconfirmation bias. Hence, customers ex-post utility used for ratings is independent of other customer actions. In other words, "past actions affect player t's inference, not her payoff" (Ifrach et al. 2019). In contrast, our work considers the disconfirmation bias. In our model, a customer's reference utility—the weighted average of the ex-post utility and disconfirmation bias—depends on her inference before the purchase decision. To filter out the direct effect of the selection bias, we assume customers are ex-ante homogeneous but ex-post heterogeneous. Nevertheless, we find that disconfirmation bias causes intertemporal ex-ante heterogeneity among customers. In other words, depending on the order in which customers arrive on the platform, they receive a different payoff because of the different information they receive from ratings. Our approach further enables us to generalize in a few directions. For example, in contrast to the literature that requires sufficiently high heterogeneity, we also consider scenarios where customer heterogeneity is small. For example, we allow situations where customers purely rate products based on their disconfirmation bias without any additional ex-ante or ex-post heterogeneity. We also do not require the intrinsic quality to have a binary domain. In this environment, we explicitly find the interplay between the disconfirmation bias, customer heterogeneity, and granularity of the rating system.

A workstream studies the firm decisions when customers exhibit various cognitive traits. Recent examples include dynamic pricing with reference effects (Chen et al. 2017), customer loss aversion (Chen and Nasiry 2020), gain-seeking behavior (Hu et al. 2016), and rational myopia (Aflaki et al. 2020). In our model, customers are Bayesian and use ratings to make rational purchase decisions. However, they exhibit disconfirmation bias when rating the product. In the context of ratings, Guan et al. (2020) consider customers who incorporate quality references when rating. In a two-period model, they study whether the firm benefits from disclosing assessments of first-period purchasers to future customers. In their model, the mere act of disclosure serves as a signal for quality. We also consider quality references in our customer rating behavior. However, our focus is on the design of rating systems and the evolution of ratings in a multi-period model.

3.3 Model

We consider a platform selling a product with unknown quality to a population of customers. Let Q be the product's intrinsic quality. The platform and customers are a priori uncertain about Q. Let $\tilde{Q}$ denote the ex-ante product quality supported on a bounded interval $\mathcal{Q}$ with density and distribution functions $f_Q(\cdot)$ and $F_Q(\cdot)$, respectively. We assume $F_Q(\cdot)$ is continuous and strictly increasing. Without loss of generality, we normalize $\mathcal{Q}$ to $[0, 1]$.

The platform has a rating system that allows customers to reflect on their experiences by rating the product on a numerical scale. We formally define this rating system in §3.3.1. Customers arrive sequentially over time. We refer to a customer who arrives at time $t \in \{1, 2, ..., \infty\}$ as Customer t. Upon observing the information provided by the platform, Customer t updates her belief about the product quality. Since customers may observe different information over time, their beliefs about the product quality may vary. Let $\mathcal{I}_t^c$ be Customer t's information set and $\tilde{Q}_t(\mathcal{I}_t^c)$ be her (posterior) quality belief upon receiving information $\mathcal{I}_t^c$. Furthermore, let $Q_t^c(\mathcal{I}_t^c)$ denote the expected value of the quality belief, i.e.,

$$Q_t^c\left(\mathcal{I}_t^c\right) = \mathbb{E}\left[\tilde{Q}_t(\mathcal{I}_t^c)\right].$$

To simplify notation, we drop argument $\mathcal{I}_t^c$ from $\tilde{Q}_t$ and Q_t^c. Additionally, we drop sub-index t when not emphasizing the sequence at which the customer has arrived.

We assume customers value the product at its quality. Thus, Customer t, purchasing the product at price p, receives "ex-ante" mean utility

$$u_t^{ea}(\mathcal{I}_t^c) = Q_t^c - p. \tag{7}$$

A customer purchases the product if she receives a positive ex-ante utility, i.e., $u_t^{ea}(\mathcal{I}_t^c) \geq 0$. Let $Z_t(p, \mathcal{I}_t^c)$ be the indicator function for Customer t's purchase decision, where $Z_t(p, \mathcal{I}_t^c) = 1$ if she purchases the product, and $Z_t(p, \mathcal{I}_t^c) = 0$, otherwise. After the purchase, customers learn the intrinsic quality of the product.

To focus on the implications of disconfirmation bias, we assume customers have homogeneous valuations a priori. However, they are intertemporally heterogeneous since they possess different information sets depending on their arrival times. This feature enables us to filter out the selection bias and, as we will see, isolate the effect of the disconfirmation bias on the evolution of the rating system. This implies that if a customer finds it optimal not to purchase the product at any time, all subsequent customers would also stop purchasing. In the presence of ex-ante heterogeneity, it is often argued that customers cannot learn from those who do not purchase (Besbes and Scarsini 2018). Therefore, the primary role of ex-ante heterogeneity would be its effect on the selection bias when customers rate the product. In the absence of ex-ante heterogeneity, we do not require such arguments. In our setting, either all customers purchase or learning stops at a given time.

Furthermore, we allow ex-post heterogeneity for non-quality related attributes. For example, a customer may find a specific color of a piece of apparel stylish, only to use it less often after purchase. Let $\tilde{\theta}$ be the ex-post idiosyncratic preferences with density and distribution functions $f_\theta(\cdot)$ and $F_\theta(\cdot)$, respectively. We assume $\tilde{\theta}$ has a zero mean and is symmetric around the mean supported on bounded interval $\left[-\overline{\theta}, \overline{\theta}\right]$ with a continuous and strictly increasing distribution function.

Upon realization Q of $\tilde{Q}$ and θ_t of $\tilde{\theta}$, Customer t receives ex-post utility

$$u_t^{ep}(\theta_t) = Q + \theta_t - p. \tag{8}$$

The ex-post heterogeneity and disconfirmation bias result in varying customer ratings over time. When there is no confusion, we drop the time index from parameter θ and utility functions. Hence, a customer receives ex-post utility $u^{ep}(\theta) = Q + \theta - p$.

Next, we discuss the details of the rating system and customers' rating behavior.

3.3.1 Rating System and Customer Disconfirmation Bias

We consider a platform allowing customers to rate the product on a K-scale system. We refer to K as the *size* of the rating system. To avoid minor ad-hoc treatments, we assume $K = 2k$ for some integer $k \geq 1$. Specifically, a customer can choose rating $r \in R = \{-k, ..., -1, 1, ..., k\}$ after purchasing the product.[1]

[1] We excluded zero from the range of the ratings to align with standard 2-scale rating systems and simplify the representation of some of our equations. However, most findings can be extended to incorporate zero as a rating option with minor modifications.

As discussed in the introduction, customers exhibit disconfirmation bias when reviewing products (Talwar et al. 2007, Ho et al. 2017). In other words, they are more likely to rate the product positively if its quality exceeds their expectations. In contrast, they may reduce their ratings if the product disappoints their prior expectations. Formally, we assume a customer reflects the following metric in her rating:

$$u^r(\alpha, \theta, Q^c) = (1-\alpha)(Q + \theta - p) + \alpha(Q - Q^c), \tag{9}$$

for $\alpha \in [0, 1]$. We refer to $u^r(\alpha, \theta, Q^c)$ as the customer *reference utility.* In this model, α captures the weight of the customer disconfirmation bias. When $\alpha = 0$, customers do not incorporate their prior expectations into ratings and do not directly exhibit disconfirmation bias. Thus, they purely rate based on their ex-post utilities, i.e., $Q + \theta - p$, as in most literature. On the other end of the spectrum, when $\alpha = 1$, customers rate the product solely based on the discrepancy between their prior beliefs and the intrinsic quality, i.e., $Q - Q^c_t$. In reality, customers may consider a combination of their ex-post utility and reference quality when rating the product, i.e., $\alpha \in (0, 1)$.

Since the platform uses a countable (possibly infinite) rating system, customers cannot precisely report their reference utility. Consider arbitrary thresholds $-\infty = \lambda_{-k} < \lambda_{-k+1} < \cdots < \lambda_0 = \lambda_1 = 0 < \cdots < \lambda_k < \lambda_{k+1} = \infty$. Customers rate the product according to the following mapping:[2]

$$r(u^r(\alpha, \theta, Q^c)) = i, \text{ if } u^r(\alpha, \theta, Q^c) \in [\lambda_i, \lambda_{i+1}) \text{ for } i \in \{-k, ..., k\}. \tag{10}$$

For instance, when $k = 1$, customers rate the product positively if $u^r(\alpha, \theta, Q^c) \geq 0$ and negatively otherwise. Also, when $\alpha = 0$ and $k = \infty$, our model captures the case where customers precisely report their ex-post utility (Besbes and Scarsini 2018).

While most of our insights hold for arbitrary thresholds λ_i, when $k \geq 2$, we further assume $\lambda_k = -\lambda_{-k+1} = 1$ and $|\lambda_{i+1} - \lambda_i| = 1/k$. In other words, thresholds λ_i, for $i \in [-k+1, k]$, partition interval $[-1, 1]$ into subintervals of equal length. This assumption allows us to control the granularity of the rating system using a single parameter k. Moreover, this

[2] Acemoglu et al. (2022) show how micro-foundations of customer rating decisions can lead to rating structure (10) under natural assumptions on customer utility.

assumption is not too restrictive as rating thresholds affect customer ratings "relative" to their reference utility, which includes a general distribution for the quality beliefs.

We study the behavior of this rating system under three information availability schemes. First, in §3.4, we consider customers who observe and analyze the entire rating history. In §3.5, we study a variation where customers only observe the frequency of the past ratings. Finally, §3.6 considers the case where customers solely use the average ratings to update their beliefs.

Table 3.1: Caption: Notation Summary (Chapter 3)

Variables	
Q	Product's intrinsic quality
$\mathcal{I}_t^c$	Customer t's information set at period t
Q_t^c	The expected value of the quality belief for Customer t given his information set
p	Price of the product
u_t^{ea}	Ex-ante utility
Z_t	The indicator function for Customer t's purchase decision, where $Z_t(p, \mathcal{I}_t^c) = 1$ if she purchases the product, and $Z_t(p, \mathcal{I}_t^c) = 0$, otherwise
θ	Ex-post heterogeneity for non-quality related attributes. Assume it has a zero mean and is symmetric around the mean supported on bounded interval $\left[-\overline{\theta}, \overline{\theta}\right]$
u_t^{ep}	Ex-post utility for Customer i
k	Granularity of the rating system. $K = 2k$ is the size of the rating system
r_t	Rating $r \in R = \{-k, ..., -1, 1, ..., k\}$ at period t
α	Captures the weight of the customer disconfirmation bias
u^r	Customer reference utility, which is a function of α, θ, and Q^c
λ_i	The threshold for rating i
π_t	The distribution of previous ratings at period t, and $\pi_t = \left(\pi_{t(-k)}, ..., \pi_{tk}\right)$, where π_{tj} is the frequency of Rating j
$\tilde{Q}^p$	The platform's quality belief
κ	The probability of Customer t rating the product at r, which is a function of r, Q, and $\tilde{Q}_t$
$\bar{r}_t$	$\bar{r}_t$ denote the average ratings until period t
$\hat{Q}_t^c$	Customer posterior belief at period t under average rating, and $\hat{Q}_t^u$ is the unbiased customer posterior belief
$V^L(Q)$	Speed of learning when customer heterogeneity is low
$V^H(Q)$	Speed of learning when customer heterogeneity is high

3.4 Complete Rating System

In a complete rating system, customers have access to the entire rating history and are capable of extracting information about product quality in a Bayesian way. In such a model, $\mathcal{I}_t^c = \{(r_1, r_2, ..., r_{t-1})\}$, where r_j is Customer j's rating for $j < t$. We refer to this regime as the *Complete (Rating) System*.

Before investigating the evolution and convergence of the ratings, we illustrate the behavior of this rating system and the implications of the disconfirmation bias using a series of examples.

Example 1: Consider a 2-scale rating system, i.e., $k = 1$. For example, this can represent a thumbs-up/down system. Suppose $Q = 2/10$, $\tilde{Q}$ is uniformly distributed on $[0, 1]$, and $\overline{\theta} = 0$, i.e., customers are ex-post homogeneous. Finally, let $\alpha = 1$ and $p = 0$. Hence, customers do not reflect on their post-purchase experiences; they purely rate based on their disconfirmation biases. In this example, the first customer forms quality belief $Q_1^c = 1/2$. This customer receives ex-ante utility $u_1^{ea} = 0.5 > 0$ and purchases the product. She then receives reference utility $u_1^r = 2/10 - 1/2 = -3/10 < 0$, and therefore, negatively rates her experience, i.e., $r_1 = -1$. Observing this rating, the second customer updates her belief about the product quality. A Bayesian customer would infer that $\tilde{Q}_2$ is uniformly distributed on $[0, 0.5]$ and $Q_2^c = 0.25$. Thus, this customer purchases and rates the product at $r_2 = -1$. Figure 3.1 plots the evolution of customer beliefs over time. We observe that the customer beliefs eventually converge to the intrinsic quality as more customers rate the product. In this example, $k = 1$, $\alpha = 1$, $Q = 2/10$, $\tilde{Q} \sim U(0, 1)$, $\overline{\theta} = 0$, and $p = 0$.

Example 2: Consider the same parameters as in Example 1, except that $\alpha = 1/3$. Hence, customers partially exhibit disconfirmation bias. Since the first customer does not observe any ratings, her prior expected belief is again $Q_1^c = 1/2$. This customer receives ex-ante utility $u_1^{ea} = 0.5 > 0$ and purchases the product. She receives reference utility $u_1^r = (1 - 1/3) \cdot 2/10 + 1/3 \cdot (2/10 - 1/2) = 1/30 > 0$ and rates positively, i.e., $r_1 = 1$. Observing this rating, the second customer updates her belief about the product quality. A Bayesian customer would infer that $\tilde{Q}$ is uniformly distributed on $[2/10, 1]$ and $Q_2^c = 6/10$. This customer purchases the product and positively rates it, i.e., $r_2 = 1$. It is straightforward

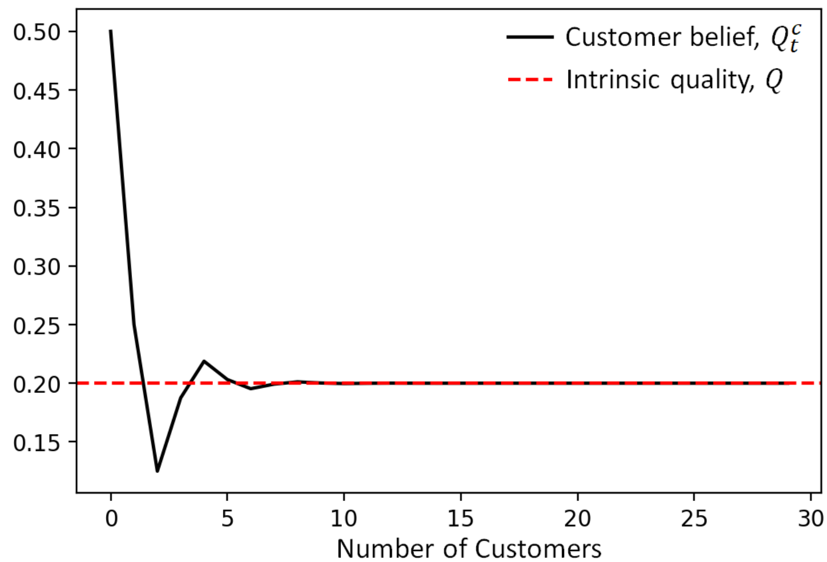

Figure 3.1: Evolution of customer quality beliefs over time in a complete system.

to see that subsequent customers would also buy and positively rate the product. Hence, $Q_2^c = Q_3^c = \cdots = Q_\infty^c = 0.6$, which is different from the product's intrinsic quality.

Figure 3.2 plots Q_∞^c for various values of k. Interestingly, we observe that Q_∞^c non-monotonically converges to the intrinsic quality as the rating system becomes more granular, i.e., the number of rating options increases. This observation motivates us to study the interplay between the disconfirmation bias and rating system's granularity and its implications for customer beliefs. In this example, $\alpha = 1$, $Q = 2/10$, $\tilde{Q} \sim U(0, 1)$, $\overline{\theta} = 0$, and $p = 0$.

Next, we formally study the convergence of the customer beliefs. We start with the special case when $\alpha = 1$, i.e., the case with maximum disconfirmation bias.

Proposition 3.1. *When $\alpha = 1$, there exists price threshold $\overline{p} > 0$ such that*
(i) If $p \leq \overline{p}$, we have $\lim_{t\to\infty} Q_t^c = Q$.
(ii) If $p > \overline{p}$, threshold $\bar{t}$ exists such that for $t > \bar{t}$, customers do not purchase the product and learning stops.

Part (i) establishes that when price is sufficiently low that customers continue purchasing and rating the product, customer beliefs converge to the intrinsic quality. This is because

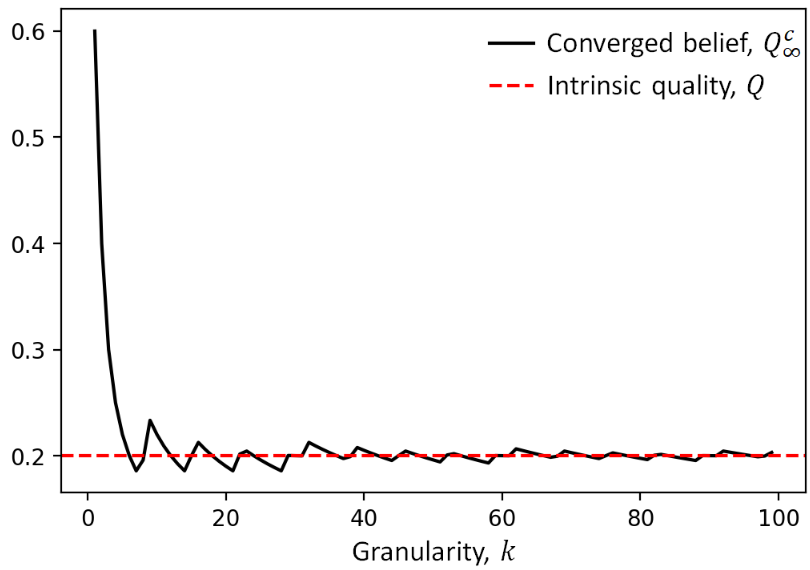

Figure 3.2: Converged customer quality beliefs as a function of the rating system's granularity in a complete system.

when Customer t rates the product at $r_t > 0$, the lower bound of the support of the next customer's quality belief, i.e., $\tilde{Q}_{t+1}$, increases. Also, a rating $r_t < 0$ would reduce the upper bound of the customer belief. This results in a sequence of contracting belief supports that eventually converge to the intrinsic quality as long as customers continue purchasing and rating the product. This convergence result does not require customers to have unbiased priors. As is typical with Bayesian updating, the prior bias eventually corrects itself as long as learning continues.

However, similar to the conventional observation in the customer ratings' literature, the product's price should be sufficiently low such that customers continue purchasing. If the price is set too high, early customers may form a belief that yields a negative utility, motivating them to stop purchasing. This would then stop ratings from being updated. Notably, this price threshold may be lower than Q. In other words, customers would have purchased the product if they knew the intrinsic quality. However, the uncertainty in the product value would discourage them from buying.

Remark 1. Throughout the paper, we find similar results for scenarios where a high price interrupts customer purchases and learning. To avoid repetition, for the remainder of the paper, we assume the price is sufficiently low such that customers continue purchasing. In each of our information schemes, we also considered a variation where the firm optimally sets the price. We show that a firm with sufficiently a high discount factor for future revenue would never set a price so high that customers would stop purchasing.

Next, we consider the model with $\alpha < 1$. In this scenario, in addition to the disconfirmation bias, customers incorporate their ex-post utility into the ratings. We establish the following result.

Proposition 3.2. *There exists threshold* $\overline{\alpha}(k, \overline{\theta})$ *such that*
(i) If $\alpha \leq \overline{\alpha}(k, \overline{\theta})$*, then* $\lim_{t\to\infty} Q_t^c = Q$.
(ii) If $\alpha > \overline{\alpha}(k, \overline{\theta})$*, there exists intrinsic quality* $Q \in \mathcal{Q}$ *such that* $\lim_{t\to\infty} Q_t^c = Q_\infty(\alpha) \neq Q$. *However,* $|Q_\infty^c - Q| \leq 2(1-\alpha)\overline{\theta}$. *As such,* $\lim_{\alpha\to 1} Q_\infty(\alpha) = Q$.
(iii) Threshold $\overline{\alpha}(k, \overline{\theta})$ *is increasing in* $\overline{\theta}$. *Furthermore, it is increasing in* k *for* $k \geq 2$*, and* $\lim_{k\to\infty} \overline{\alpha}(k, \overline{\theta}) = 1$ *and* $\lim_{\overline{\theta}\to\infty} \overline{\alpha}(k, \overline{\theta}) = 1$.

Part (i) shows that when customers use the entire rating history, their beliefs converge to the intrinsic quality as long as the disconfirmation bias is not strong, i.e., when $\alpha \leq \overline{\alpha}(k, \overline{\theta})$. This result is not a priori trivial: When α is positive, customers report their ex-post utility with a "bias." This bias distorts the ratings for future customers, which can eventually result in the non-convergence or convergence to a biased value. However, in a complete system, Bayesian customers can correctly account for this bias when α is small. In this setting, customers can rationally hypothesize the thought process of previous customers from the rating history and eventually distinguish the effect of disconfirmation bias from customers' idiosyncratic heterogeneity.

As indicated in Part (iii), $\overline{\alpha}(k, \overline{\theta})$ is increasing in $\overline{\theta}$. Therefore, as customer heterogeneity increases, the convergence to the intrinsic quality is guaranteed for a larger range of the disconfirmation bias. In the extreme case where customers are infinitely heterogeneous ($\overline{\theta} \to \infty$), the beliefs converge to the intrinsic quality for all values of α. This is somewhat counterintuitive because a higher heterogeneity adds to the ambiguity in ratings, particularly in the presence of the disconfirmation bias: A customer may rate a product positively (negatively) because she receives a high (low) ex-post utility or had a low (high) expectation which resulted in a high positive (negative) disconfirmation. However, high heterogeneity ensures that extreme ratings will eventually arise that enable customers to refine their beliefs. This result resembles the findings of Ifrach et al. (2019) and Acemoglu et al. (2022), where the convergence to the intrinsic quality is guaranteed if the customer heterogeneity possesses sufficient "Richness."

Part (ii) illustrates that, in any rating system with a given granularity, the convergence to the intrinsic quality is not guaranteed if the disconfirmation bias is significant. This is because the relative heterogeneity in customers' reference utility is governed by $(1-\alpha)\theta$. Thus, an increase in α reduces the relative richness of heterogeneity and, consequently, the possibility of observing sufficiently low and high ratings. This can result in convergence to a belief that deviates from the intrinsic quality. However, as the proposition shows, the converged belief is always within the $2(1-\alpha)\overline{\theta}$ neighborhood of the intrinsic quality, which is decreasing in α. An important implication is that a sufficiently significant disconfirmation bias warrants convergence to an arbitrarily close neighborhood of the intrinsic quality. Customers' increased emphasis on the discrepancy between the intrinsic quality and their prior expectations reduces the weight of their private subjective preferences when rating. This enables future customers to form beliefs closer to the intrinsic quality. In the special case when $\alpha = 1$, this result replicates the finding of Proposition 3.2.

Interestingly, Part (iii) shows that $\overline{\alpha}(k, \overline{\theta})$ is increasing in k with $\lim_{k \to 1} \overline{\alpha}(k, \overline{\theta}) = 1$. As such, for any heterogeneity $\overline{\theta}$ and disconfirmation bias α, the platform can ensure correct learning by designing sufficiently a granular rating system, i.e., by setting a high value of k. A more granular system works as a magnifier for customers to account for the previous customer biases. We formalize this discussion in the following proposition.

Proposition 3.3. *In a rating system with* $k \geq 1 + \frac{1}{2\overline{\theta}(1-\alpha)}$*, customer beliefs converge to the intrinsic quality.*

We observe that the granularity required to guarantee correct learning increases in α proportional to $\frac{1}{2\overline{\theta}(1-\alpha)}$. Hence, while the platform must make the system more granular to guarantee correct convergence when α increases, the increase in granularity does not need to be significant as long as α is not close to one. For example, when $Q = 1$ and $\overline{\theta} = 1$, a change in the value of α from 0 to 0.5 does not require any change in granularity for the ratings to converge to the intrinsic quality. Additionally, a further increase to $\alpha = 0.75$ only requires changing from a 4-scale to a 6-scale system. In this case, even when $\alpha > 0.75$, the customer quality beliefs would always be within 25% of the intrinsic quality without further increasing the required granularity, as shown in Part (ii) of Proposition 3.2.

3.5 Aggregate Rating System

In this section, we consider a variation where customers only observe the frequency of past ratings to form beliefs about quality. This can occur because either the platform does not disclose individual ratings or customers solely rely on rating statistics to reduce the cognitive cost of processing the complete history. We refer to this scenario as the "Aggregate (Rating) System." In an aggregate system, Customer t observes $\pi_t = \left(\pi_{t(-k)}, ..., \pi_{tk}\right)$, where π_{tj} is the frequency of Rating j.[3] Hence, $\mathcal{I}_t^c = \{\pi_t\}$.

The following example illustrates the behavior of such a rating system.

Example 3: We consider the same parameters as in Example 1, i.e., $k = 1$, $Q = 2/10$, $\tilde{Q} \sim U[0,1]$, $\overline{\theta} = 0$, $\alpha = 1$, and $p = 0$. In this example, customers only observe the ratings' frequency. Suppose Customer 3 observes $\pi_3 = (1,1)$. If $\alpha = 0$ and customers only reflected their ex-post utilities in their ratings, observing rating distribution $\pi_3 = (1,1)$ would not be possible: In this case, all customers receive a positive utility, and the only feasible rating frequency would be $\pi_3 = (0,2)$. However, when $\alpha = 1$, $\pi_3 = (1,1)$ is feasible and Customer 3 can deduce that one of the previous customers rated the product positively and one negatively. Since this customer cannot observe the ratings' order, she must account for all possibilities that led to such a ratings' frequency. Specifically, two rating paths can lead to $\pi_t = (1,1)$: $(r_1, r_2) = (+1, -1)$ and $(r_1, r_2) = (-1, +1)$.

[3]In our model, only customers who purchase can rate. Thus, we can equivalently work with the distribution of rating options by dividing each option's frequency by the number of purchasing customers. For clarity, we focus on rating frequencies.

First, consider $(r_1, r_2) = (+1, -1)$. The first customer rates the product positively if $Q > 1/2$. Therefore, the second customer would believe the intrinsic quality is uniformly distributed on $[1/2, 1]$, and Customer 2 would rate negatively if the product quality realizes below 0.75. In this scenario, Customer 3 would form a belief that $Q \sim U[0.5, 0.75]$. This happens with ex-ante probability $\mathbb{P}(Q \in [0.5, 0.75]) = 1/4$.

Next, consider $(r_1, r_2) = (-1, 1)$. Similar to the previous scenario, Customer 3 would believe that $Q \sim U[0.25, 0.5]$, which happens with ex-ante probability $\mathbb{P}(Q \in [0.25, 0.5]) = 1/4$. Since it is equally likely to observe $(r_1, r_2) = (-1, +1)$ and $(r_1, r_2) = (+1, -1)$, we find that $\mathbb{P}((r_1, r_2) = (-1, +1) | \pi_t = (1, 1)) = \mathbb{P}((r_1, r_2) = (+1, -1) | \pi_t = (1, 1)) = 1/2$ and $Q_t^c = 1/2$. As more customers rate the product, customers must account for more possibilities. Figure 3.3 illustrates the convergence of beliefs in this example. In this example, $k = 1$, $\alpha = 1$, $Q = 2/10$, $\tilde{Q} \sim U(0, 1)$, $\overline{\theta} = 0$, and $p = 0$.

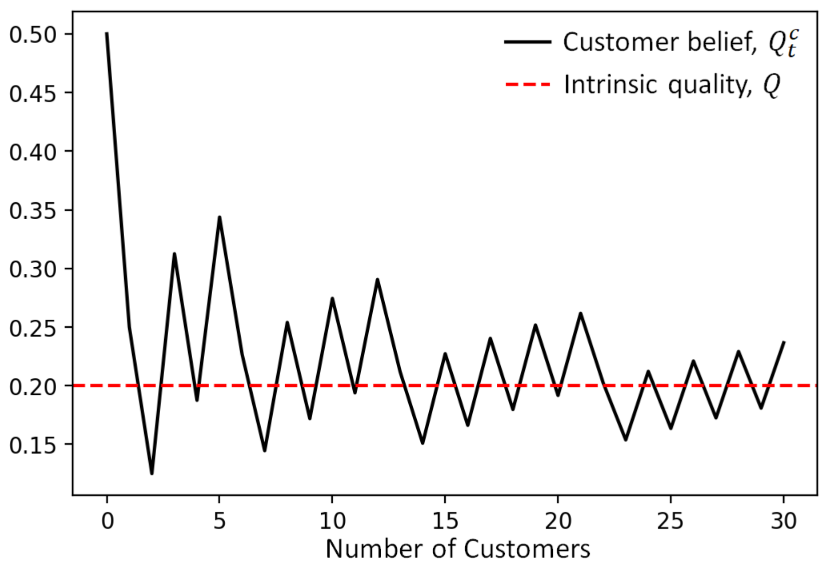

Figure 3.3: Evolution of customer quality beliefs over time in an aggregate system.

We observe that ratings oscillate below and above the intrinsic quality, and the size of these oscillations non-monotonically decreases. With this example, we next study the beliefs' convergence in an aggregate system. We first focus on the special case when $\alpha = 1$.

3.5.1 Aggregate System: $\alpha = 1$

We define the following properties to be used in the subsequent analysis.

Definition 1. *Let Supp* $\left(\tilde{Q}|\pi\right)$ *be the support of the quality belief distribution conditioned on observing rating frequency* π*, and let Int* $\left(\tilde{Q}|\pi\right)$ *be the interior of Supp* $\left(\tilde{Q}|\pi\right)$*. Also, let* $e_i = (0, ..., 1, ...0)$ *be the standard basis and define partial order* $>_o$ *such that* $\pi + e_i >_o \pi + e_j$ *for* $i \geq j$*. We define the following properties.*

(i) Feasibility: *Rating frequency* π_t *is feasible if it can be generated by a sequence of Bayesian customers with non-zero probability.*

(ii) Separation: *Two feasible rating frequencies* $\pi_t^1 \neq \pi_t^2$ *are separate if Int* $\left(\tilde{Q}|\pi_t^1\right) \cap$ *Int* $\left(\tilde{Q}|\pi_t^2\right) = \emptyset$.

(iii) Partition: *A set of rating frequencies* $\pi_t^1, \pi_t^2, ..., \pi_t^n$ *partition the quality belief space if they are separate and* $\bigcup_{i=1}^{n}$ *Supp* $\left(\tilde{Q}|\pi_t^i\right) = \mathcal{Q}$*, where* $\mathcal{Q}$ *is the support of* $\tilde{Q}$*.*

(iv) Monotonicity: *An aggregate rating system satisfies the monotonicity property if for any two feasible rating frequencies* $\pi_t^1 >_o \pi_t^2$*, we have* $\mathbb{E}\left[\tilde{Q}|\pi_t^1\right] > \mathbb{E}\left[\tilde{Q}|\pi_t^2\right]$.

This definition introduces a set of intuitive properties. A rating frequency is *feasible* if it can arise from customer rating choices. For example, in a binary rating system, one cannot have more than t positive or negative reviews at time t. Additionally, some combinations of ratings may not arise, as discussed in Example 3. The *Separation* of two rating frequencies implies that, upon observing them, rational agents can form quality beliefs with non-overlapping domains (except for the boundaries). Hence, they induce completely separate quality beliefs. A *partition* of the quality beliefs enables us to represent the belief space in terms of separate rating frequencies. Finally, *monotonicity* asserts that the expected quality belief increases upon observing a rating frequency deemed better than another, e.g., an arriving customer likes the product instead of disliking it.

While these are intuitive properties, the existence of a monotonic partition in each period is not a priori trivial. The following lemma establishes the existence of such a partition.

Lemma 3.1. *When* $\alpha = 1$*, for any* $t \geq 1$*, the set of all feasible rating frequencies partitions* $\mathcal{Q}$*. Furthermore, this partition can be uniquely ordered according to partial order* $>_o$*. In other words, this partition can be represented as tuple* $(\pi_t^1, ..., \pi_t^{n(t)})$ *such that* $\pi_t^i >_o \pi_t^{i+1}$*.*

This lemma illustrates that any two different rating frequencies that can feasibly arise from customer rating choices induce disjoint supports for customer quality beliefs except for the boundaries. Consequently, customers form distinct beliefs about the intrinsic quality upon observing different rating frequencies. This also implies that all feasible rating frequencies partition the quality belief space. In other words, upon observing a rating frequency, customers can immediately refine their beliefs to a subinterval of their prior beliefs.

In addition, Lemma 3.1 shows that the set of feasible rating frequencies can be partially ordered. This property is important because Customer t can only change the current rating frequency by e_i for some $i \in R$. Hence, each arriving customer forms a distinct belief from the previous customer. This suggests that the learning continues as long as customers continue purchasing. However, this learning is not monotonic, as observed in Figure 3.3. This result assists us in formalizing the following Lemma.

Lemma 3.2. *Given rating frequency π_t at time t, there exists real numbers $\overline{Q}_t > \underline{Q}_t$ such that*

(i) In an aggregate system, the support of the customer beliefs, i.e., Supp $\left(\tilde{Q}|\pi_t\right)$ is the interval $\left[\underline{Q}_t, \overline{Q}_t\right)$. Furthermore, this interval contains the intrinsic quality, i.e., $Q \in \left[\underline{Q}_t, \overline{Q}_t\right)$.

(ii) Let $\tilde{Q}^p$ denote the platform's quality belief. Then,

$$\mathit{Supp}\left(\tilde{Q}^p|(r_1, ..., r_t)\right) = \left[\max_{\tau<t}\{\underline{Q}_\tau\}, \min_{\tau<t}\{\overline{Q}_\tau\}\right) \subset \left[\underline{Q}_t, \overline{Q}_t\right).$$

Furthermore, $Q \in \left[\max_{\tau<t}\{\underline{Q}_\tau\}, \min_{\tau<t}\{\overline{Q}_\tau\}\right)$.

Part (i) establishes that customer quality beliefs always belong to an interval that contains the intrinsic quality. In other words, upon observing the ratings, customers rule out the quality realizations above and below certain thresholds. While the customer quality belief intervals evolve and include the intrinsic quality, they are not necessarily unbiased and may not contract. So, the convergence of the ratings is not apparent from this result.

In an aggregate system, there is information asymmetry between the platform and customers: while customers do not observe the sequence of ratings, the platform has access to the complete rating history. Part (ii) shows that the platform's quality beliefs are more refined than those inferred by customers. In other words, the support of the platform's belief is a subset of the customers' belief support. Furthermore, unlike customers' beliefs, the platform's belief constitutes a sequence of contracting intervals over time. Consequently, if customer quality beliefs converge under the aggregate system, the platform's beliefs would also converge, potentially at a faster speed. This creates the opportunity to influence the evolution of ratings through various mechanisms, such as information provisioning. Although we do not explore such dynamics in this paper, they could be an interesting avenue for future research.

Figure 3.4 plots the dynamics of the quality beliefs and intervals discussed in Lemma 3.2. In this example, $k = 1$, $\alpha = 1$, $Q = 2/10$, $\tilde{Q} \sim U[0, 1]$, $\overline{\theta} = 0$, and $p = 0$.

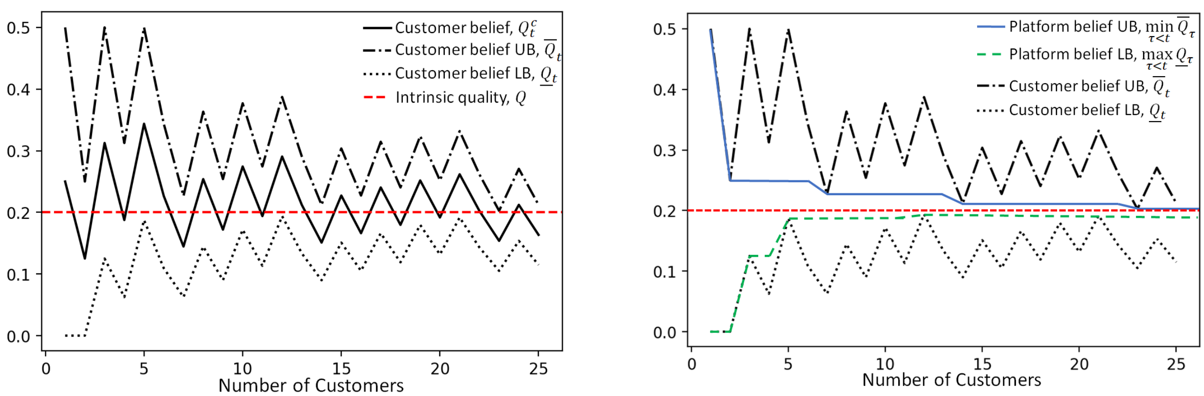

(a) The expected value and support of customer beliefs.

(b) Comparison between the support of customer and platform beliefs.

Figure 3.4: Evolution of customer and platform quality beliefs over time in an aggregate system. UB: Upper Bound; LB: Lower Bound.

Next proposition shows the convergence of the quality beliefs when $\tilde{Q}$ is uniformly distributed and $\alpha = 1$. We will return to the case with a general distribution for $\tilde{Q}$ in §3.5.2.

Proposition 3.4. *Suppose* $\tilde{Q} \sim U[0, 1]$ *and* $\alpha = 1$. *Then,* $\lim_{t\to\infty} Q_t^c = Q$ *in an aggregate system.*

This result illustrates that quality beliefs eventually converge to the intrinsic quality in the presence of significant disconfirmation bias even when customers do not observe the complete rating history. Notably, this finding does not require customers to have unbiased prior beliefs. Prior beliefs play an important role particularly when customers solely rate based on the deviation between the realized quality and their prior perception. If their prior perception is incorrect, their ratings would also be biased. However, Proposition 3.4 shows that these biases correct themselves as more customers rate the product. This is because the disconfirmation effect serves as an instrument to create heterogeneity among customers based on their arrival times. This heterogeneity leads to continued learning until quality beliefs eventually converge to the intrinsic quality. As Figure 3.4a shows, customer biases may not contract monotonically when customer information is incomplete.

3.5.2 Aggregate System: $\alpha < 1$

In this section, we consider $\alpha \in [0, 1)$. When $\alpha = 1$, customers solely rate using $Q - Q_t^c$. Thus, the ex-post heterogeneity (θ) does not weigh in the ratings. In this situation, as shown in Lemma 3.1, different rating frequencies induce disjoint supports for quality beliefs. However, when $\alpha < 1$, customers not only exhibit a disconfirmation bias, but they also reflect on their heterogeneous ex-post utilities. This creates further uncertainty, leading to overlapping supports for quality beliefs conditioned on different rating frequencies. Notably, the separation property defined in Definition 1 would no longer hold in the presence of ex-post heterogeneity. However, this does not imply that customers cannot refine their beliefs upon observing new ratings: As long as customers can form sufficiently differentiated quality distributions from different ratings, they may eventually uncover the intrinsic quality. As such, we require a weaker notion of separation.

To facilitate discussions, let $\kappa\left(r; Q, \tilde{Q}_t\right)$ denote the probability of Customer t rating the product at r given intrinsic quality Q and belief $\tilde{Q}_t$. Also, define $\mathcal{F}_t$ to be the space of all feasible belief distributions. Motivated by the definition used in Acemoglu et al. (2022), we define the following, which we refer to as *separation in distribution.*

Definition 2. *(Separation in Distribution) For arbitrary quality values $Q_1 > Q_2$, a rating system satisfies weak separation in distribution if there exists a subset of rating options $S \subseteq \mathbf{R}$, independent of Q_1 and Q_2, such that*

$$\inf_{\tilde{q}_t \in \mathcal{F}_t} \sum_{i \in S} \kappa(i; Q_1, \tilde{q}_t) \geq \sup_{\tilde{q}_t \in \mathcal{F}_t} \sum_{i \in S} \kappa(i; Q_2, \tilde{q}_t), \tag{11}$$

or

$$\inf_{\tilde{q}_t \in \mathcal{F}_t} \sum_{i \in S} \kappa(i; Q_2, \tilde{q}_t) \geq \sup_{\tilde{q}_t \in \mathcal{F}_t} \sum_{i \in S} \kappa(i; Q_1, \tilde{q}_t). \tag{12}$$

A rating system satisfies the strict separation condition if the above inequalities are strict.

Separation in distribution ensures that for any two intrinsic quality values, at least one rating option exists that, if selected by the customers, can induce beliefs with distinct supports. The notion of separation defined in Definition 2 is stronger than the one used in Acemoglu et al. (2022). In their paper, the distribution of the quality beliefs is binary. Thus, the set of rating options S used in the definition is fixed. However, since we consider a general distribution for the intrinsic quality, Definition 2 requires that this set be the same for any two quality values. Additionally, we need to have Infimum and Supremum on the space of all feasible probability distributions which is not necessary when the intrinsic quality is supported on a binary space. The following Lemma finds a sufficient condition for the strict separation in distribution in an aggregate system.

Lemma 3.3. *In an aggregate system, there exists threshold $\overline{\alpha}_I(k, \overline{\theta})$ such that if $\alpha \leq \overline{\alpha}_I(k, \overline{\theta})$, then any two distinct intrinsic quality values can be separated in distribution.*

In contrast to the complete system, different rating frequencies do not necessarily induce completely disjoint quality beliefs. However, when α is small, "some" rating options would enable future customers to learn about the intrinsic quality. In other words, for any two quality values, customers can identify with high confidence which one is more likely if they observe a separating set of ratings. Therefore, convergence to the intrinsic quality depends on whether the set of separating ratings would eventually arise. The following proposition addresses this question.

Proposition 3.5. *There exist thresholds $\overline{\alpha}(k, \overline{\theta})$ and $\overline{\alpha}_I(k, \overline{\theta})$ such that*
(i) If $\alpha < \overline{\alpha}_I(k, \overline{\theta})$, then $\lim_{t\to\infty} Q_t^c = Q$.
(ii) If $\overline{\alpha}_I(k, \overline{\theta}) \leq \alpha \leq \overline{\alpha}(k, \overline{\theta})$ and $\lim_{t\to\infty} Q_t^c = Q_\infty^c(\alpha)$ exists, then $Q_\infty^c(\alpha) = Q$.
(iii) If $\alpha > \overline{\alpha}(k, \overline{\theta})$, there exists intrinsic quality Q such that either $Q_\infty^c(\alpha)$ does not exist or $Q_\infty^c(\alpha) \neq Q$.
(iv) Threshold $\overline{\alpha}_I(k, \theta)$ is non-decreasing in $\overline{\theta}$. Furthermore, it is non-decreasing in k for $k \geq 2$, and $\lim_{k\to\infty} \overline{\alpha}_I(k, \overline{\theta}) = 1/2$ and $\lim_{\overline{\theta}\to\infty} \overline{\alpha}_I(k, \overline{\theta}) = 1/2$.

Part (i) shows that when disconfirmation bias is small ($\alpha < \overline{\alpha}_I(k, \overline{\theta})$), customer beliefs converge to the intrinsic quality in an aggregate system similar to the complete system. In this case, ratings would eventually arise that allow future customers to refine their quality beliefs. This refinement continues until customer beliefs converge to the intrinsic quality. Furthermore, when $\alpha \leq \overline{\alpha}(k, \overline{\theta})$, correct learning is guaranteed as long as $Q^c_\infty(\alpha)$ exists. The proof proceeds by showing that if beliefs converge, they would be "strictly" increasing in the intrinsic quality. In other words, a higher quality product would induce higher quality beliefs. Thus, Bayesian customers can infer the intrinsic quality.

For large disconfirmation bias, i.e., $1 > \alpha > \overline{\alpha}(k, \overline{\theta})$, either the beliefs do not converge, or they would converge to a biased value undetectable by customers. Similar to the complete system, an increase in α reduces the richness in customer ex-post heterogeneity, which prevents sufficiently low or high ratings to arise. Hence, customers may not be able to continue refining their beliefs.

Part (iv) implies that the platform can design a sufficiently granular rating system to guarantee convergence in an aggregate system as long as $\alpha \leq 1/2$. Moreover, we note that threshold $\overline{\alpha}(k, \overline{\theta})$ in Part (ii) is the same as the one found in Proposition 3.2 and $\lim_{k\to\infty} \overline{\alpha}(k, \overline{\theta}) = 1$. As such, if the beliefs converge, the platform can guarantee correct learning by offering a sufficiently large number of rating options to the customers for all values of α.

In sum, all the results in a complete system extend to an aggregate system, with the caveat that the convergence of the beliefs is not guaranteed for large values of the disconfirmation bias.

3.6 Average Ratings

As discussed earlier, although Bayesian customers may infer the intrinsic quality from the complete and aggregate systems, it is cognitively expensive. In a complete system, this inference requires a large memory of tracking the entire rating history. In an aggregate system, it involves combinatorially increasing number of scenario analyses. Therefore, customers may naturally use simpler summary statistics, such as the average of past ratings.

Let $\bar{r}_t$ denote the average ratings until period t, i.e., $\bar{r}_t = \left(\sum_{\tau=1}^{t} r_\tau\right)/t$. To focus on the asymptotic behavior of the beliefs, we follow a similar approach as Shin et al. (2023) and consider the following customer posterior belief.

$$\hat{Q}^c_{t+1} \equiv \frac{1}{\gamma t+1}\mathbb{E}(\tilde{Q}) + \frac{\gamma t}{\gamma t+1}\bar{r}_t, \tag{13}$$

for $\gamma > 0$ and expected prior quality belief $\mathbb{E}(\tilde{Q})$. Intuitively, customers initially give significant weight to their prior beliefs when the product has few ratings. As more customers rate the product, the weight of the ratings increases in customer beliefs. Parameter γ controls the decay in customer weights for their prior beliefs. A higher γ indicates a higher sensitivity to average ratings when forming beliefs about the intrinsic quality. Nevertheless, it is straightforward to see that if $\bar{r}_\infty$ exists, then $\hat{Q}^c_\infty = \bar{r}_\infty$ is independent of γ. In this section, we continue to use mapping (10) for customer ratings. However, customers use (13) to form beliefs about the product quality when calculating their reference utility.

Let $\overline{\alpha}(k,\overline{\theta})$ be the threshold defined in Propositions 3.2 and 3.5. The following result establishes the asymptotic behavior of the customer quality beliefs for a general α.

Proposition 3.6. $\lim_{t\to\infty}\hat{Q}^c_t = \hat{Q}^c_\infty(Q,\alpha)$ *exists. Furthermore,*
(i) If $\alpha \leq \overline{\alpha}(k,\overline{\theta})$*, then* $\hat{Q}^c_\infty(Q,\alpha)$ *is strictly increasing in* Q*.*
(ii) If $\alpha \in \left(\overline{\alpha}(k,\overline{\theta}),1\right)$*, then* $\hat{Q}^c_\infty(Q,\alpha)$ *is weakly increasing in* Q*.*
(iii) If $\alpha = 1$*, then* $\hat{Q}^c_\infty(Q,\alpha) = Q$*.*

This Lemma shows that customer quality beliefs converge when customers use the average ratings. However, as Part (ii) illustrates, the converged belief is not "strictly" monotonic in the intrinsic quality when the disconfirmation bias is large. While the weak monotonicity property implies that a higher-quality product would never receive a lower average rating, it does not guarantee that customers would be able to distinguish two different quality products. In other words, a product with a higher quality may receive the same rating as an inferior product when the disconfirmation bias is large. This finding is illustrated in Figure 3.5a. In this example, all products with an intrinsic quality higher than 0.82 receive the same average rating, making them indistinguishable.

Part (ii) holds only when $\alpha \neq 1$. When $\alpha = 1$, the customer beliefs converge to the intrinsic quality even when customers use average ratings. In this case, customers are so biased toward the deviations from their prior expectations that the ex-post utility does not play a role in their ratings. In such a scenario, customers can attribute the variations in ratings only to the disconfirmation bias, and a simple summary statistic such as average ratings enables customers to eventually learn the intrinsic quality.

When the disconfirmation bias is small ($\alpha \leq \overline{\alpha}$), Part (i) demonstrates a one-to-one mapping between the intrinsic quality and average ratings. Specifically, higher-quality products would receive higher average ratings. This enables customers to correctly "rank" the quality of the products. However, as illustrated in Figure 3.5a, the converged beliefs do not precisely match the intrinsic quality. In this example, $\tilde{Q} \sim U(0, 1)$, $\tilde{\theta} \sim U(-0.4, 0.4)$, and $p = 0$. In particular, we observe a systematic error between customer beliefs and the intrinsic quality. Although such an error cannot exist with Bayesian customers as they incorporate it into their belief updating, in this section, customers follow a simple heuristic, making them prone to these errors.

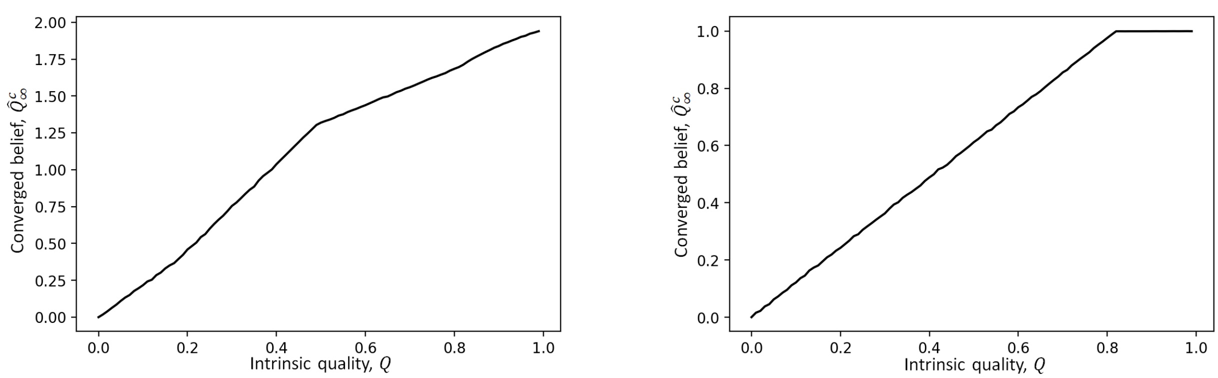

(a) $\alpha = 0.1$

(b) $\alpha = 0.8$

Figure 3.5: Converged customer quality beliefs as a function of the intrinsic quality in a system with average ratings.

Next, we explore a bias-correcting rule for customer beliefs. Since $\hat{Q}^c_\infty(Q, \alpha)$ is strictly increasing in Q when α is small, function $\phi_\alpha(\cdot)$ exists such that $\hat{Q}^c_\infty(Q, \alpha) = \overline{r}_\infty = \phi_\alpha(Q)$.

Define the following belief updating rule when $\alpha \leq \overline{\alpha}(k, \overline{\theta})$:

$$\hat{Q}^u_{t+1} = \frac{1}{\gamma t + 1}\mathbb{E}(\tilde{Q}) + \frac{\gamma t}{\gamma t + 1} \cdot \phi_\alpha^{-1}(\bar{r}_t). \tag{14}$$

The following corollary formally establishes the convergence of the above rule.

Corollary 2. *When customers update their beliefs according to* (14) *and* $\alpha \leq \overline{\alpha}$, $\lim_{t\to\infty} \hat{Q}^c_t = \hat{Q}^u_\infty(Q, \alpha)$ *exists. Furthermore,* $\hat{Q}^u_\infty(Q, \alpha) = Q$.

When the disconfirmation bias is small, the bias correcting rule (14) guarantees convergence to the intrinsic quality.

3.7 Learning Speed

As discussed in the previous sections, the rating system's granularity, customers' disconfirmation bias, and ex-post heterogeneity play critical roles in the customer belief evolution. Particularly, convergence to the intrinsic quality is guaranteed as long as the disconfirmation bias is small. In this section, we study the learning speed in the aggregate system. To comprehensively analyze the learning speed, we separately consider the cases of low and high heterogeneity in §3.7.1 and §3.7.2.

To further refine our insights, for the remainder of the paper, we assume the prior quality beliefs follow a discrete distribution with support $\mathcal{Q} = \{Q_1, ..., Q_M\}$ for $M \geq 2$.

3.7.1 Learning Speed: Low Customer Heterogeneity

In this section, we analyze the case where $\overline{\theta} < \frac{1}{1-\alpha} + Q - p$. Recall that $\kappa(r; Q, \tilde{Q}_t)$ denotes the probability of Customer t rating the product at r given intrinsic quality Q and customer belief $\tilde{Q}_t$. We define the following.

Definition 3. *(Separation Divergence) Given customer belief* $\tilde{Q}_t$*, two arbitrary quality values* Q_1 *and* Q_2 *are separation divergent if there exists at least one rating* $r(Q_1, Q_2)$ *such that*

$$\begin{aligned} &\kappa\left(r(Q_1, Q_2); Q_1, \tilde{Q}_t\right) \cdot \kappa\left(r(Q_1, Q_2); Q_2, \tilde{Q}_t\right) = 0, \\ \textit{and } &\kappa\left(r(Q_1, Q_2); Q_1, \tilde{Q}_t\right) + \kappa\left(r(Q_1, Q_2); Q_2, \tilde{Q}_t\right) > 0. \end{aligned} \tag{15}$$

When heterogeneity is not sufficiently rich, some ratings may never arise given the intrinsic quality. This definition imposes that certain ratings exist such that they can only arise for some intrinsic quality while they would never arise for others. In this circumstance, we define the learning speed as the expected number of ratings for separation divergence between all $Q_i \in \mathcal{Q} \setminus Q$ and the intrinsic quality.

Let $\tau(Q_1, Q_2)$ be the minimum number of customers until rating $r(Q_1, Q_2)$ arises to separate Q_1 and Q_2. In other words,

$$\tau(Q_1, Q_2) = \min\{t : \kappa\left(r(Q_1, Q_2); Q_1, \tilde{Q}_t\right) = 1\}. \tag{16}$$

Then, for the intrinsic quality Q, the learning speed is given by

$$V^L(Q) \equiv \frac{1}{\max\limits_{Q_i \in \mathcal{Q} \setminus \{Q\}} \mathbb{E}[\tau(Q, Q_i)]}. \tag{17}$$

Thus, $V^L(Q)$ controls the rate at which beliefs converge to the intrinsic quality. With these preliminaries, we characterize the learning speed when heterogeneity is low.

Proposition 3.7. *Let Q be the intrinsic quality and $\overline{\theta} \leq \frac{1}{1-\alpha} + Q - p$. Conditional on correct learning,*
(i) there exists threshold $\bar{k}$ such that Q is separation divergent from any other quality value for $k \geq \bar{k}$.
(ii) Suppose $k \geq \bar{k}$. For $\epsilon > 0$, let $\underline{Q}_\epsilon \equiv \max\{Q - \epsilon, 0\}$ and $\overline{Q}_\epsilon \equiv \min\{Q + \epsilon, 1\}$. Then, for arbitrary $Q' \in \mathcal{Q}$, we have

$$\mathbb{E}[\tau(Q, Q')] \leq \frac{1}{\bar{F}_\theta\left(\overline{\theta} - \frac{|Q-Q'| - \alpha\left(\overline{Q}_\epsilon - \underline{Q}_\epsilon\right) - \frac{1}{k-1}}{1-\alpha}\right)}. \tag{18}$$

(iii) Furthermore, let $u \equiv \min\limits_{Q_i \in \mathcal{Q} \setminus \{Q\}} |Q - Q_i|$. For all Q,

$$V^L(Q) \geq \bar{F}_\theta\left(\overline{\theta} - \frac{u - \alpha\left(\overline{Q}_\epsilon - \underline{Q}_\epsilon\right) - \frac{1}{k-1}}{1-\alpha}\right). \tag{19}$$

Part (i) echoes our earlier finding that a sufficiently granular rating system guarantees convergence to the intrinsic quality. Parts (ii) and (iii) further illustrate that the learning speed depends on the distribution of customer heterogeneity, disconfirmation bias, and the rating system's granularity. Specifically, we observe that as customer heterogeneity stochastically increases, the learning speed increases. This is because an increase in the likelihood of observing extreme ratings accelerates the rise of separation divergent ratings, enabling customers to distinguish different quality values.

Additionally, when customer beliefs are close to the intrinsic quality (ϵ is small), an increase in the disconfirmation bias (increase in α) can speed up learning. In this case, disconfirmation bias serves as intertemporal ex-ante heterogeneity, which increases the likelihood of ratings that induce separation divergence. Interestingly, however, disconfirmation bias slows down convergence when customer beliefs are far from the intrinsic quality. In this situation, customer biases, i.e., $Q - Q^c$, are substantial, and an increase in α significantly skews the reference utility, which in turn slows down learning.

Finally, an increase in the rating system's granularity increases the learning speed. A more refined rating system allows customers to reflect on their experiences more accurately, making separating ratings more likely to arise.

Next, we study the learning speed when customer heterogeneity is large.

3.7.2 Learning Speed: High Customer Heterogeneity

In this section, we analyze the case where $\overline{\theta} > \frac{1}{1-\alpha} + Q - p$. In this environment, the existence and rise of a rating that induces separation divergence discussed in the previous section is not guaranteed. Consequently, we use the Kullback-Leibler (KL) divergence to measure the distance between two rating frequencies arising from different intrinsic quality values. Formally, for two discrete probability distributions $\mu = (\mu_1, ..., \mu_m)$ and $\upsilon = (\upsilon_1, ..., \upsilon_m)$ with $\mu_i, \upsilon_i > 0$, KL divergence is defined as

$$D(\mu||\upsilon) \equiv \sum_{i=1}^{m} \mu_i \log(\frac{\mu_i}{\upsilon_i}). \tag{20}$$

Intuitively, KL divergence measures the average difference between the information derived from two different probability distributions. Let $\vec{\kappa}(Q, \tilde{Q}_t)$ be the vector of the probabilities of Customer t choosing each rating option. In other words,

$$\vec{\kappa}(Q, \tilde{Q}_t) = \left(\kappa(r; Q, \tilde{Q}_t) : r \in R\right). \tag{21}$$

These definitions equip us with the tools to measure the expected difference between the probability of customer actions (ratings) for different intrinsic quality values and customer beliefs. For example, $D\left(\vec{\kappa}(Q = Q_1, \tilde{Q}_t = Q_1)||\vec{\kappa}(Q = Q_1, \tilde{Q}_t = Q_2)\right)$ at time t and intrinsic quality Q_1 measures the average information difference between the rating probability distributions that arise when customers believe the intrinsic quality is Q_1 compared to when they believe it is Q_2. The larger the distance, the more quickly the two rating options can be separated.

With this background, we define the learning speed as

$$V^H(Q) \equiv \lim_{\tau \to \infty} \frac{\log\left(1 - q_\tau(Q)\right)}{\tau}, \tag{22}$$

where $q_\tau(Q) \equiv \mathbb{P}\left(\tilde{Q}_\tau = Q\right)$. Hence, if this limit exists, the convergence of the beliefs to the intrinsic quality is exponentially fast. With these preliminaries, we state the following proposition.

Proposition 3.8. *Let Q be intrinsic quality and $\overline{\theta} > \frac{1}{1-\alpha} + Q - p$. Conditional on correct learning, the learning speed is exponentially fast. Specifically,*

$$-\max_{\underline{Q} \in \mathcal{Q} \setminus \{Q\}} D(\vec{\kappa}(Q, q_\infty(Q) = 1)||\vec{\kappa}(\underline{Q}, q_\infty(\underline{Q}) = 1)) \leq V^H(Q) \tag{23}$$

and

$$V^H(Q) \leq -\min_{\overline{Q} \in \mathcal{Q} \setminus \{Q\}} D(\vec{\kappa}(Q, q_\infty(Q) = 1)||\vec{\kappa}(\overline{Q}, q_\infty(\overline{Q}) = 1)) \tag{24}$$

Convergence to the intrinsic quality is exponentially fast for any disconfirmation bias as long as it does not prohibit correct learning. This proposition also provides a lower bound on the learning speed based on the KL divergence between the intrinsic quality and one of the possible quality realizations. Intuitively, we can break down the problem of identifying the underlying quality into multiple binary hypothesis tests where customers proceed by separating all possible quality dyads. However, as the proposition shows, the time required to separate all these dyadic problems is not additive. Particularly, these dyadic separation problems would eventually nest such that customers can distinguish all quality values by separating the intrinsic quality from another quality value that is the hardest to determine. However, one should note that, as discussed in §3.5, the support of customer beliefs does not monotonically contract over time. Hence, it is not necessarily easiest/hardest to separate a quality value that is farthest from/closest to the intrinsic quality.

Proposition 3.8 also provides an upper bound for the learning speed. In this case, the learning speed is bounded above by the KL divergence between the intrinsic quality and the fastest quality value to separate from. Intuitively, distinguishing the intrinsic quality involves separating "all" possible quality values, which cannot be faster than the KL divergence between the intrinsic quality and any other values. In the special case where $M = 2$, i.e., the intrinsic quality can belong to only two possible values, the bounds of Proposition 3.8 sandwich the learning speed, and we have

$$V^H(Q) = -D(\vec{\kappa}(Q, q_\infty(Q) = 1)||\vec{\kappa}(\mathcal{Q} \setminus Q, q_\infty(\mathcal{Q} \setminus Q) = 1)). \tag{25}$$

This result replicates the learning speed in Acemoglu et al. (2011), where the quality distribution is binary. While their model incorporates the selection bias, it does not explicitly include the disconfirmation bias. We find a similar learning speed because the problem of distinguishing quality involves validating or rejecting the hypotheses that the intrinsic quality is low or high. In the case of the independently drawn samples, the exponential decay in the error term is governed by KL divergence between the sample probability distributions conditioned on each hypothesis (Cover and Thomas 2012). However, in our papers, the samples are not independently drawn due to different biases. Yet, we both show that the effect of these different sources of bias can be bounded under correct learning, and the convergence rate is governed by the KL divergence between two relevant distributions.

While the learning speed is governed by Equation (25), this equation depends on function $\vec{\kappa}(\cdot,\cdot)$, which changes based on the behavior of the reference utility. In this example, $f_\theta(x) = -x/\overline{\theta}^2$, if $x \in [-\overline{\theta}, 0)$, and $f_\theta(x) = x/\overline{\theta}^2$, if $x \in [0, \overline{\theta}]$. Thus, the disconfirmation bias affects the rate at which beliefs converge to the intrinsic quality, as illustrated in Figure (3.6a). Overall, we observe that an increase in α increases the learning speed. The disconfirmation bias affects the learning speed in two ways: first, it adds additional noise to the ratings, distorting the perception of the intrinsic quality. This distortion slows down learning. Second, it reduces the weight of customers' subjective preferences, which accelerates learning. When customer heterogeneity is large, the latter dominates the former, and the speed of learning increases when customers exhibit more disconfirmation bias.

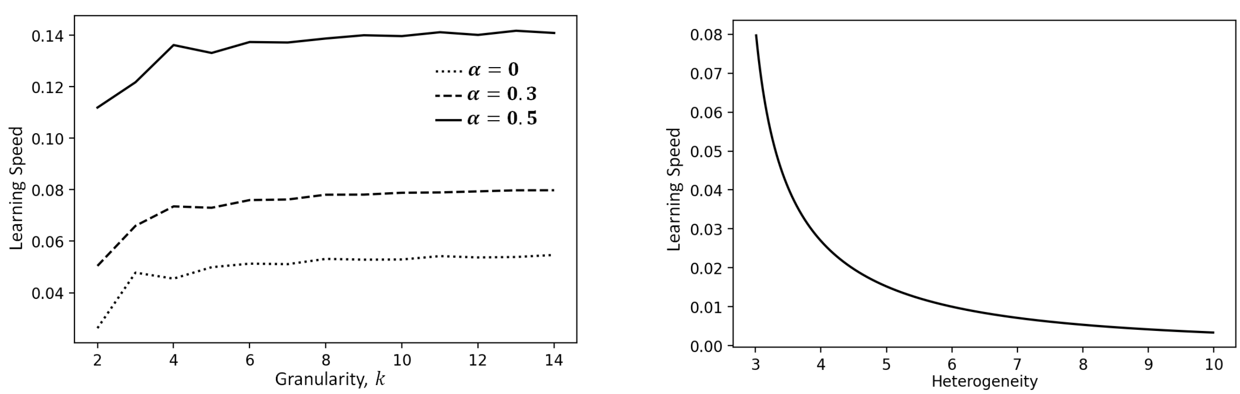

(a) $Q = Q_1 = 0.25$, $Q_2 = 0.75$, and $p = 0$.

(b) $Q = Q_1 = 0$, $Q_2 = 1$, $\alpha = 0.5$, $k = 10$, and $p = 0$.

Figure 3.6: Learning Speed

Furthermore, Figure (3.6a) illustrates that the learning speed may not be monotonic in the rating system's granularity. On the one hand, increasing granularity allows customers to reflect on their experiences more accurately. On the other hand, it reduces the likelihood of customers reflecting the same experience, thus reducing the strength of each rating as an information piece. This, in turn, requires more customers to rate the product to uncover its intrinsic quality. However, a significant increase in granularity eventually speeds up learning due to allowing the possibility of significantly informative rating options.

Finally, Figure (3.6b) demonstrates that an increase in customer heterogeneity slows down learning. This finding aligns with our initial intuition: while an increase in heterogeneity enhances the likelihood of observing more comprehensive ratings and ensures convergence to the intrinsic quality, it adds to the time required for such ratings to arise.

3.8 Conclusion

We studied the asymptotic behavior of the ratings in the presence of the customer disconfirmation bias, where customers reflect the discrepancy between their prior expectations and post-purchase experience in their ratings. In this environment, high ratings from past customers induce high expectations, making the product more likely to offer a lower-than-expected quality and receive lower ratings from new customers. These lower ratings, in turn, can reduce expectations and lead to higher future ratings. In this situation, the convergence of the ratings to the intrinsic quality is unclear. We investigated this convergence in three rating systems: a system where customers observe individual past customer ratings (complete system), only observe the frequency of each rating option (aggregate system), and only observe the average of past ratings.

We show that this convergence depends on customer heterogeneity, disconfirmation bias, and the rating system's granularity. Specifically, we can divide the effect of the disconfirmation bias into three regions: when the disconfirmation bias is small, the ratings converge to the intrinsic quality in both the complete and aggregate systems. In this case, customers decouple the effect of heterogeneous experiences and their cognitive biases and eventually learn the intrinsic quality. When the disconfirmation bias is large, the ratings may converge to a value different from the intrinsic quality. However, the gap between the customer beliefs and intrinsic quality will be small. Particularly, when customers only rate based on their disconfirmation, this gap will be zero, and customers learn the intrinsic quality correctly. When the disconfirmation bias is intermediate, the discrepancy between beliefs and the intrinsic quality can be large. This corresponds to the case where the reference utility has the least information regarding whether the ratings result from customer heterogeneity or disconfirmation bias. Hence, the case of intermediate disconfirmation bias can be most problematic for customers and platforms.

We find similar insights when customers only observe the average ratings with some nuances. Specifically, while the beliefs converge to a value that is strictly increasing in the intrinsic quality when the disconfirmation bias is small, this converged value has a systematic error. In other words, customers can correctly rank the products but cannot precisely uncover the intrinsic quality. In this circumstance, a belief-correcting rule can fix this error. However, when the disconfirmation bias is large, the converged value only weakly increases in the intrinsic quality. Hence, customers cannot vertically rank the products.

Although correct learning may not be possible for a "given" rating system and disconfirmation bias, we show that the platform can guarantee correct learning by increasing the granularity of the rating system. A more granular system allows customers to more precisely reflect on their reference utility, which enables future customers to infer further information about the intrinsic quality. We also find that the platform requires higher granularity for correct learning as customer heterogeneity and the disconfirmation bias increase. However, the number of required rating options increases proportionally to $1/(1-\alpha)$. Hence, a significant increase in granularity is only required when α (i.e., the weight of the disconfirmation bias in customer reference utility) is large. In this case, the error in the converged belief would be small. Thus, the platform can maintain granularity without sacrificing much accuracy in learning the intrinsic quality. This finding illustrates the significance of the platform design on the evolution of ratings in the presence of the disconfirmation bias.

We also studied the learning speed in the aggregate system. While the literature primarily focuses on the case of high heterogeneity, we characterize the learning speed for various heterogeneity levels. When the heterogeneity is small, we find a lower bound for the learning speed as a function of customer heterogeneity, disconfirmation bias, and rating system granularity. This lower bound is increasing in the rating system's granularity. Hence, the platform can make convergence faster by giving more rating options to customers. However, the learning speed is not necessarily monotonic in the disconfirmation bias. Specifically, when customer beliefs are close to the intrinsic quality, disconfirmation bias can speed up learning. However, disconfirmation bias may slow down learning when customer beliefs are far from the intrinsic quality. In the case of high customer heterogeneity, we bound the convergence rate by the Kullback-Leibler (KL) divergence of relevant probability distributions and establish exponentially fast learning. When the domain of quality realizations is binary, we show that the bounds are tight, and the convergence rate precisely follows the KL divergence between the probability of customer actions when the belief about quality is correct vs. when it is incorrect.

We consider the following managerial implications. First, the study underscores the critical role of rating system design in shaping customer perceptions and learning about product quality. Businesses can leverage these findings by implementing more granular rating systems to ensure customers can accurately reflect their experiences. By increasing the number of rating options, companies can help customers express nuanced opinions, which mitigates the effects of disconfirmation bias and enhances the accuracy of the perceived product quality. For example, a more detailed rating system on e-commerce platforms can help prospective buyers distinguish between products more effectively, leading to better purchase decisions and higher customer satisfaction. This strategy is particularly beneficial in markets where customer heterogeneity is high, as it ensures diverse customer experiences are accurately captured and conveyed. Second, from a managerial perspective, understanding the dynamics of disconfirmation bias and its impact on customer learning offers valuable insights for strategic decision making. Businesses can utilize these insights to design marketing and customer engagement strategies that foster positive customer experiences and encourage accurate ratings. For instance, companies can implement follow-up surveys to gather detailed feedback and provide corrective measures for customers who had negative experiences. Additionally, businesses can strategically manage their product listings and pricing based on the observed patterns of customer ratings and feedback. By anticipating how ratings evolve and influence customer perceptions, managers can proactively adjust their strategies to maintain a favorable market position and improve long-term customer loyalty.

In summary, we show that the disconfirmation bias has significant implications for the evolution of ratings and platform design. This opens up various avenues for future research. For example, we find that ratings' evolution differs for the customers and platforms. While platforms observe the individual ratings, customers may not have access to this information in an aggregate system. This creates information asymmetry and potential for information provisioning. Furthermore, the platform's pricing strategy affects customer reference utility. The interplay between prices and disconfirmation bias is an interesting future research direction. Other problems, such as the platform's assortment display and customers' multidimensional rating of various attributes of the purchase, are other problems worth investigating in the presence of the disconfirmation bias.

4.0 Conclusion

This thesis considered the information and platform design problems from the firm's and customer's perspectives. It bridges the gap between theoretical models and practical applications, offering guidelines for firms to enhance customer trust and market efficiency through informed decision-making. We explored these issues in two settings: the firm's optimal information revelation to alleviate customer concerns about Personalized Pricing (PP), and customers' social learning to infer intrinsic quality using ratings with disconfirmation bias.

Chapter 2 studied the firm's optimal information provision problem around PP, where the firm customizes prices for individual customers. Specifically, we explored whether the prices can informatively signal PP to customers and how firms should adjust pricing strategies in response to customer reactions. We also investigated whether disclosing inventory information could benefit firms and customers, ultimately advocating for increased transparency in PP practices. We find that firms reduce the first-period price to persuade high-valuation customers to purchase in the first period even when they do not intend to implement PP. This is because the mere presence of PP risk makes customers reluctant to reveal their identity. The firm then must "compensate" customers to persuade them to reveal their valuations. We show that the price alone cannot perfectly signal the firm's PP intention, hurting all stakeholders. However, an inventory signal where the firm reveals whether inventory availability is high or low, can improve the firm revenue when customers believe the firm conducts PP with a sufficiently low probability. In this case, an inventory signal alleviates customer PP concerns and allows the firm to set higher prices, creating a win-win outcome for all stakeholders. With the growing interest in PP regulation, we propose that requiring firms to disclose inventory availability information could be a viable policy to make PP more transparent and credibly reduce customer concerns.

Chapter 3 considered the rational customers' social learning problem in the presence of disconfirmation bias, where customers reflect the discrepancy between their prior expectations and post-purchase experience in their ratings. In this social learning context, whether customers can correctly infer the intrinsic quality from the ratings is unclear. We investigated this problem in three rating systems: a system where customers observe individual past customer ratings (complete system), only observe the frequency of each rating option (aggregate system), and only observe the average of past ratings. We divide the effect of the disconfirmation bias into three regions: when the disconfirmation bias is small, the ratings converge to the intrinsic quality; when the disconfirmation bias is large, the ratings may converge to a value different from the intrinsic quality. However, the gap between the customer beliefs and intrinsic quality will be small. When the disconfirmation bias is intermediate, the discrepancy between beliefs and the intrinsic quality can be large. In addition to building a general framework for customer learning from the ratings and identifying the conditions for complete learning of the quality of a product, our analysis has the following contributions. First, it shows that the platform can guarantee correct learning by increasing the granularity of the rating system. A more granular system allows customers to more precisely reflect on their reference utility, which enables future customers to infer further information about the intrinsic quality. Second, we also studied the learning speed in the aggregate system. While the literature primarily focuses on the case of high heterogeneity, we characterize the learning speed for various heterogeneity levels. When the heterogeneity is small, we find a lower bound for the learning speed as a function of customer heterogeneity, disconfirmation bias, and rating system granularity. In the case of high customer heterogeneity, we bound the convergence rate by the Kullback-Leibler (KL) divergence of relevant probability distributions and establish exponentially fast learning. In summary, we show that the disconfirmation bias has significant implications for the evolution of ratings and platform design.

This thesis contributes to the growing literature on information, platform, and mechanism design applications in operations management. The findings underscore the importance of transparency, strategic information disclosure, and platform design in modern business practices.

Appendix A Chapter 2

A.1 Proofs from Chapter 2

A.1.1 Proof of Proposition 2.1

Part (i). Myopic customer i purchases the product in period $t \in \{1, 2\}$ if and only if $p_{it} \leq v_i$. Therefore, the expected demand function for each period under uniform price p_t is given by

$$D_t(p_t) = \int_{i \in M} a^m_{it}(p_t) = \begin{cases} 1, & \text{if } p_t \leq v_L \\ \alpha, & \text{if } p_t \in (v_L, v_H] \\ 0, & \text{if } p_t > v_H \end{cases} \tag{26}$$

Given this demand function and the revenue maximization problem (4), we observe that the firm must either set $p_t = v_L$ or $p_t = v_H$ for all t, depending on the inventory realization I. If the firm sets $p_t = v_L$, then it sells all its inventory in both periods and earns total revenue $2Iv_L$. If the firm sets $p_t = v_H$, then it sells only to a fraction α of the market of customers in each period and earns total revenue $2\alpha v_H$. The comparison of these two revenue values yield the result.

Part (ii). A P-type firm can personalize the prices in the second period. Similar to the previous part, myopic consumer i purchases the product in period t if and only if $p_{it} \leq v_i$. Given any inventory realization I, if the firm sets $p_1 = v_L$, then the revenue-to-go at the beginning of period 1 is given by $v_L I + \max(v_H \alpha, v_L I)$. The updated Bayesian belief follows the following rule.

$$\mathbb{P}\left(v_i = v_H | a_{it}(p_{it}, s) = 1\right) = \begin{cases} 1 & \text{if } p_{it} > v_L, \\ \alpha & \text{otherwise} \end{cases} \tag{27}$$

and

$$\begin{aligned} &\mathbb{P}\left(v_i = v_H | a_{it}(p_{it}, s) = 0\right) \\ &= \frac{\mathbb{P}\left(a_{it}(p_{it}, s) = 0 | v_i = v_H\right) \mathbb{P}\left(v_i = v_H\right)}{\mathbb{P}\left(a_{it}(p_{it}, s) = 0 | v_i = v_H\right) \mathbb{P}\left(v_i = v_H\right) + \mathbb{P}\left(a_{it}(p_{it}, s) = 0 | v_i = v_L\right) \mathbb{P}\left(v_i = v_L\right)}. \end{aligned} \tag{28}$$

If it sets $p_1 = v_H$, then due to the updated belief characterized in the above Equation (27), the firm implements PP, i.e., it sets $p_{i2} = v_H$ if $a_{i1}(p_1) = 1$, and $p_{i2} = v_L$ if $a_{i1}(p_1) = 0$. Thus, the revenue from setting $p_1 = v_H$ is given by $v_H\alpha + v_H\alpha + (I - \alpha)v_L$.

The existence of unique threshold I^m follows from the fact that equation $v_H\alpha + v_H\alpha + (I - \alpha)v_L = v_L I + \max(v_H\alpha, v_L I)$ has a unique solution in I, namely I^{m*}. Then, we define $I^m = \min(I^{m*}, 1)$. □

A.1.2 Proof of Lemma 2.1

In the last period, all customers purchase the product if they receive a non-negative utility. Also, L-type customers only purchase when they receive price v_L. Therefore, we focus on the first-period decisions of H-type customers. If H-type customer i accepts the offer at price $p_{i1} > v_L$ in period 1, then in period 2, she either receives personalized price $p_{i2} = v_H$ if the firm is P-type or uniform price $p_2 = v_H$ if the firm is U-type. Thus, the expected utility-to-go is given by

$$u_{i1}(p_{i1}, p_{i2}, \mathcal{H}_{it}^c | a_{i1} = 1) = v_H - p_1 + 0.$$

Alternatively, if this customer does not purchase at price $p_1 > v_L$ in the first period, i.e., if she hides her type, then in period 2, she will receive personalized price $p_{i2} = v_L$ from a P-type firm. A U-type that charges $p_1 > v_L$ in the first period, in equilibrium, charges v_H in the second period. Thus, the customer's expected utility-to-go is given by

$$u_{i1}(p_{i1}, p_{i2}, \mathcal{H}_{it}^c | a_{i1} = 0) = 0 + Pr(\omega = 1 | p_1 > v_L)\xi_{i2}^c(\mathcal{H}_2^c | \omega = 1 \ \& \ p_1 > v_L)(v_H - v_L) + 0.$$

Hence, Customer i's period 1 decision must satisfy

$$a_{i1} = \begin{cases} 1 & \text{if } u_{i1}(p_{i1}, p_{i2}, \mathcal{H}_{it}^c | a_{i1} = 1) \geq u_{i1}(p_{i1}, p_{i2}, \mathcal{H}_{it}^c | a_{i1} = 0) \\ 0 & \text{if otherwise} \end{cases}$$

From the discussion above, we have $u_{i1}(p_{i1}, p_{i2}, \mathcal{H}_{it}^c | a_{i1} = 1) \geq u_{i1}(p_{i1}, p_{i2}, \mathcal{H}_{it}^c | a_{i1} = 0)$ if and only if

$$v_H - p_1 \geq Pr(\omega = 1 | p_1 > v_L)\xi_{i2}^c(\mathcal{H}_2^c | \omega = 1 \ \& \ p_1 > v_L)(v_H - v_L). \tag{29}$$

First, we prove the following claim.

Claim 1. Define $g(I) \equiv v_H\alpha + v_L(I-\alpha) - \max(\alpha v_H, v_L I)$ and $l(I) \equiv v_L I - p_1\alpha$. There exist unique thresholds $\overline{I} = \min(\frac{p_1\alpha}{v_L} + \frac{(v_H-v_L)\alpha}{v_L}, 1)$ and $\underline{I} = \min(\frac{p_1\alpha}{v_L}, 1)$ such that for a P-type firm,

$$p_{i1}(1, I) = \begin{cases} p_1 & \text{if } I \le \overline{I} \\ v_L & \text{if otherwise} \end{cases}$$

and for a U-type firm,

$$p_{i1}(0, I) = \begin{cases} p_1 & \text{if } I \le \underline{I} \\ v_L & \text{if otherwise} \end{cases}$$

Proof of Claim 1. Note that function $g(I)$ is the revenue gain from PP compared to UP. Also, $l(I)$ is the revenue loss due to PP in the first period. Since $g(I)$ and $l(I)$ are continuous and differentiable almost everywhere, $g(\alpha) = 0$, $l(\alpha) < 0$, and $l'(I) \ge g'(I)$ for $\forall I$, then $g(I) - l(I)$ has a single crossing property. The value of the unique crossing inventory is $\frac{p_1\alpha}{v_L} + \frac{(v_H-v_L)\alpha}{v_L}$. As such, $g(I) \ge l(I)$ if and only if $I \le \frac{p_1\alpha}{v_L} + \frac{(v_H-v_L)\alpha}{v_L}$. In other words, if inventory is above this threshold, a P-type firm does not benefit from implementing PP. Since we may have $\frac{p_1\alpha}{v_L} + \frac{(v_H-v_L)\alpha}{v_L} > 1$, we define $\overline{I} = \min(\frac{p_1\alpha}{v_L} + \frac{(v_H-v_L)\alpha}{v_L}, 1)$. Hence, a P-type firm sets price $p_{i1}(1, I) = p_1$ if $I \le \overline{I}$ and price v_L, otherwise. This proves the claim for a P-type firm.

Next, we prove the claim for a U-type firm. Given price p_1, a U-type firm's revenue for the first period is $\max(p_1\alpha, v_L I)$. Therefore, the firm sets price p_1 iff $p_1\alpha \ge v_L I$, or equivalently, iff $I \le \frac{p_1\alpha}{v_L}$. We define $\underline{I} = \min(\frac{p_1\alpha}{v_L}, 1)$. This completes the proof of the claim.

Having this result, now we continue the proof of Lemma 2.1. First, we show that $\overline{p} \in (v_L, v_H]$ exists such that it satisfies the equality in Equation (29). Note that the left side of the inequality in Equation (29) is continuous in price on domain $(v_L, v_H]$ and has the range $[0, v_H - v_L)$.

Furthermore, from Claim 1, the right-hand side of the inequality in Equation (29) is given by

$$
\begin{aligned}
&Pr(\omega = 1|p_1 > v_L)\xi^c_{i2}(\mathcal{H}^c_2|\omega = 1 \;\&\; p_1 > v_L)(v_H - v_L) \\
&= \frac{\tau(\overline{I} - \alpha)}{\tau(\overline{I} - \alpha) + (1-\tau)(\underline{I} - \alpha)} \int_{\alpha}^{\overline{I}} \frac{I - \alpha}{1-\alpha} \frac{1}{\overline{I} - \alpha} dF(I)(v_H - v_L) \\
&= \frac{\tau(\overline{I} - \alpha)^2 (v_H - v_L)}{2\left(\tau(\overline{I} - \alpha) + (1-\tau)(\underline{I} - \alpha)\right)(1-\alpha)} < \frac{v_H - v_L}{2},
\end{aligned}
\tag{30}
$$

where the last inequality followed from noting that $\frac{\tau(\overline{I}-\alpha)^2}{[\tau(\overline{I}-\alpha)+(1-\tau)(\underline{I}-\alpha)](1-\alpha)} < 1$. Therefore, the range of the right-hand side of Equation (29) is $(0, \frac{v_H - v_L}{2})$ and $v_H - \frac{\tau(\overline{I}-\alpha)^2(v_H-v_L)}{2[\tau(\overline{I}-\alpha)+(1-\tau)(\underline{I}-\alpha)](1-\alpha)} \in [\frac{v_H+v_L}{2}, v_H]$. Hence, from the Brower's fixed point theorem, this equation has a fixed point. Consequently, there exists at least a solution $\overline{p}$ such that Equation (29) holds with equality.

Next, we show the uniqueness of this solution. Define

$$
y(p_1) \equiv 2(v_H - p_1)\left(\tau(\overline{I} - \alpha) + (1-\tau)(\underline{I} - \alpha)\right)(1-\alpha) - \tau(\overline{I} - \alpha)^2(v_H - v_L). \tag{31}
$$

Then, the equality holds in Equation (29) if and only if $y(p_1) = 0$. We consider three cases: (a) $\underline{I} < \overline{I} \leq 1$, (b) $\underline{I} < 1 = \overline{I}$, and (c) $\underline{I} = \overline{I} = 1$.

If solution $\overline{p}$ satisfies Case (a), then

$$
\begin{aligned}
y_a(p_1) \equiv y(p_1) =& 2(v_H - p_1)(1-\alpha)[\tau(\frac{p_1\alpha}{v_L} + \frac{(v_H - v_L)\alpha}{v_L}) + (1-\tau)\frac{p_1\alpha}{v_L} - \alpha] \\
&- (v_H - v_L)\tau(\frac{p_1\alpha}{v_L} + \frac{(v_H - v_L)\alpha}{v_L} - \alpha)^2 = 0.
\end{aligned}
\tag{32}
$$

Note that $y_a(p_1)$ is concave and quadratic in p_1. Also, $y_a(v_H) = -(v_H - v_L)\tau(\frac{v_H\alpha}{v_L} + \frac{(v_H-v_L)\alpha}{v_L} - \alpha)^2 < 0$ and $y_a(v_L) > 0$. Therefore, a unique value of $\overline{p}$ is guaranteed in the region where Case (a) holds. Moreover, we have $y_a'(\overline{p}) < 0$. Also, $\overline{I}$ is nondecreasing in p_1. Let $p_o = \inf\{p : \overline{I} = 1\}$. Hence, as p_1 increases beyond p_o, $y(p_1)$ switches from $y_a(p_1)$ to

$$
y_b(p_1) \equiv 2(v_H - p_1)(1-\alpha)[\tau + (1-\tau)\frac{p_1\alpha}{v_L} - \alpha] - (v_H - v_L)\tau(1-\alpha)^2, \tag{33}
$$

where $y_b(p_o) = y_a(p_o)$. Additionally, we observe that $y_b(p_1)$ is concave and quadratic in p_1. If $\overline{p}$ satisfies Case (a), then we must have $y_b'(p_o) \leq 0$, which shows that a solution does not exist that satisfies Case (b). Also, $\underline{I}$ is nondecreasing in p_1. Let $p_e = \inf\{p : \underline{I} = 1\}$. Hence, as p_1 increases beyond p_e, $y(p_1)$ switches from $y_b(p_1)$ to $y_c(p_1) = 2(v_H - p_1)(1-\alpha)^2 - (v_H - v_L)\tau(1-\alpha)^2$, which is linearly decreasing in p_1 with $y_c(p_e) = y_b(p_e)$. Hence, a solution cannot satisfy

Case (c). In sum, if a solution satisfies Case (a), it is unique. Following similar steps, one can argue the uniqueness of the solution if it belongs to Cases (b) and (c). This completes the uniqueness of solution $\overline{p}$.

Next, we prove the monotonicity of $\overline{p}$ in τ. First, we consider $\overline{p}$ that satisfies $y_a(\overline{p}) = 0$. From implicit differentiation with respect to τ, we obtain

$$\begin{aligned}\frac{\partial y_a(p_1)}{\partial p_1}\frac{\partial p_1}{\partial \tau}(\overline{p}) =&(v_H - v_L)(\frac{\overline{p}\alpha}{v_L} + \frac{(v_H - v_L)\alpha}{v_L} - \alpha)^2 - 2(v_H - \overline{p})(1-\alpha)\frac{(v_H - v_L)\alpha}{v_L} \\ =&(v_H - v_L)(\overline{I} - \alpha)^2\left(1 - \frac{\overline{I} - \underline{I}}{\overline{I} - \alpha}\frac{\tau(\overline{I} - \alpha)}{\tau(\overline{I} - \alpha) + (1-\tau)(\underline{I} - \alpha)}\right),\end{aligned} \tag{34}$$

where the last equality followed from substituting $(v_H - \overline{p})$ with $\frac{\tau(\overline{I}-\alpha)^2(v_H - v_L)}{2\left(\tau(\overline{I}-\alpha)+(1-\tau)(\underline{I}-\alpha)\right)(1-\alpha)}$ from Equation (30). Since $\frac{\overline{I}-\underline{I}}{\overline{I}-\alpha} < 1$ and $\frac{\tau(\overline{I}-\alpha)}{\tau(\overline{I}-\alpha)+(1-\tau)(\underline{I}-\alpha)} < 1$, Equation (34) is positive, which implies that $\frac{\partial p_1}{\partial \tau}(\overline{p}) < 0$.

Similarly for Cases (b) and (c), we obtain

$$\begin{aligned}\frac{\partial y_b(p_1)}{\partial p_1}\frac{\partial p_1}{\partial \tau}(\overline{p}) &= (v_H - v_L)(1-\alpha)^2 - 2(v_H - \overline{p})(1-\alpha)(1-\underline{I}) \\ &= (v_H - v_L)(1-\alpha)^2\left(1 - \frac{1-\underline{I}}{1-\alpha}\frac{\tau(1-\alpha)}{\tau(1-\alpha) + (1-\tau)(\underline{I} - \alpha)}\right) > 0,\end{aligned}$$

and

$$\frac{\partial y_c(p_1)}{\partial p_1}\frac{\partial p_1}{\partial \tau}(\overline{p}) = (1-\alpha)^2(v_H - v_L) > 0.$$

Since $\overline{p}$ is continuous in τ in each case discussed above and it is monotonic in τ, the result follows. □

A.1.3 Proof of Proposition 2.2

The result follows from the proofs of Proposition 2.1 and Lemma 2.1. Also, demand is constant in price for values greater than v_L until it drops to 0, as price increases beyond a threshold. Since the firm revenue is increasing in price for a given demand, both firm types set the maximum price that will be accepted by the H-type customers, if they decide not to sell to the L-type customers. The maximum price that will be accepted by the H-type customers is $\overline{p}$ defined in Lemma 2.1. $\square$

A.1.4 Proof of Corollary 1

The proof follows from the discussion before the corollary and Proposition 2.2. $\square$

A.1.5 Proof of Proposition 2.3

Part (i). From Proposition 2.2, $\overline{I} = \min(\frac{\overline{p}\alpha}{v_L} + \frac{(v_H - v_L)\alpha}{v_L}, 1)$, and $\underline{I} = \min(\frac{\overline{p}\alpha}{v_L}, 1)$. Thus, $\gamma(\tau, \alpha) = \overline{I} - \underline{I} = \min(\frac{\overline{p}\alpha}{v_L} + \frac{(v_H - v_L)\alpha}{v_L}, 1) - \min(\frac{\overline{p}\alpha}{v_L}, 1)$. Consider the Cases (a), (b), and (c) defined in the proof of Lemma 2.1. Since $\overline{p}$ is non-increasing in τ, as τ increases, the conditions switch from Case (c) to Case (b) and then Case (a). For sufficiently small τ, Case (c) holds and $\gamma(\tau, \alpha) = 0$. For an intermediate value of τ, Case (b) holds and $\gamma(\tau, \alpha) = 1 - \frac{\overline{p}\alpha}{v_L}$, which is increasing in τ. For sufficiently large value of τ, we have $\gamma(\tau, \alpha) = \frac{(v_H - v_L)\alpha}{v_L}$, which is independent of τ. Since $\gamma(\tau, \alpha)$ is continuous in τ, Part (i) follows.

Part (ii). First, we characterize the behavior of $\underline{I}$ in α in the following lemma.

Lemma A.1. *$\underline{I}$ is monotonically non-decreasing in α.*

Proof of Lemma A.1. If α is such that Case (a) holds, then $v_H - \overline{p} = \frac{(v_H - v_L)\tau(\overline{I} - \alpha)^2}{2(1-\alpha)[\tau\overline{I} + (1-\tau)\underline{I} - \alpha]}$.

Multiplying both side of the equation by $\frac{\alpha}{v_L}$ and replacing $\frac{\overline{p}\alpha}{v_L}$ by $\underline{I}$, we derive

$$\frac{v_H\alpha}{v_L} - \underline{I} = \frac{(v_H - v_L)\tau\alpha(\overline{I} - \alpha)^2}{2v_L(1-\alpha)[\tau\overline{I} + (1-\tau)\underline{I} - \alpha]}.$$

Using this equations, from implicit differentiation we find

$$\begin{aligned}
\frac{\partial \underline{I}}{\partial \alpha} &= [-v_H(1-\alpha)(\underline{I} + \frac{(v_H - v_L)\tau\alpha}{v_L} - \alpha) + 2(\alpha v_H - \underline{I}v_L)(\underline{I} + \frac{(v_H - v_L)\tau\alpha}{v_L} - \alpha) \\
&- 2(\alpha v_H - \underline{I}v_L)(1-\alpha)(\frac{(v_H - v_L)\tau}{v_L} - 1) + \tau(v_H - v_L)(\underline{I} + \frac{(v_H - v_L)\alpha}{v_L} - \alpha)^2 \\
&+ 2\tau\alpha\frac{v_H - 2v_L}{v_L}(v_H - v_L)(\underline{I} + \frac{(v_H - v_L)\alpha}{v_L} - \alpha)]/ \\
&\left(2[-v_L(1-\alpha)(\underline{I} + \frac{(v_H - v_L)\alpha\tau}{v_L} - \alpha) + (\alpha v_H - \underline{I}v_L)(1-\alpha) - (v_H - v_L)\tau\alpha(\underline{I} + \frac{(v_H - v_L)\alpha}{v_L} - \alpha)]\right).
\end{aligned} \tag{35}$$

With some algebraic manipulations we derive the following for the denominator of Equation (35):

$$\begin{aligned}
&- v_L(1-\alpha)(\underline{I} + \frac{(v_H - v_L)\alpha\tau}{v_L} - \alpha) + (\alpha v_H - \underline{I}v_L)(1-\alpha) - (v_H - v_L)\tau\alpha(\underline{I} + \frac{(v_H - v_L)\alpha}{v_L} - \alpha) \\
&= [v_L(\alpha - 1) + (v_L - v_H)\alpha]\frac{(v_H - v_L)\alpha\tau}{v_L} + (v_L - v_H)\tau\alpha(\underline{I} - \alpha) + (1-\alpha)\alpha(v_H + v_L - 2\overline{p}).
\end{aligned}$$

Since the terms $(\alpha - 1)$, $(v_L - v_H)$, and $(v_H + v_L - 2\overline{p})$ are negative and the rest of the terms in the above equation are positive, the denominator of $\frac{\partial \underline{I}}{\partial \alpha}$ defined in Equation (35) is negative. For the numerator of Equation (35), we have

$$\begin{aligned}
&- v_H(1-\alpha)(\underline{I} + \frac{(v_H - v_L)\tau\alpha}{v_L} - \alpha) + 2(\alpha v_H - \underline{I}v_L)(\underline{I} + \frac{(v_H - v_L)\tau\alpha}{v_L} - \alpha) \\
&- 2(\alpha v_H - \underline{I}v_L)(1-\alpha)(\frac{(v_H - v_L)\tau}{v_L} - 1) + \tau(v_H - v_L)(\underline{I} + \frac{(v_H - v_L)\alpha}{v_L} - \alpha)^2 \\
&+ 2\tau\alpha\frac{v_H - 2v_L}{v_L}(v_H - v_L)(\underline{I} + \frac{(v_H - v_L)\alpha}{v_L} - \alpha) \\
=&2\tau\frac{(v_H - v_L)\alpha}{v_L}v_H(\overline{I} - 1) + 3\tau(v_H - v_L)\alpha(\alpha - \overline{I}) + \alpha\tau(v_H - v_L)(\overline{I} - \alpha)(\frac{\overline{I} - \alpha}{1-\alpha} - 1) \\
&+ 2(1-\alpha)\alpha\frac{v_H v_L - \overline{p}^2}{v_L}.
\end{aligned}$$

Since $(\overline{I} - 1)$, $(\alpha - \overline{I})$, $(\frac{\overline{I}-\alpha}{1-\alpha} - 1)$, and $(v_H v_L - \overline{p}^2)$ are negative, we find that the numerator of Equation (35) is also negative. Combining the results above, we have $\frac{\partial \underline{I}}{\partial \alpha} > 0$. Thus, Lemma

A.1 holds when α satisfies Case (a). Moreover, note that $\overline{I} = \underline{I} + \frac{(v_H - v_L)\alpha}{v_L}$. Therefore, $\overline{I}$ is also strictly increasing when Case (a) holds.

When Case (b) holds, then $v_H - \overline{p} = \frac{(v_H - v_L)\tau(1-\alpha)}{2[\tau + (1-\tau)\underline{I} - \alpha]}$. Similar to the previous case, from implicit differentiation we have

$$\frac{\partial \underline{I}}{\partial \alpha} = \frac{\overline{p}}{v_L} + \frac{\alpha}{v_L}\frac{\partial p_1}{\partial \alpha}(\overline{p}),$$

where

$$\begin{aligned}\frac{\partial p_1}{\partial \alpha}(\overline{p}) =& \frac{2(v_H - \overline{p})[\tau(1-\alpha) + (1-\tau)(\underline{I} - \alpha)] - 2(1-\alpha)(v_H - \overline{p})[(1-\tau)\frac{\overline{p}}{v_L} - 1] - 2(v_H - v_L)\tau(1-\alpha)}{\frac{\partial y_b(p)}{\partial p}} \\ =& \frac{-\frac{1-\alpha}{\alpha}\left[\frac{(v_H - v_L)\tau}{2} - (v_H - \overline{p})\tau\right]}{\frac{\partial y_b(p)}{\partial p}} > 0,\end{aligned}$$

where the positive sign followed from noting that $\frac{\partial y_b(p)}{\partial p}(\overline{p}) < 0$. This shows that when Case (b) holds, $\underline{I}$ is also strictly increasing in α. Furthermore, in this case $\overline{I} = 1$.

When Case (c) holds, $\underline{I} = \overline{I} = 1$, which are constant in α.

Finally, note that as α increases, the cases switch from Case (a) to Case (b), and then Case (c). Given the behavior of $\underline{I}$ and $\overline{I}$ in α discussed above, and from noting that $\gamma(\tau, \alpha) = \overline{I} - \underline{I}$, the result of Part (ii) follows.

Part (iii). Follows from the proofs of Parts (i) and (ii). □

A.1.6 Proof of Lemma 2.2

In the last period, customers purchase the product if they receive a nonnegative utility.

In period 1, when the firm sends the LI signal, customers update their beliefs about the inventory to a uniform distribution on $[\alpha, I_c]$. The proof of the existence $\overline{I}_b$ and $\underline{I}_b$ then follows by repeating the steps in the proof of Lemma 2.1, and redefining $\overline{I} = \min(\frac{p_b\alpha}{v_L} + \frac{(v_H - v_L)\alpha}{v_L}, I_c)$ and $\underline{I} = \min(\frac{p_b\alpha}{v_L}, I_c)$, where p_b is the first-period price when the firm sends the LI signal.

Furthermore, by dividing the space into the following three cases (a) $I_c \leq \underline{I}_b$, (b) $\underline{I}_b < I_c \leq \overline{I}_b$, and (c) $\overline{I}_b < I_c$, we can characterize $\overline{p}_b$ given in the Lemma. However, we note that the characterization of $\overline{p}_b$ is an implicit function when Case (b) occurs. We prove that in this case, $\overline{p}_b$ is well-defined and possesses a unique value.

Note that $\underline{I}_b$ is a function of p_b. In this proof, we denote this dependence by using $\underline{I}_b(p_b)$. With this dependence in mind, in Case (b), $\overline{p}_b(I_c)$ is the solution to the implicit Equation $p_b = v_H - \frac{v_H - v_L}{2}\frac{I_c-\alpha}{1-\alpha}\frac{\tau(I_c-\alpha)}{\tau I_c+(1-\tau)\underline{I}_b(p_b)-\alpha}$. From the definition of $\underline{I}_b(p_b)$, then p_b is the solution to

$$\begin{aligned} &p_b = v_H - \frac{v_H - v_L}{2}\frac{I_c-\alpha}{1-\alpha}\frac{\tau(I_c-\alpha)}{\tau I_c+(1-\tau)\frac{p_b\alpha}{v_L}-\alpha}, \text{or equivalently,} \\ &2(v_H - p_b)(1-\alpha)\left(\tau I_c + (1-\tau)\frac{p_b\alpha}{v_L} - \alpha\right) = (v_H - v_L)(I_c-\alpha)^2\tau. \end{aligned} \tag{36}$$

The above function defines an Ellipse, where its upper curve crosses the point $(I_c, p_b) = (\alpha, v_H)$. Solving this equation in p_b yields that it has a unique solution in the interval $(\overline{p}, v_H]$. This completes the proof for the uniqueness and existence of $\overline{p}_b(I_c)$.

Next, we consider the case where the firm sends the HI signal. In this case, customers update their inventory beliefs to a uniform distribution on $[I_c, 1]$. Following the steps of Proof of Lemma 2.1 and by redefining $\overline{I} = \min(\frac{p_a\alpha}{v_L} + \frac{(v_H-v_L)\alpha}{v_L}, 1)$, where p_a is the first-period price when the firm sends the HI signal, we can show the existence of $\overline{I}_a$.

To prove the existence of $\underline{I}_a$, note that if such threshold exists, it must satisfy $\underline{I}_a \geq \underset{\sim}{I}_a \equiv \min(\frac{p_a\alpha}{v_L}, 1)$, because $\underset{\sim}{I}_a$ is the threshold below which a U-type firm sets a price higher than v_L. We divide the space into the following three cases: (a) $I_c \leq \underline{I}_a$, (b) $\underline{I}_a < I_c \leq \overline{I}_a$, and (c) $\overline{I}_a < I_c$. In all these cases, we note that $\overline{p}_a(I_c)$ is implicitly defined. In Case (b), $\overline{p}_a(I_c)$ is the solution to

$$p_a = v_H - \frac{I_c + \overline{I}_a(I_c) - 2\alpha}{2(1-\alpha)}(v_H - v_L) \equiv v_H - J(p_a, I_c). \tag{37}$$

Since $\overline{I}_a$ is nondecreasing in p_a, function $v_H - J(p_a, I_c)$ is nonincreasing in p_a, which suggests that it has a unique fixed point. Solving the equation gives us the value of the price characterized in the Lemma for Case (b).

In Case (c), $\overline{p}_a(I_c)$ is the solution to $p_a = v_H - \frac{2\overline{I}_a(I_c)-2\alpha}{2(1-\alpha)}(v_H - v_L)$. Once again, the right-hand side of the equation is nonincreasing in p_a, and therefore, has a unique fixed point. Solving for this fixed point, gives the price characterized in the lemma for Case (c).

When Case (a) occurs, $\overline{p}_a(I_c)$ is the solution to

$$p_a = v_H - \frac{\tau(v_H - v_L)(\overline{I}_a(I_c) - I_c)}{(1-\tau)\underline{I}_a(I_c) + \tau\overline{I}_a(I_c) - I_c} \cdot \frac{I_c + \overline{I}_a(I_c) - 2\alpha}{2(1-\alpha)}. \tag{38}$$

If a unique solution $\overline{p}_a(I_c)$ to the Equation (38) exists such that it is higher than the one found in Case (b), then $\underline{I}_a(I_c) = \underset{\sim}{I}_a(I_c)$ and Equation (38) turns into:

$$p_a = v_H - \frac{\tau(v_H - v_L)(\overline{I}_a(I_c) - I_c)}{(1-\tau)\underset{\sim}{I}_a(I_c) + \tau\overline{I}_a(I_c) - I_c} \cdot \frac{I_c + \overline{I}_a(I_c) - 2\alpha}{2(1-\alpha)} \equiv v_H - H(p_a, I_c). \tag{39}$$

Since both $\overline{I}_a(I_c)$ and $\underset{\sim}{I}_a(I_c)$ are minimization functions truncated at 1, we further divide Case (a) into three subcases: Case a1) $\overline{I}_a(I_c) < 1$, Case a2) $\underset{\sim}{I}_a(I_c) < 1 \leq \overline{I}_a(I_c)$, and Case a3) $\underline{I}_a(I_c) = 1$. For Case (a3), the right-hand side of Equation (39) is independent of p_a, and therefore $\overline{p}_a(I_c)$ is explicitly defined by the equation and is decreasing in I_c. Furthermore, this solution is the same as the solution we find in Case (b). Hence, $\underline{I}_a(I_c)$ exists and it is equal to $\underset{\sim}{I}_a(I_c)$.

For Cases (a1) and (a2), we can rewrite the Equation (39) as follows:

$$\tau(v_H - v_L)(I_c - N_1)^2 - 2(1-\alpha)p_a I_c - \left(2(1-\alpha)\frac{\alpha}{v_L} + \tau(v_H - v_L)\frac{\alpha^2}{v_L^2}\right)(p_a - N_2)^2 = C, \tag{40}$$

where N_1, N_2, and C are independent of p_a and I_c. Note that Equation (40) represent an angled hyperbola in $(I_c, \overline{p}_a)$ space. Furthermore, this hyperbola crosses the point $(I_c, p_a) = (\alpha, \overline{p})$. Hence, for any threshold I_c, there exists at most two prices $\overline{p}_a$ that satisfy Equation (39). Also, note that

$$\frac{\partial \overline{p}_a}{\partial I_c}|_{\{I_c=\alpha\}} = \frac{\tau(v_H - v_L)\left(I_c - \alpha - \frac{(\overline{I}_a - I_c)}{(\overline{I}_a - \frac{(1-\tau)(v_H - v_L)\alpha}{v_L} - I_c)} \cdot \frac{(\overline{I}_a + I_c - 2\alpha)}{2}\right)}{\frac{\alpha}{v_L}\left((1-\alpha)(2p_a - I_c\frac{v_L}{\alpha} - v_H) + \tau(v_H - v_L)(1 - \alpha + \overline{I}_a - \alpha)\right)}|_{\{I_c=\alpha\}}.$$

Since $\frac{(\overline{I}_a - I_c)}{(\overline{I}_a - \frac{(1-\tau)(v_H - v_L)\alpha}{v_L} - I_c)} > 1$, we have $I_c - \alpha - \frac{(\overline{I}_a - I_c)}{(\overline{I}_a - \frac{(1-\tau)(v_H - v_L)\alpha}{v_L} - I_c)} \cdot \frac{(\overline{I}_a + I_c - 2\alpha)}{2} < I_c - \alpha - \frac{(\overline{I}_a + I_c - 2\alpha)}{2} < 0$. Hence, the numerator of the above equation is negative. We also know that $1 - \alpha + \overline{I}_a - \alpha > 0$. Furthermore, at $I_c = \alpha$, we have $2\overline{p}_a - v_L - v_H = 2\overline{p} - v_L - v_H > 0$; therefore, the denominator of the above equation is positive. As such, $\frac{\partial \overline{p}_a}{\partial I_c} < 0$ at $I_c = \alpha$. This shows that the hyperbola defined by Equation (40) is decreasing and is less than or equal to $\overline{p}$ at $I_c = \alpha$. Since the equation defined in (40) is a hyperbola, it is either concave or convex in I_c at $I_c = \alpha$. Using the notations of functions $J(p_a, I_c)$ and $H(p_a, I_c)$ defined in Equations (37) and (39), we have

$$\frac{H(p_a, I_c)}{J(p_a, I_c)} = \frac{\tau(\overline{I}_a - I_c)}{\tau(\overline{I}_a - I_c) + (1-\tau)(\underset{\sim}{I}_a - I_c)} = \tau + \frac{\tau(1-\tau)\frac{(v_H - v_L)\alpha}{v_L}}{\overline{I}_a - (1-\tau)\frac{(v_H - v_L)\alpha}{v_L} - I_c} \in (\tau, 1].$$

We observe that $\frac{H(p_a,I_c)}{J(p_a,I_c)}$ is increasing in I_c. Since $J(p_a, I_c)$ is linearly increasing in I_c, we find that $\overline{p}_a$ is concave in I_c at $I_c = \alpha$.

Let $\tilde{p}_a(I_c)$ be the price when I_c satisfies Case (a) and $\hat{p}_a(I_c)$ be the price when I_c satisfies Case (b). Since we showed that $\hat{p}_a(I_c)$ is linearly nonincreasing in I_c, it crosses $\tilde{p}_a(I_c)$ at most once. If this crossing point occurs above the vertex of the hyperbola, then $\underline{I}_a(I_c)$ exists and is equal to $\underset{\sim}{I}_a(I_c)$. However, if the crossing point occurs below the vertex, then $\underline{I}_a(I_c)$ might not be equal to $\underset{\sim}{I}_a(I_c)$. In this case, the customer believes that a U-type firm would never charge a price higher than v_L if I_c is greater than the x-coordinate of the vertex. Thus, $\underline{I}_a(I_c)$ would be the x-axis value of the vertex of the hyperbola.

Finally, we prove the monotonicity of $\overline{p}_b$ and $\overline{p}_a$ in I_c. First, we show that $\overline{p}_b(I_c)$ is piece-wise nonincreasing in I_c. In Case (a), $\overline{p}_b(I_c)$ is linearly decreasing in I_c. In Case (c), $\overline{p}_b(I_c) = \overline{p}$, which is constant in I_c. In Case (b), we have

$$\begin{aligned}
\frac{\partial \overline{p}_b}{\partial I_c} =& \frac{\tau(v_H - v_L)(I_c - \alpha) - \tau(v_H - p_b)(1-\alpha)}{(1-\alpha)\Big((v_H - p_b)(1-\tau)\frac{\alpha}{v_L} - \tau(I_c - \alpha) - (1-\tau)(p_b\frac{\alpha}{v_L} - \alpha)\Big)} \\
=& \frac{(v_H - v_L)(I_c - \alpha)(1 - \frac{\tau(I_c - \alpha)}{2[\tau(I_c-\alpha) + (1-\tau)(\underline{I}_b - \alpha)]})}{(1-\alpha)\Big((1-\tau)\frac{\alpha}{v_L}(v_H + v_L - 2p_b) - \tau(I_c - \alpha)\Big)}.
\end{aligned}$$

From Equation (36), we observe that $v_H + v_L - 2p_b < 0$. Therefore, the denominator is negative. Also, the numerator is positive. Consequently, $\frac{\partial \overline{p}_b}{\partial I_c} < 0$. Hence, $\overline{p}_b(I_c)$ is piecewise nonincreasing in I_c. Since $\overline{p}_b(I_c)$ is piecewise nonincreasing in I_c, we find that $\overline{I}_b$ and $\underline{I}_b$ are also nonincreasing in I_c. This combined with the fact that $\overline{p}_b(I_c)$ is continues in I_c, prove the nonincreasing property of $\overline{p}_b(I_c)$ in I_c.

Next, we show that $\overline{p}_a(I_c)$ is nonincreasing in I_c. From the discussion above, we showed that $\overline{p}_a(I_c)$ is piecewise nonincreasing. Also, $\overline{p}_a(I_c)$ continuously transitions from Case (b) to Case (c). Furthermore, the arguments above show that $\overline{p}_a(I_c)$ is either continuous or jumps down when it transitions from Case (a) to Case (b). This completes the proof of nonincreasing property of $\overline{p}_a(I_c)$.

Finally, since $\overline{p}_b(I_c)$ is nonincreasing in I_c and $\overline{p}_b(\overline{I}_b) = \overline{p}$, we have $\overline{p}_b(I_c) \geq \overline{p}$. Also, since $\overline{p}_a(I_c)$ is nonincreasing in I_c and $\overline{p}_a(\alpha) = \overline{p}$, we have $\overline{p} \geq \overline{p}_a(\alpha)$. This completes the proof of the last part of the lemma. □

A.1.7 Proof of Proposition 2.4

First, we show that the optimal signal cut-off does not belong to the interval $[\alpha, \underline{I}_a]$. To show this, we argue that there exists a signal cut-off $\hat{I}_c \in [\underline{I}, \underline{I}_b)$ that yields an strictly higher revenue than the case when the firm chooses any cut-off $I_c \in [\alpha, \underline{I}_a]$. Denote the revenue with signal cut-off $I_c \in [\alpha, \underline{I}_a]$ by $R_1^C(I_c)$, the revenue with signal cut-off $\hat{I}_c \in [\underline{I}, \underline{I}_b)$ with $R_2^C(\hat{I}_c)$, and the expected revenue in base model with R^B. Then,

$$
\begin{aligned}
R_1^C(I_c) - R^B =& \tau \int_{\overline{I}_a(I_c)}^{\overline{I}} \left(v_L I - (v_H - v_L)\alpha + \frac{(I-\alpha)(v_H - v_L)\alpha}{1-\alpha} - v_H \alpha \right) dF(I) \\
&+ (1-\tau) \int_{\underline{I}_a(I_c)}^{\underline{I}} (v_L I - v_H \alpha) dF(I) \\
<& \tau \int_{\overline{I}_a(\hat{I}_c)}^{\overline{I}} \left(v_L I - (v_H - v_L)\alpha + \frac{(I-\alpha)(v_H - v_L)\alpha}{1-\alpha} - v_H \alpha \right) dF(I) \\
&+ (1-\tau) \int_{\underline{I}_a(I_c)}^{\underline{I}} (v_L I - v_H \alpha) dF(I) \\
\leq& \tau \int_{\overline{I}_a(\hat{I}_c)}^{\overline{I}} \left(v_L I - (v_H - v_L)\alpha + \frac{(I-\alpha)(v_H - v_L)\alpha}{1-\alpha} - v_H \alpha \right) dF(I) \\
&+ (1-\tau) \int_{\underline{I}}^{\hat{I}_c} (v_H \alpha - v_L I) dF(I) \\
=& R_2(\hat{I}_c) - R^B,
\end{aligned}
$$

where the first inequality followed from the fact that the term

$$
\tau \int_{\overline{I}_a(I_c)}^{\overline{I}} \left(v_L I - (v_H - v_L)\alpha + \frac{(I-\alpha)(v_H - v_L)\alpha}{1-\alpha} - v_H \alpha \right) dF(I)
$$

is increasing in I_c and $\hat{I}_c > I_c$. The second inequality followed from noting that $(1 - \tau)\int_{\underline{I}_a(I_c)}^{\underline{I}}(v_L I - v_H\alpha)dF(I) < 0$ and $(1-\tau)\int_{\underline{I}}^{\hat{I}_c}(v_H\alpha - v_L I)dF(I) \geq 0$. Hence, the optimal signal cut-off cannot belong to interval $[\alpha, \underline{I}_a]$.

Since for $I_c \notin [\alpha, \underline{I}_a]$, $\overline{p}_a$ and $\overline{p}_b$ from Lemma 2.2 are continuous in I_c, the firm revenue is also a continuous function. Thus, on the compact interval $(\underline{I}_a, 1]$, the firm revenue achieves its maximum in I_c, which shows the existence of I_c^*.

Part (i) and Part (ii). Follow from Lemma 2.2 and the existence of I_c^*.

Part (iii). It is straightforward to see that for any I_c^*, we must have $\overline{I}_b^* \geq \underline{I}_b^*$ and $\overline{I}_a^* \geq \underline{I}_a^*$, where the inequalities are strict when $\underline{I}_b^* < 1$ and $\underline{I}_a^* < 1$, respectively. □

A.1.8 Proof of Proposition 2.5

From the proof of Proposition 2.4 recall that the optimal cut-off cannot belong to the interval $[\alpha, \underline{I}_a]$. Thus, we only focus on the interval $(\underline{I}_a, 1]$. Let $\underline{I}$ be the threshold for the base model without the inventory signal. First, we prove that if $\underline{I} = 1$, then there does not exist any $I_c^* < 1$ such that an inventory signal can improve the firm's revenue. Let R^B and $R^C(I_c)$ be the expected revenues in the base model and the model with inventory signal, respectively. In this case, i.e., when $\underline{I} = 1$, we have

$$R^C(I_c) - R^B = (1-\tau)\int_{I_c}^{\underline{I}}(v_L I - v_H\alpha)dF(I) < 0,$$

which is what we desired to show.

Next, consider the case where $\underline{I} < 1$. We prove there exists a unique $I_c^* \in [\underline{I}, 1)$ such that the model with the inventory signal yields strictly a higher revenue than the base model.

Consider the following four cases: (i) $\underline{I}_a < I_c < \underline{I}_b$, (ii) $\underline{I}_b \leq I_c \leq \overline{I}_a$, (iii) $\overline{I}_a < I_c \leq \overline{I}_b$, and (iv) $\overline{I}_b < I_c \leq 1$. We observe that the ex-ante expected revenue is equivalent to the base model when I_c is in (iv).

We first show that, for Case (i), $R^C(I_c) - R^B$ is strictly increasing in I_c. In this case, we have

$$
\begin{aligned}
R^C(I_c) - R^B =&\tau \int_{\overline{I}_a}^{\overline{I}} \left(v_L I - (v_H - v_L)\alpha + \frac{(I-\alpha)(v_H - v_L)\alpha}{1-\alpha} - v_H \alpha \right) dF(I) \\
&+ (1-\tau) \int_{\underline{I}}^{I_c} (v_H \alpha - v_L I) dF(I).
\end{aligned}
$$

From the proof of Proposition 2.4, we know that first line of the equation above is strictly increasing in I_c if $\overline{I}_a < 1$, and zero otherwise. We show that the second line is also strictly increasing in I_c. The first derivative of the second line is $v_L I_c - v_H \alpha$. Since $I_c < \frac{v_H \alpha}{v_L}$ for all $I_c \in (\underline{I}_a, \underline{I}_b)$, we have $v_L I_c - v_H \alpha > 0$. Therefore, $R^C(I_c) - R^B$ is strictly increasing in Case (i).

Next, we prove that the expected revenue in Case (iii) is strictly decreasing in I_c. The expected revenue is Case (iii) is

$$
\begin{aligned}
R^C(I_c) =& \left(\tau(I_c - \alpha) + (1-\tau)(\underline{I}_b - \alpha)\right) v_H \alpha - \frac{\tau\alpha(v_H - v_L)(I_c - \alpha)^2}{2(1-\alpha)} + \tau \int_{\alpha}^{I_c} (v_H - v_L)\alpha dF(I) \\
&+ \tau \int_{I_c}^{1} v_L I dF(I) + (1-\tau) \int_{\underline{I}_b}^{1} v_L I dF(I)
\end{aligned}
$$

Thus, for any I_c from Case (iii), the difference is given by $R^C(I_c) - R^B = \tau v_L \int_{\overline{I}_a}^{I_c} (I - \overline{I}_a) dF(I)$, which is positive for any I_c in Case (iii). Taking the first derivative of this equation yields:

$$
\begin{aligned}
&\frac{\partial (R^C(I_c) - R^B)}{\partial I_c} = \tau \left(\frac{\alpha(v_H - v_L)(1 - I_c)}{1-\alpha} + v_H \alpha - v_L I_c \right) \\
=&\tau \left(\frac{\alpha(v_H - v_L)(1 - \overline{I}_a)}{1-\alpha} + v_H \alpha - v_L \overline{I}_a + \frac{\alpha(v_H - v_L)(\overline{I}_a - I_c)}{1-\alpha} + v_L(\overline{I}_a - I_c) \right) \\
=& \left(\frac{\alpha(v_H - v_L)}{2(1-\alpha)} + v_L \right) (\overline{I}_a - I_c) < 0,
\end{aligned} \tag{41}
$$

where the last inequality followed from noting that $I_c > \overline{I}_a$ in Case (iii). Therefore, $R^C(I_c) - R^B$ is strictly decreasing in I_c in Case (iii).

We claim that the optimal cut-off must belong to Case (ii) by showing that, in this case, the revenue function $R^C(I_c)$ is concave and decreasing at $I_c = \overline{I}_a$. The expected revenue in Case (ii) is

$$
\begin{aligned}
R^C(I_c) =& \left(\tau(\overline{I}_a - \alpha) + (1-\tau)(\underline{I}_b - \alpha)\right) v_H \alpha - \frac{\tau\alpha(v_H - v_L)(\overline{I}_a - \alpha)^2}{2(1-\alpha)} \\
&+ \tau \int_{\alpha}^{\overline{I}_a} (v_H - v_L)\alpha dF(I) + \tau \int_{\overline{I}_a}^{1} v_L I dF(I) + (1-\tau) \int_{\underline{I}_b}^{1} v_L I dF(I).
\end{aligned}
$$

The first derivative of the revenue function is given by

$$\frac{\partial R^C(I_c)}{\partial I_c} = \tau\left(\frac{\alpha(v_H - v_L)(I_c - \overline{I}_a)}{2(1-\alpha)}\right)\frac{\partial \overline{I}_a}{\partial I_c} + (1-\tau)(v_H\alpha - v_L\underline{I}_b)\frac{\partial \underline{I}_b}{\partial I_c}. \tag{42}$$

where on the interval $[\underline{I}_b, \overline{I}_a]$, we have

$$\begin{aligned}
\overline{I}_a(I_c) =& \frac{2v_H\alpha(1-\alpha) + 2(v_H - v_L)\alpha - I_c\alpha(v_H - v_L)}{2(1-\alpha)v_L + \alpha(v_H - v_L)}, \text{ and} \\
\underline{I}_b(I_c) =& \left(2(1-\tau)v_L(1-\alpha)\right)^{-1}\alpha(1-\alpha)(v_H + v_L - v_H\tau) - v_L(1-\alpha)\tau I_c \\
&+ (1-\alpha)^{\frac{1}{2}}[(2\alpha v_L(1-\tau)(\alpha^2(v_H(2-\tau) + v_L\tau) + \tau I_c^*(v_H(2-I_c) + v_L I_c) - 2\alpha(v_H + v_L\tau I_c)) \\
&+ (1-\alpha)(\alpha(v_H + v_L - v_H\tau) - v_L\tau I_c)^2]^{\frac{1}{2}}
\end{aligned}$$

Plugging these values in Equation (41), we find that $\tau\left(\frac{\alpha(v_H-v_L)(I_c-\overline{I}_a)}{2(1-\alpha)}\right)\frac{\partial\overline{I}_a}{\partial I_c} > 0$. Furthermore, $\tau\left(\frac{\alpha(v_H-v_L)(I_c-\overline{I}_a)}{2(1-\alpha)}\right)\frac{\partial\overline{I}_a}{\partial I_c}$ is strictly decreasing in I_c, because $\frac{\partial\overline{I}_a}{\partial I_c} < 0$ is constant and $\frac{\alpha(v_H-v_L)(I_c-\overline{I}_a)}{2(1-\alpha)} < 0$ is increasing in I_c. Moreover, since $v_H\alpha - v_L\underline{I}_b > 0$ is increasing in I_c, and $\frac{\partial \underline{I}_b}{\partial I_c} < 0$ is decreasing in I_c, we observe that $(1-\tau)(v_H\alpha - v_L\underline{I}_b)\frac{\partial\underline{I}_b}{\partial I_c} < 0$ and it is strictly decreasing in I_c. Therefore, the first order derivative is strictly decreasing, and $R^C(I_c)$ is a concave function in the range $[\underline{I}_b, \overline{I}_a]$.

Furthermore, the first derivative of the revenue function at $I_c = \overline{I}_a$ is given by

$$\begin{aligned}
\frac{\partial R^C(I_c)}{\partial I_c}|_{\{I_c=\overline{I}_a\}} =& \tau\left(\frac{\alpha(v_H - v_L)(I_c - \overline{I}_a)}{2(1-\alpha)}\right)\frac{\partial \overline{I}_a}{\partial I_c}|_{\{I_c=\overline{I}_a\}} + (1-\tau)(v_H\alpha - v_L\underline{I}_b)\frac{\partial \underline{I}_b}{\partial I_c}|_{\{I_c=\overline{I}_a\}} \\
=& (1-\tau)(v_H\alpha - v_L\underline{I}_b)\frac{\partial \underline{I}_b}{\partial I_c}|_{\{I_c=\overline{I}_a\}} < 0,
\end{aligned}$$

where the last equality followed from noting that $\frac{\alpha(v_H-v_L)(I_c-\overline{I}_a)}{2(1-\alpha)} = 0$ at $I_c = \overline{I}_a$, and the last inequality followed from the facts that $(v_H\alpha - v_L\underline{I}_b) > 0$ and $\frac{\partial\underline{I}_b}{\partial I_c} < 0$. Therefore, at $I_c = \overline{I}_a$, the expected revenue is strictly decreasing. Since the revenue function in this region is concave and the expected revenue is continuous on $(\underline{I}_a, 1]$, the optimal solution I_c^* must exist in the interval $[\underline{I}_b, \overline{I}_a)$, and I_c^* is unique. □

A.1.9 Proof of Proposition 2.6

Let $I_0 \equiv \min(\frac{\alpha v_H}{v_L}, 1)$. To emphasize the dependence of the functions $\overline{I}_a$ and $\underline{I}_b$ on I_c, we denote them by $\overline{I}_a(I_c)$ and $\underline{I}_b(I_c)$. From the definition of the H-type customer surplus and $I_c^* \in [\underline{I}_b, \overline{I}_a)$, we have

$$CS^n = \tau \left(\int_\alpha^{\underline{I}} (v_H - \overline{p}) dF(I) + \int_{\underline{I}}^{I_0} (v_H - v_L) I dF(I) + \int_{I_0}^1 2(v_H - v_L) I dF(I) \right) + (1-\tau) \left(\int_\alpha^{\overline{I}} (v_H - \overline{p}) dF(I) + \int_{\overline{I}}^1 2(v_H - v_L) I dF(I) \right),$$

and,

$$CS^s(I_c) = \tau \left(\int_\alpha^{\underline{I}_b(I_c)} (v_H - \overline{p}_b(I_c)) dF(I) + \int_{\underline{I}_b(I_c)}^{I_0} (v_H - v_L) I dF(I) + \int_{I_0}^1 2(v_H - v_L) I dF(I) \right) + (1-\tau) \left(\int_\alpha^{I_c} (v_H - \overline{p}_b(I_c)) dF(I) + \int_{I_c}^{\overline{I}_a(I_c)} (v_H - \overline{p}_b(I_c)) dF(I) + \int_{\overline{I}_a(I_c)}^1 2(v_H - v_L) I dF(I) \right).$$

Therefore,

$$CS^s(I_c) - CS^n = \frac{\tau}{2(1-\alpha)} \left((\overline{I}_a(I_c) - \alpha)^2 - (\overline{I} - \alpha)^2 \right) + \tau(\overline{I}^2 - \overline{I}_a(I_c)^2) - (1-\tau) \frac{\underline{I}_b(I_c)^2 - \underline{I}^2}{2}. \tag{43}$$

We first show that the function $CS^s(I_c) - CS^n$ is strictly increasing in I_c. The first-order derivative of this function with respect to I_c is

$$\frac{\partial (CS^s(I_c) - CS^n)}{\partial I_c} = \tau \frac{-\alpha - \overline{I}_a + 2\overline{I}_a \alpha}{1-\alpha} \frac{\partial \overline{I}_a}{\partial I_c} - (1-\tau) \underline{I}_b \frac{\partial \underline{I}_b}{\partial I_c} \tag{44}$$

From Lemma 2.2, we know both $\overline{p}_a$ and $\overline{p}_b$ are non-increasing in I_c, therefore, both $\overline{I}_a(I_c)$ and $\underline{I}_b(I_c)$ are also non-increasing in I_c. Consequently, from Equation (44), $\frac{\partial (CS^s(I_c) - CS^n)}{\partial I_c} > 0$, which is what we desired to show. Since $I_c^* \geq \underline{I}_b$, we have

$$CS^s(I_c^*) - CS^n \geq CS^s(\underline{I}_b) - CS^n.$$

Hence, $CS^s(\underline{I}_b) - CS^n$ is a lower bound for $CS^s(I_c^*) - CS^n$. Plugging $I_c = \underline{I}_b$, we have

$$CS^s(\underline{I}_b) - CS^n = -\frac{\tau}{2(1-\alpha)} (\overline{I} - \alpha)^2 + \tau \overline{I}^2 + \frac{\tau}{2(1-\alpha)} (\overline{I}_a(\underline{I}_b) - \alpha)^2 - \tau \overline{I}_a(\underline{I}_b)^2 - (1-\tau) \frac{\underline{I}_b(\underline{I}_b)^2 - \underline{I}^2}{2}.$$

Note that $\underline{I}_b(\underline{I}_b) = \alpha + \frac{2(1-\alpha)(v_H - v_L)\alpha}{2(1-\alpha)v_L + \tau\alpha(v_H - v_L)}$ and $\overline{I}_a(\underline{I}_b) = \frac{2(1-\alpha)v_H\alpha + \alpha(v_H - v_L)(2 - \underline{I}_b)}{2(1-\alpha)v_L + \alpha(v_H - v_L)}$.

Following the proof of Proposition 2.3, we know that $\underline{I}$ and $\overline{I}$ are strictly increasing in α when $\underline{I} < 1$ and $\overline{I} < 1$, respectively. Following the same steps as the proof of Proposition 2.3, one can show that $\overline{I}_a(\underline{I}_b)$ and $\underline{I}_b(\underline{I}_b)$ are strictly increasing in α when $\overline{I}_a(\underline{I}_b) < 1$ and $\underline{I}_b(\underline{I}_b) < 1$, respectively. Therefore, the pattern for $CS^s(\underline{I}_b) - CS^n$ can be divided into 4 cases: Case (i) $\overline{I} < 1$, Case (ii) $\overline{I} = 1$ and $\overline{I}_a(\underline{I}_b) < 1$, Case (iii) $\overline{I}_a(\underline{I}_b) = 1$ and $\underline{I}_b(\underline{I}_b) < 1$, Case (iv) $\underline{I} = 1$.

The first-order derivative of the lower bound $CS^s(\underline{I}_b) - CS^n$ with respect to α is given by

$$\begin{aligned}
&\frac{\partial(CS^s(\underline{I}_b) - CS^n)}{\partial\alpha} = \\
&\tau\left(\frac{\overline{I}_a\frac{\partial\overline{I}_a}{\partial\alpha} - \overline{I}\frac{\partial\overline{I}}{\partial\alpha}}{1-\alpha} + \frac{\overline{I} - \overline{I}_a}{1-\alpha} + \frac{(\overline{I}_a - \alpha)^2 - (\overline{I} - \alpha)^2}{2(1-\alpha)^2} + 2(\overline{I}\frac{\partial\overline{I}}{\partial\alpha} - \overline{I}_a\frac{\partial\overline{I}_a}{\partial\alpha}) + \frac{\frac{\partial\overline{I}}{\partial\alpha} - \frac{\partial\overline{I}_a}{\partial\alpha}}{1-\alpha}\right) \\
&+ (1-\tau)(\underline{I}\frac{\partial\underline{I}}{\partial\alpha} - \underline{I}_b\frac{\partial\underline{I}_b}{\partial\alpha}).
\end{aligned} \tag{45}$$

By plugging the formulas for $\overline{I}_a(\underline{I}_b)$ and $\underline{I}_b(\underline{I}_b)$, we find that $CS^s(\underline{I}_b) - CS^n$ behaves in α as follows: strictly increases when Case (i) holds, strictly decreases when Case (ii) holds, and strictly increases when Case (iii) holds.

Define $\alpha_1 \equiv \inf\{\alpha|\overline{I} = 1\}$, $\alpha_2 \equiv \inf\{\alpha|\overline{I}_a(\underline{I}_b) = 1\}$, and $\alpha_3 \equiv \inf\{\alpha|\underline{I} = 1\}$. Note that since $CS^s(\underline{I}_b) - CS^n = 0$ for $\alpha = 0$, $CS^s(\underline{I}_b) - CS^n > 0$ for $\alpha = \alpha_1$, $CS^s(\underline{I}_b) - CS^n < 0$ for $\alpha = \alpha_2$, and $CS^s(\underline{I}_b) - CS^n = 0$ for $\alpha = \alpha_3$, there exists a unique threshold $\underset{\sim}{\alpha} \in (0, \alpha_3)$ such that $CS^s(\underline{I}_b) - CS^n = 0$ when $\alpha = \underset{\sim}{\alpha}$.

Next we show the existence and uniqueness of $\underline{\alpha}$. Define $\alpha_2^* = \inf\{\alpha|\overline{I}_a(I_c^*) = 1\}$. Firstly, since $CS^s(I_c^*) - CS^n \geq CS^s(\underline{I}_b) - CS^n > 0$ when $\alpha = \alpha_1$, and $CS^s(I_c^*) - CS^n < 0$ when $\alpha = \alpha_2^*$, there exists at least one $\underline{\alpha}$ such that $CS^s(I_c^*) - CS^n = -(1-\tau)\frac{\underline{I}_b(I_c^*)^2 - \underline{I}^2}{2} < 0$. Next, we show this threshold is unique.

When $I_c = I_c^*$, we rearrange Equation (43) as follows:

$$\begin{aligned}
CS^s(I_c^*) - CS^n = &- \frac{\tau}{2(1-\alpha)}(\overline{I} - \alpha)^2 + \tau\overline{I}^2 + \frac{\tau}{2(1-\alpha)}(\overline{I}_a(I_c^*) - \alpha)^2 - \tau\overline{I}_a(I_c^*)^2 \\
&- (1-\tau)\frac{\underline{I}_b(I_c^*)^2 - \underline{I}^2}{2}.
\end{aligned} \tag{46}$$

Recall that on the interval $[\alpha_1, \alpha_2^*]$, we have

$$\begin{aligned}
\overline{I} =& 1, \\
\overline{I}_a(I_c^*) =& \frac{2v_H\alpha(1-\alpha) + 2(v_H - v_L)\alpha - I_c^*\alpha(v_H - v_L)}{2(1-\alpha)v_L + \alpha(v_H - v_L)}, \\
\underline{I}_b(I_c^*) =& (2(1-\tau)v_L(1-\alpha))^{-1}\,\alpha(1-\alpha)(v_H + v_L - v_H\tau) - v_L(1-\alpha)\tau I_c^* \\
& + (1-\alpha)^{\frac{1}{2}}[(2\alpha v_L(1-\tau)(\alpha^2(v_H(2-\tau) + v_L\tau) + \tau I_c^*(v_H(2-I_c^*) + v_L I_c^*) - 2\alpha(v_H + v_L\tau I_c^*)) \\
& + (1-\alpha)(\alpha(v_H + v_L - v_H\tau) - v_L\tau I_c^*)^2]^{\frac{1}{2}}
\end{aligned} \tag{47}$$

Note that either $I_c^* > \underline{I}_b(I_c^*)$, and therefore, I_c^* satisfies the first-order condition $\frac{\partial R^C(I_c)}{\partial I_c}|_{\{I_c=I_c^*\}}=0$, or $I_c^* = \underline{I}_b(I_c^*)$. Plugging the values from Equation (47) into

$$\frac{\partial R^C(I_c)}{\partial I_c}|_{\{I_c=I_c^*\}}= \tau\left(\frac{\alpha(v_H - v_L)(I_c - \overline{I}_a(I_c^*))}{2(1-\alpha)}\right)\frac{\partial \overline{I}_a}{\partial I_c}|_{\{I_c=I_c^*\}}+(1-\tau)(v_H\alpha - v_L\underline{I}_b(I_c^*))\frac{\partial \underline{I}_b}{\partial I_c}|_{\{I_c=I_c^*\}}= 0,$$

and solving for I_c^*, we find I_c^* as a function of α. Then, one can prove that the first line in Equation (46) is strictly decreasing and convex in α on the interval $[\alpha_1, \alpha_2^*]$, and the second line in Equation (46) is strictly increasing and convex in α on the interval $[\alpha_1, \alpha_3]$. Therefore, $CS^s(I_c^*) - CS^n$ is a convex function on the interval $[\alpha_1, \alpha_2^*] \cup [\alpha_1, \alpha_3] = [\alpha_1, \alpha_2^*]$. As such, $\underline{\alpha}$ is unique.

For the L-type customer, we have

$$Q_L^n = \tau\left(\int_\alpha^{\overline{I}} \frac{I-\alpha}{1-\alpha} dF(I) + \int_{\overline{I}}^1 2IdF(I)\right) + (1-\tau)\left(\int_{\underline{I}}^{I_0} IdF(I) + \int_{I_0}^1 2IdF(I)\right),$$

and

$$Q_L^s(I_c) = \tau\left(\int_\alpha^{\overline{I}_a} \frac{I-\alpha}{1-\alpha} dF(I) + \int_{\overline{I}_a}^1 2IdF(I)\right) + (1-\tau)\left(\int_{\underline{I}_b}^{I_0} IdF(I) + \int_{I_0}^1 2IdF(I)\right).$$

Therefore, $CS^s(I_c) - CS^n = Q_L^s(I_c) - Q_L^n$, and the results for $Q_L^s(I_c) - Q_L^n$ follow from the above analysis for $CS^s(I_c) - CS^n$.

From Proposition 2.5, we know that when $\underline{I} = 1$, the model with inventory signal is equivalent to the base model, and therefore, an inventory signal does not change the customer surplus compared to the base model. Also, recall that $\underline{I}$ is strictly increasing in α when $\underline{I} < 1$. Define $\overline{\alpha} = \alpha_3 \equiv \inf\{\alpha | \underline{I} = 1\}$ to obtain the result of the proposition. □

A.2 Full-Disclosure Inventory Signal

Since the binary signaling mechanism (6) does not always improve the firm revenue, one may conjecture that a more refined signaling mechanism can benefit the firm by sending more information to the customers. To investigate this possibility, we consider an alternative signaling mechanism

$$\Sigma^r(I) = s_I \text{ for } I \in [\alpha, 1]. \tag{48}$$

With this mechanism, the firm fully discloses the inventory level to the customers, resolving the inventory uncertainty. However, customers would remain uncertain about the firm type and use the inventory and price information to update their beliefs about the possibility of future PP. The following proposition studies whether such signaling mechanism outperforms the binary signal structure studied in the main body of the paper for $\tau \leq \overline{\tau}(\alpha)$.

Proposition A.1. *When $\tau \leq \overline{\tau}(\alpha)$, the firm revenue with signaling mechanism $\Sigma^r(I)$ is no larger than the base model revenue without any inventory signal.*

Proof of Proposition A.1. Let $\underline{I}^F \equiv \min(\alpha\frac{v_H(1-\alpha)+\tau\alpha(v_H-v_L)}{v_L(1-\alpha)+\tau\alpha(v_H-v_L)}, 1)$ and $\overline{I}^F \equiv \min(\frac{v_H\alpha(1-\alpha)+(v_H-v_L)\alpha}{v_L(1-\alpha)+(v_H-v_L)\alpha}, 1)$. Since the firm fully discloses the inventory information to the customers, for any given inventory level, the first-period equilibrium prices are:

(i) U-type firm sets,

$$p_1^*(0, I) = \begin{cases} v_H - \frac{\tau(v_H-v_L)(I-\alpha)}{1-\alpha} & \text{if } I \leq \underline{I}^F \\ v_L & \text{otherwise.} \end{cases}$$

(ii) P-type firm sets,

$$p_1^*(1, I) = \begin{cases} v_H - \frac{\tau(v_H-v_L)(I-\alpha)}{1-\alpha} & \text{if } I \leq \underline{I}^F \\ v_H - \frac{(v_H-v_L)(I-\alpha)}{1-\alpha} & \text{if } \underline{I}^F < I \leq \overline{I}^F \\ v_L & \text{otherwise.} \end{cases}$$

The second-period equilibrium prices would be the same as the ones found in Proposition 2.4.

When $\tau \leq \overline{\tau}(\alpha)$, we have $\frac{v_H\alpha}{v_L} > \underline{I} = 1$. Consequently, $\overline{I}^F = 1$. Then, the expected revenue in the base model is:

$$\begin{aligned}R^B =&\tau \left(\int_\alpha^1 (v_H - \frac{\tau(v_H - v_L)}{2})\alpha dF(I) + \int_\alpha^1 (v_H\alpha + v_L(I-\alpha))dF(I)\right)\\ &+ (1-\tau)\left(\int_\alpha^1 (v_H - \frac{\tau(v_H - v_L)}{2})\alpha dF(I) + \int_\alpha^1 v_H\alpha dF(I)\right).\end{aligned}$$

Furthermore, the expected revenue with full-disclosure inventory signal is given by

$$\begin{aligned}R^F =&\tau \left(\int_\alpha^{\underline{I}^F} (v_H - \tau(v_H - v_L)\frac{I-\alpha}{1-\alpha})\alpha dF(I) + \int_{\underline{I}^F}^1 (v_H - (v_H - v_L)\frac{I-\alpha}{1-\alpha})\alpha dF(I)\right)\\ &+ \tau\left(\int_\alpha^1 (v_H\alpha + v_L(I-\alpha))dF(I)\right)\\ &+ (1-\tau)\left(\int_\alpha^{\underline{I}^F} (v_H - \tau(v_H - v_L)\frac{I-\alpha}{1-\alpha})\alpha dF(I) + \int_{\underline{I}^F}^1 v_L I dF(I) + \int_\alpha^1 v_H\alpha dF(I)\right).\end{aligned}$$

Therefore,

$$R^B - R^F = (1-\tau)(\int_{\underline{I}^F}^1 (v_H\alpha - v_L I) \geq 0,$$

which proves the proposition. □

As such, when the binary inventory signal does not improve the firm revenue (i.e., when $\tau \leq \overline{\tau}(\alpha)$), a full-disclosure signal also does not increase the revenue over the base model. As discussed in §2.4, a low value of τ corresponds to a higher likelihood of having a U-type firm. In this case, customers associate a high price with inventory scarcity rather than the firm's intention to implement PP in the future. By committing to a full-disclosure inventory signal, the firm risks skewing the customer beliefs toward PP when the inventory realization is high. The firm then must reduce the first-period price to persuade H-type customers to purchase, which hurts the firm revenue. Additionally, when inventory realizes low, an inventory signal is not much beneficial, as customers already believe the firm is U-type with a high probability when τ is small. Hence, the revenue loss from committing to a full-disclosure signal when inventory realizes high outweighs the benefit when inventory realization is low. Consequently, both firm types are a priori worse off by committing to a full-disclosure inventory signal.

The two inventory signaling mechanisms considered in this paper are the extreme cases in the space of all possible mechanisms: With the binary signal, the inventory support is divided into two intervals. With the full-disclosure signal, the firm precisely informs the customers about the state of inventory; in other words, the inventory support is divided into infinite intervals. While these two mechanisms are commonly used in practice, we do not study the "optimal" mechanism in the space of "all" possible mechanisms. In a general signaling framework, Guo and Shmaya (2019) find that a nested structure such as dividing the inventory space into multiple intervals and sending a separate signal for each interval is optimal. Arieli et al. (2019) show that when the state-space is continuous (such as inventory in our model), a bi-pooling signal is optimal where the sender divides the sate-space into two disjoint intervals and separately signals the receivers for each interval. This is similar to the binary signal studied in our paper. However, to prove the result, they require that the sender's utility be state-independent, which is not the case in our application. Nevertheless, the optimal structural design of the inventory signal is a fruitful avenue for future research.

A.3 Supporting Results for Section 2.6.1: Demand Uncertainty and Optimal Inventory

When there is no confusion, we use the same notation in this section as the main model with inventory uncertainty. For example, we use p_{it} for the price set for customer i in period t. However, one should keep in mind that the precise equilibrium values would be different with inventory uncertainty than demand uncertainty. With this note in mind, we first replicate the result regarding the customer behavior.

Lemma A.2. *In any SBPE, we have $a_{i2}(p_{i2}) = \mathbb{1}(v_i \geq p_{i2})$. Furthermore, there exists a unique threshold $\overline{p} \in (v_L, v_H)$ such that H-type customers purchase in the first period if and only if $p \leq \overline{p}$. Thus, we have*

$$a_{i1}(p_{i1}) = \begin{cases} \mathbb{1}(p_{i1} \leq v_L) & \textit{if } v_i = v_L \\ \mathbb{1}(p_{i1} \leq \overline{p}) & \textit{if } v_i = v_H. \end{cases}$$

Proof. Similar to the proof of Lemma 2.1, we know that if H-type customer i purchases at price $p_{i1} > v_L$ in period 1, she may receive personalized price $p_{i2} = v_H$ in period 2. Therefore, the incentive compatible condition to motivate H-type customers to purchase at p_{i1} is

$$v_H - p_1 \geq Pr(\omega = 1|p_1 > v_L)\xi^c_{i2}(\mathcal{H}^c_2|\omega = 1 \ \& \ p_1 > v_L)(v_H - v_L). \tag{49}$$

We assert the following claim.

Claim A1. Define $g(\alpha) \equiv v_H \min(\alpha, I) + v_L \max(I - \alpha, 0) - \max(\min(\alpha, I)v_H, v_L I)$ and $l(\alpha) \equiv v_L I - p_1 \min(\alpha, I)$. There exist unique thresholds $\underline{\alpha} = \frac{I v_L}{p_1 + v_H - v_L}$ and $\overline{\alpha} = \frac{I v_L}{p_1}$ such that for a P-type firm,

$$p_{i1}(1, \alpha) = \begin{cases} p_1 & \text{if } \alpha \geq \underline{\alpha} \\ v_L & \text{otherwise,} \end{cases}$$

and for a U-type firm,

$$p_{i1}(0, \alpha) = \begin{cases} p_1 & \text{if } \alpha \geq \overline{\alpha} \\ v_L & \text{otherwise.} \end{cases}$$

For $\alpha > I$, we have $g(\alpha) = 0$ and $l(\alpha) < 0$. For $\alpha \leq I$, $g(\alpha)$ and $l(\alpha)$ are continuous and differentiable almost everywhere, and $l'(\alpha) \geq g'(\alpha)$ for each α. Therefore, $g(\alpha) - l(\alpha)$ has the single-crossing property. The rest of the proof of Claim A1 follows similarly to the proof of Claim 1 in Lemma 2.1.

Using this result, the right-hand side of Equation (49) is given by

$$\begin{aligned} &Pr(\omega = 1|p_1 > v_L)\xi^c_{i2}(\mathcal{H}^c_2|\omega = 1 \ \& \ p_1 > v_L)(v_H - v_L) \\ &= \frac{\tau(1 - \underline{\alpha})}{\tau(1 - \underline{\alpha}) + (1 - \tau)(1 - \overline{\alpha})} \int_{\underline{\alpha}}^{I} \frac{I - \alpha}{1 - \alpha} \frac{1}{I - \underline{\alpha}} dF(I)(v_H - v_L), \end{aligned} \tag{50}$$

Since the functional form of Equation (50) is similar to Equation (30), the rest of the proof follows from same steps as in the proof of Lemma 2.1. □

Proposition A.2. *A unique SBPE to the game exists when α is uncertain. Define $\overline{\alpha} = \frac{I v_L}{\overline{p}}$ and $\underline{\alpha} = \frac{I v_L}{v_H + \overline{p} - v_L}$, where $\underline{\alpha}$ and $\overline{\alpha}$ are given in Lemma A.2. Let $p^n_1(\omega, \alpha)$ and $p^n_2(\omega, \alpha, \mathcal{H}^f_2)$ be the equilibrium first and second-period prices. Then,*

(i) A U-type firm sets prices

$$p^n_1(0, \alpha) = \begin{cases} \overline{p} & \text{if } \alpha \geq \overline{\alpha} \\ v_L & \text{otherwise,} \end{cases} \quad \text{and} \quad p^n_2(0, \alpha, \mathcal{H}^f_2) = \begin{cases} v_H & \text{if } \alpha \geq \dfrac{I v_L}{v_H} \\ v_L & \text{otherwise.} \end{cases}$$

(ii) A P-type firm sets prices

$$p_1^n(1,\alpha) = \begin{cases} \overline{p} & \textit{if } \alpha \geq \underline{\alpha} \\ v_L & \textit{otherwise,} \end{cases} \quad \textit{and} \quad p_{i2}^n(1,\alpha,\mathcal{H}_2^f) = \begin{cases} a_{i1}(\overline{p})v_H + (1-a_{i1}(\overline{p}))v_L & \textit{if } \alpha \geq \underline{\alpha} \\ v_L & \textit{otherwise.} \end{cases}$$

Proof. Similar to the proof of Proposition 2.2. □

The P-type firm implements PP only if the demand popularity, α, is sufficiently large. Furthermore, both firm types set the same equilibrium price $\overline{p}$ in the first period if they intend to sell only to the H-type customers. Therefore, an H-type customer cannot fully distinguish between the two firm types upon observing the price. However, since $\underline{\alpha} \leq \overline{\alpha}$, a price $p_1 > v_L$ is more likely to be associated with a P-type firm. Thus, price can only partially signal the firm type.

Similar to the main model, we consider a binary class of inventory signals, i.e., signaling mechanisms of the form

$$\Sigma^d(I) = \begin{cases} s_b^d & \text{if } \alpha \leq \alpha_c \\ s_a^d & \text{otherwise.} \end{cases} \tag{51}$$

Parallel to Lemma 2.2, we characterize the following result.

Lemma A.3. *In any SBPE, Customer i who receives price p_{i2} and signal $s \in \{s_a^d, s_b^d\}$ makes decision $a_{i2}(p_{i2}, s) = \mathbb{1}(v_i \geq p_{i2})$ in the second period. Additionally, there exist unique thresholds $\overline{p}_a(\alpha_c) \in (v_L, v_H)$ and $\overline{p}_b(\alpha_c) \in (v_L, v_H]$ such that H-type customers purchase the product in the first period if and only if $p_1 \leq \overline{p}_b(\alpha_c)$ when they receive signal s_b^d and $p_1 \leq \overline{p}_a(\alpha_c)$ when they receive signal s_a^d. Hence,*

$$a_{i1}(p_1, s) = \begin{cases} \mathbb{1}(p_1 \leq v_L) & \textit{if } v_i = v_L \\ \mathbb{1}(p_1 \leq \overline{p}_b(\alpha_c)) & \textit{if } v_i = v_H \textit{ and } s = s_b^d \\ \mathbb{1}(p_1 \leq \overline{p}_a(\alpha_c)) & \textit{if } v_i = v_H \textit{ and } s = s_a^d. \end{cases}$$

Furthermore, for any signaling mechanism $\Sigma^c(\alpha)$, there exist thresholds $\underline{\alpha}_b(\alpha_c)$, $\overline{\alpha}_b(\alpha_c)$,

$\underline{\alpha}_a(\alpha_c)$, and $\overline{\alpha}_a(\alpha_c)$ as functions of α_c such that

$$\overline{p}_b(\alpha_c) =$$

$$\begin{cases} v_L & \text{if } \alpha_c \in [0, \underline{\alpha}_b(\alpha_c)) \\ v_H - (1 + \frac{1-I}{\alpha_c - \underline{\alpha}_b(\alpha_c)} \cdot \log(\frac{1-\alpha_c}{1-\underline{\alpha}_b(\alpha_c)})) \cdot (v_H - v_L) & \text{if } \alpha_c \in [\underline{\alpha}_b(\alpha_c), \overline{\alpha}_b(\alpha_c)) \\ v_H - \frac{\tau(\alpha_c - \underline{\alpha}_b(\alpha_c))}{\tau(\alpha_c - \underline{\alpha}_b(\alpha_c)) + (1-\tau)(\alpha_c - \overline{\alpha}_b(\alpha_c))}(1 + \frac{1-I}{\alpha_c - \underline{\alpha}_b(\alpha_c)} \cdot \log(\frac{1-\alpha_c}{1-\underline{\alpha}_b(\alpha_c)})) \cdot (v_H - v_L) & \text{if } \alpha_c \in [\overline{\alpha}_b(\alpha_c), I) \\ v_H - \frac{\tau(\alpha_c - \underline{\alpha}_b(\alpha_c))}{\tau(\alpha_c - \underline{\alpha}_b(\alpha_c)) + (1-\tau)(\alpha_c - \overline{\alpha}_b(\alpha_c))}(1 + \frac{1-I}{I - \underline{\alpha}_b(\alpha_c)} \cdot \log(\frac{1-I}{1-\underline{\alpha}_b(\alpha_c)})) \cdot (v_H - v_L) & \text{if } \alpha_c \in [I, 1], \end{cases}$$

where $\overline{p}$ is given in Lemma A.2. Also,

$$\overline{p}_a(\alpha_c) = \begin{cases} \overline{p} & \text{if } \alpha_c \in [0, \underline{\alpha}_a(\alpha_c)) \\ v_H - \frac{\tau(1-\alpha_c)}{\tau(1-\alpha_c) + (1-\tau)(1-\overline{\alpha}_a(\alpha_c))}(1 + \frac{1-I}{I-\alpha_c} \cdot \log(\frac{1-I}{1-\alpha_c})) \cdot (v_H - v_L) & \text{if } \alpha_c \in [\underline{\alpha}_a(\alpha_c), \overline{\alpha}_a(\alpha_c)) \\ v_H - \tau(1 + \frac{1-I}{I-\alpha_c} \cdot \log(\frac{1-I}{1-\alpha_c})) \cdot (v_H - v_L) & \text{if } \alpha_c \in [\overline{\alpha}_a(\alpha_c), I) \\ v_H & \text{if } \alpha_c \in [I, 1]. \end{cases}$$

Finally, $\overline{p}_b(\alpha_c)$ and $\overline{p}_a(\alpha_c)$ are non-decreasing in α_c, and $\overline{p}_b(\alpha_c) \leq \overline{p} \leq \overline{p}_a(\alpha_c)$.

Proof. The proof follows similar to the proof of Lemma 2.2 with two main differences: First, the signal mechanism reveals information about α instead of inventory, I. Second, in contrast to the model with inventory uncertainty, α may realize in a way that $\alpha > I$. However, we note that the Low-α signal corresponds to the HI signal in the main model, and the High-α signal corresponds to the LI signal. Intuitively, a product low on inventory has the same availability as a popular product with high demand and vice versa, inducing the same availability risk as in our main model.

In period 1, when the firm sends the Low-α signal, customers update their beliefs about α to a uniform distribution supported on $[0, \alpha_c)$. The proof of the existence $\underline{\alpha}_b$ then follows from repeating the steps in the proof of Lemma 2.2 with the HI signal and defining $\underline{\alpha} = \min(\frac{Iv_L}{\overline{p}_b + v_H - v_L}, \alpha_c)$. In addition, the proof of the existence of $\overline{\alpha}_b$ follows from similar steps to the proof of the existence of $\underline{I}_a$ in the main model. Given the existence of $\underline{\alpha}_b$ and $\overline{\alpha}_b$, we can divide the space into the following four subcases (a) $\alpha_c \leq \underline{\alpha}_b$, (b) $\underline{\alpha}_b < \alpha_c \leq \overline{\alpha}_b$, (c) $\overline{\alpha}_b < \alpha_c \leq I$, and (d) $I < \alpha_c$, and repeat the analysis in the proof of Lemma 2.2 for each case.

When the firm sends the Low-α signal, one can also prove the existence and uniqueness of $\overline{p}_a(\alpha_c) \in [\overline{p}, v_H]$ following similar analysis in Lemma 2.2 for the HI inventory signal.

When the firm sends the High-α signal, the existence and uniqueness of $\overline{p}_a(\alpha_c)$ in the interval $[\overline{p}, 1]$ follows from similar analysis to the Proof of Lemma 2.2 with the LI signal and defining $\underline{\alpha} = \min(\frac{Iv_L}{\overline{p}_a + v_H - v_L}, \alpha_c)$ and $\overline{\alpha} = \min(\frac{Iv_L}{\overline{p}_a}, \alpha_c)$.

To prove the monotonicity results, similar to the proof of Lemma 2.2, we can show that price $\overline{p}_a(\alpha_c)$ is constant in α_c for $\alpha_c \in [0, \underline{\alpha}_a(\alpha_c))$ and $\alpha_c \in [I, 1]$, and monotonically increases for $\alpha_c \in [\underline{\alpha}_a(\alpha_c), \overline{\alpha}_a(\alpha_c))$ and $\alpha_c \in [\overline{\alpha}_a(\alpha_c), I)$. Furthermore, $\overline{p}_b(\alpha_c)$ is piecewise non-decreasing, and continuously transitions from Case (a) to Case (b) and from Case (c) to Case (d). Additionally, when α_c transitions from Case (b) to Case (c), $\overline{p}_b(\alpha_c)$ is either continuous or jumps up. □

Except for the logarithmic terms in $\overline{p}_a(\alpha_c)$ and $\overline{p}_b(\alpha_c)$, these equilibrium values derived from above equations possess similar functional forms as those found in Lemma 2.2. This lemma enables us to study the existence and structure of the optimal signaling mechanism. Let α_c^* be the optimal signal cut-off in (51) and $\Sigma^*(\alpha)$ be the corresponding signaling mechanism. We establish the following result.

Lemma A.4. *In any SBPE, there exist at most two local maxima, one in the interval* $[\underline{\alpha}_b, I)$ *and one at* $\alpha_c^* = I$*.*

Proof. First, when $\alpha_c \in [0, I]$, the availability risk and expected revenue with demand signal will have equivalent functional forms to the main model with inventory uncertainty for some $I^c > \alpha$. Therefore, an argument similar to the proof of Proposition 2.5 shows the existence of at most one interior optimal cut-off $\alpha_c^* \in [\underline{\alpha}_b, I)$. This cut-off would be one of the local maxima for the problem.

Next, we consider the case when $\alpha_c \in [I, 1]$. In this case, the expected revenue is given by

$$R^C(\alpha_c) = \tau[\int_0^{\underline{\alpha}_b} 2v_L I dF(\alpha) + \int_{\underline{\alpha}_b}^{I} ((\overline{p}_b(\alpha_c) + v_H - v_L)\alpha + v_L I)\, dF(\alpha) + \int_I^{\alpha_c} (\overline{p}_b(\alpha_c) + v_H)\, I dF(\alpha)$$
$$+ \int_{\alpha_c}^{1} 2v_H I dF(\alpha)] + (1-\tau)[\int_0^{\overline{\alpha}_b} (v_L I + \max(v_H \alpha, v_L I))\, dF(\alpha) + \int_{\overline{\alpha}_b}^{I} (\overline{p}_b(\alpha_c) + v_H)\alpha dF(\alpha)$$
$$+ \int_I^{\alpha_c} (\overline{p}_b(\alpha_c) + v_H)\, I dF(\alpha) + \int_{\alpha_c}^{1} 2v_H I dF(\alpha)]. \tag{52}$$

We can show that the derivative of $R^C(\alpha_c)$ with respect to α_c is negative for $\alpha_c \in [I, 1]$. Thus, $R_C(\alpha_c)$ is decreasing in this interval.

Furthermore, when $\alpha_c \in [\overline{\alpha}_b, I)$, the expected revenue is given by

$$R^C(\alpha_c) = \tau[\int_0^{\underline{\alpha}_b} 2v_L I dF(\alpha) + \int_{\underline{\alpha}_b}^{\alpha_c} ((\overline{p}_b(\alpha_c) + v_H - v_L)\alpha + v_L I)\, dF(\alpha)$$
$$+ \int_{\alpha_c}^{I} ((\overline{p}_a(\alpha_c) + v_H - v_L)\alpha + v_L I)\, dF(\alpha) + \int_I^1 (v_H + \overline{p}_a(\alpha_c)) I dF(\alpha)]$$
$$+ (1-\tau)[\int_0^{\overline{\alpha}_b} (v_L I + \max(v_H \alpha, v_L I))\, dF(\alpha) + \int_{\overline{\alpha}_b}^{\alpha_c} (\overline{p}_b(\alpha_c) + v_H)\alpha dF(\alpha)$$
$$+ \int_{\alpha_c}^{I} (\overline{p}_a(\alpha_c) + v_H)\, \alpha dF(\alpha) + \int_I^1 (\overline{p}_a(\alpha_c) + v_H) I dF(\alpha)]. \tag{53}$$

One can show that when $\alpha_c < I$, $R^C(\alpha_c)$, defined in Equation (53), is increasing in α_c. Hence, $\alpha_c = I$ is another local maximum, proving the existence of at most two local maxima. Since the expected revenue function is upper semi-continuous with at most one jump up, the global maximum exists and must be one of the two local maxima. □

Figure 2.4 in the main body of the paper visualizes Lemma A.4.

A.4 Supporting Results for Section 2.6.2: Customer Non-Stationary Valuations and a $T-$Period Model

Given the model setup discussed in §2.6.2, the customer utility is given by

$$u_{it}(p_{i1}, p_{i2}, \mathcal{H}_{it}^c) = \mathbb{E}_{I,\omega}\left[\sum_{j=t}^{T} a_{ij}\,(p_{ij}, s)\, \xi_{ij}^c(\mathcal{H}_{ij}^c)(v_i - p_{ij}) \middle| \mathcal{H}_{it}^c\right], \tag{54}$$

and customers solve

$$\max_{a_{it}\in\{0,1\}} u_{it}(p_{i1}, p_{i2}, \mathcal{H}_{it}^c) \text{ for } t \in \{1, ..., T\}. \tag{55}$$

Similarly, the firm revenue-to-go function at the beginning of period t is given by

$$R_t(\mathcal{H}_t^f, p_t) = \mathbb{E}_{I,\omega}\left[\sum_{j=t}^{T} \int_{i\in M} a_{ij}\left(p_{ij}, \Sigma(I)\right) \xi_{ij}^f(\mathcal{H}_j^f) p_{ij}\, di \middle| \mathcal{H}_t^f\right], \tag{56}$$

and the firm solves

$$\begin{aligned} \max_{p_t} \quad & R_t(\mathcal{H}_t^f, p_t) \\ \text{s.t.} \quad & p_1(\omega, I) \in \arg\max_{p_1(\omega,I)} R_1(\mathcal{H}_t^f, p_t), \\ & (1-\omega)(p_{it} - p_{lt}) = 0,\ \forall i, l \in M,\ \omega \in \Omega,\ \&\ t \in \{1, 2, ..., T\}. \end{aligned} \tag{57}$$

First, we replicate the result of Proposition 2.1 with myopic customers.

Proposition A.3. *A unique SBPE to the game with myopic customers exists. Let* $p_t^m(\omega, I, \mathcal{H}_t^f)$ *be the equilibrium prices in period* t *with myopic customers. Then,*

(i) A U-type firm sets prices

$$p_t^m(0, I, \mathcal{H}_t^f) = \begin{cases} v_H & \text{if } I \leq \dfrac{\alpha\lambda v_H}{v_L} \\ v_L & \text{otherwise.} \end{cases}$$

(ii) A unique threshold t^m *exists such that, for* $t \leq t^m$ *(possibly* $t^m = 0$*), a P-type firm sets prices* $p_t^m(1, I, \mathcal{H}_t^f) = v_H$*. Furthermore, for* $t > t^m$*, it sets prices*

$$p_t^m(1, I, \mathcal{H}_t^f) = \max_{k<t^m}\{a_{ik}(p_k^m)\} v_H + \left(1 - \max_{k<t^m}\{a_{ik}(p_k^m)\}\right) v_L.$$

Proof. Part (i). Myopic customers purchase the product if they receive a non-negative utility. In this case, a U-type firm either sets price v_H to all customers and obtains revenue $\alpha\lambda v_H$ in each period, or sets price v_L and receives the revenue $I v_L$ in each period. Hence, the optimal pricing strategy for the U-type firm follows from solving $\max(\alpha\lambda v_H, I v_L)$. This completes the proof of Part (i).

Part (ii). Since $\lambda v_H > v_L$, it would be optimal to set price v_H for the rest of the periods to the identified H-type customers. Therefore, the P-type firm either charges v_H to all customers, or charges v_H to "identified" H-type customers and v_L to the remaining customers.

Next, we show that a P-type firm is always better-off to explore in the earlier stages than later. Given any period t, suppose a P-type firm has explored for s periods, where $s \leq t$. Then, the size of the identified H-type customers is $\alpha\lambda \sum_{j=0}^{s-1}(1-\lambda)^j$. Moreover, the size of the unidentified H-type customers is given by $\alpha(1-\lambda)^s$. Finally, a fraction $1-\alpha$ of customers are L-type. We have

$$\begin{aligned}\max_{p_j, j\in\{t,\ldots,T\}} R_t(\mathcal{H}_t^f, p_t) \geq & v_H\alpha\lambda(1-\lambda)^s + v_H\alpha\lambda(1-(1-\lambda)^s) \\ & + \alpha\lambda(1-(1-\lambda)^{s+1})v_H + (I - \alpha\lambda(1-(1-\lambda)^{s+1}))v_L \\ & + \max_{p_j, j\in\{t+2,\ldots,T\}} R_{t+2}(\mathcal{H}_{t+2}^f, p_{t+2}) \\ \geq & \alpha\lambda(1-(1-\lambda)^s)v_H + (I - \alpha\lambda(1-(1-\lambda)^s))v_L \\ & + v_H\alpha\lambda(1-\lambda)^s + \alpha\lambda(1-(1-\lambda)^s)v_H \\ & + \max_{p_j, j\in\{t+2,\ldots,T\}} R_{t+2}(\mathcal{H}_{t+2}^f, p_{t+2}).\end{aligned}$$

The term after the first inequality corresponds to the case when the P-type firm explores in period t and exploits in period $t+1$. In contrast, the term after the second inequality corresponds to the case when the firm exploits in period t and explores in period $t+1$. Note that the information sets in period $t+2$ would be identical for both cases. The second inequality followed because $\lambda v_H > v_L$.

Next, we show that if the firm stops exploration in a period, then it would never be optimal to explore after that period. Since a fraction $\alpha(1-(1-\lambda)^s)$ of customers are identified as H-type, the remaining inventory to be divided among the remaining customers is $I - \alpha\lambda(1-(1-\lambda)^s)$. Suppose the firm stops exploring in period t_s. For this strategy to be optimal, we must have

$$\begin{aligned}& (T-t_s)v_L \min\left(I - \alpha\lambda((1-(1-\lambda)^{t_s}), \alpha(1-\lambda)^s + (1-\alpha)\right) \\ \geq & v_H\alpha(1-\lambda)^{t_s}\lambda + v_H\alpha\lambda^2(1-\lambda)^{t_s}(T-s_t-1) \\ & + \min\left(I - \alpha\lambda(1-(1-\lambda)^{t_s+1}), \alpha(1-\lambda)^{t_s+1} + (1-\alpha)v_L(T-t_s-1)\right).\end{aligned}$$

Consequently, we must have $(1-\alpha)v_L \geq (\lambda(T-t_s-1)+1)\alpha(1-\lambda)^{t_s}(\lambda v_H - v_L)$ or $v_L(I-\alpha\lambda) \geq (\lambda(T-t_s-1)+1)\alpha\lambda(1-\lambda)^{t_s}(v_H - v_L)$. Note that the right-hand side of both inequalities are decreasing in t_s while the left-hand side is constant. As such, it would never be optimal to explore again for any $t \geq t_s$. This completes the proof of Part (ii). □

The P-type firm first explores customer types early on during the selling season by charging price v_H to all customers. This enables the firm to identify some of the customers who are H-type. However, since some of the H-type customers may not have possessed valuation v_H until period t^m, the firm cannot identify all H-type customers. After period t^m, the firm exploits its knowledge about the customers by personalizing the prices.

Also, whether the firm benefits from PP depends on the demand popularity realization. Particularly, if the popularity realization is low, the firm would not benefit from PP, and we would have $t^m = 0$.

We next extend the results of Lemma 2.1 and Proposition 2.2 in the main body to any arbitrary finite number of selling periods with strategic customers.

Lemma A.5. *In any SBPE in the model with T periods, we have $a_{iT}(p_{iT}) = \mathbb{1}(v_{it} \geq p_{iT})$. Furthermore, there exists a unique sequence of thresholds $\overline{p}_t \in [v_L, v_H)$ such that unidentified H-type customers purchase in period t if and only if $p \leq \overline{p}_t$. Thus, we have*

$$a_{it}(p_{it}) = \begin{cases} \mathbb{1}(p_{it} \leq v_L) & \text{if } v_{it} = v_L \\ \mathbb{1}(p_{it} \leq \overline{p}_t) & \text{if } v_{it} = v_H. \end{cases}$$

Proof. In any period t, an unidentified H-type customer with valuation v_H can purchase the product at price $\overline{p}_t > v_L$ to receive a positive utility in period t at the cost of receiving personalized prices v_H for the remaining periods if the firm is P-type. Alternatively, the customer may forgo the purchase and wait until the next period she will have valuation v_H for a similar decision. Similar to the proof of Lemma 2.1, The monotonic indifference conditions in each period guarantee the existence and uniqueness of $\overline{p}_t$. □

Intuitively, only H-type customers who possess a high valuation v_H face the problem of deciding to reveal their types. If such a customer chooses to hide her type, she only benefits during future periods when she will have a high valuation. Thus, the value of the hiding type in period t is bounded above by $\lambda(T-t)(v_H - v_L)$. However, the customer can learn more about the firm type by hiding her type for a longer period.

We assert the following lemma to be used for characterizing the SBPE of the game. To facilitate the rest of the analyses, we assume that $I < (1-\alpha)+\alpha\lambda$. Additionally, with abuse of notation, we let p_t represent the price charged to the unidentified customers.

Lemma A.6. *Let t_1 and t_2 be two arbitrary periods with $t_1 < t_2$. Also, let $p^*_{it_1}(1, I)$ and $p^*_{it_2}(1, I)$ be the corresponding equilibrium price mappings set for unidentified Customer i when the firm is P-type. If $p^*_{it_2}(1, I) \in (v_L, v_H)$ and $p^*_{it_1}(1, I) \in (v_L, v_H)$, then for any $t_1 \leq t \leq t_2$, we have $p^*_{it}(1, I) \in (v_L, v_H)$.*

Proof. We show that once the firm starts exploring, it would never be optimal to pause the exploration and restart it later. To show this, it is sufficient to show that a one-period pause would not be optimal. For any period $t \in [t_1, t_2]$, suppose the P-type firm has explored the customer types for s periods, where $s \leq t$. Note that a firm intending to explore the high-valuation customers sets the prices such that customers with valuation v_H would purchase the product. Thus, the size of the identified H-type customers is given by $\alpha\lambda\sum_{j=0}^{s-1}(1-\lambda)^j$. Moreover, the size of the unidentified H-type customers is given by $\alpha(1-\lambda)^s$. Suppose $\overline{p}_t$ be the thresholds proved in Lemma A.5 for the case when the firm first explores and then exploits. Also, we use superscript $'$ for parameters corresponding to the case when the firm first exploits and then explores. For example, $\overline{p}'_t$ are the price thresholds when the firm first exploits and then explores. We have

$$
\begin{aligned}
\max_{p_j, j\in\{t,\ldots,T\}} R_t(\mathcal{H}^f_t, p_t) \geq & \overline{p}_t(\mathcal{H}^f_t)\alpha\lambda(1-\lambda)^s + \alpha\lambda(1-(1-\lambda)^s)v_H \\
& + \alpha\lambda(1-(1-\lambda)^{s+1})v_H + (I - \alpha\lambda(1-(1-\lambda)^{s+1}))v_L \\
& + \max_{p_j, j\in\{t+2,\ldots,T\}} R_{t+2}(\mathcal{H}^f_{t+2}, p_{t+2}|p_t, p_{t+1}) \\
\geq & \alpha\lambda(1-(1-\lambda)^s)v_H + (I - \alpha\lambda(1-(1-\lambda)^s))v_L \\
& + \overline{p}'_{t+1}(\mathcal{H}^f_{t+1}|p_t = v_L)\alpha\lambda(1-\lambda)^s + \alpha\lambda(1-(1-\lambda)^s)v_H \\
& + \max_{p_j, j\in\{t+2,\ldots,T\}} R_{t+2}(\mathcal{H}^{f\prime}_{t+2}, p_{t+2}|p'_t, p'_{t+1})
\end{aligned}
$$

The term after the first inequality represents the revenue-to-go under the strategy that $p_t = \overline{p}_t$ and $p_{t+1} = v_L$, i.e., exploration in period t and exploitation in period $t+1$. The term after the second inequality represents the revenue-to-go under the strategy that $p_t = v_L$ and $p_{t+1} = \overline{p}'_{t+1}$, i.e., exploitation in period t and exploration in period $t+1$. The second inequality follows if and only if $\overline{p}_t + \lambda(v_H - v_L) \geq \overline{p}'_{t+1}$. We prove that this inequality holds. Parameter $\overline{p}_t$ is the price that compensates H-type customers to reveal their types for the remaining $T - t$ periods. Also, $\overline{p}'_{t+1}$ is the price that motivates the same H-type customers to reveal their types for the remaining $T - t - 1$ periods. These prices are influenced by three factors: (1) posterior belief about firm's type, (2) posterior belief about inventory availability, and (3) the number of the remaining periods. When comparing $\overline{p}_t$ and $\overline{p}'_{t+1}$, the first two factors remain the same. However, the number of the remaining periods differ. When there are more periods left, the customers should receive a lower price to reveal their types because they may be targeted by personalized prices for a longer period. An H-type customer's potential benefit from forgoing a purchase for one additional period is bounded above by $\lambda(v_H - v_L)$. This is because, to benefit from forgoing a purchase, the H-type customer's valuation must realize to be v_H, in which case she would gain extra utility $v_H - v_L$. This extra utility, however, realizes only if the customer is assigned a unit. As such, $\overline{p}_t + \lambda(v_H - v_L) \geq \overline{p}'_{t+1}$, which is what we desired to show. □

Lemma A.6 states that if a P-type firm benefits from further exploring the unidentified H-type customers in a given period by charging a price higher than v_L, it also benefits from doing so in the earlier periods once exploration starts. With the above lemma, we characterize the game's equilibrium in the next proposition.

Proposition A.4. *A unique SBPE to the game with strategic customers exists. Let $p^s_t(\omega, I, \mathcal{H}^f_t)$ be the equilibrium prices in period t with strategic customers. Unique thresholds t^s_1 and t^s_2 exist such that*

(i) U-type firm sets

$$p^s_t(0, I, \mathcal{H}^f_t) = \begin{cases} \overline{p}_t & \text{if } I \leq \dfrac{\alpha\lambda\overline{p}_t}{v_L} \text{ and } t \geq t^s_1 \\ v_L & \text{otherwise.} \end{cases}$$

(ii) P-type firm's pricing strategy spans into three stages as follows:

Dormancy stage: *For $t < t^s_1$, $p^s_t(0, I, \mathcal{H}^f_t) = v_L$.*

Exploration stage: *For $t_1^s \le t < t_2^s$,*

$$p_t^s(1, I, \mathcal{H}_t^f) = \max_{k \le t}\{a_{ik}(p_k^m)\}v_H + \left(1 - \max_{k \le t}\{a_{ik}(p_k^m)\}\right)\overline{p}_t.$$

Exploitation Stage: *For $t_2^s \le t$,*

$$p_t^s(1, I, \mathcal{H}_t^f) = \max_{k \le t_2^s}\{a_{ik}(p_k^m)\}v_H + \left(1 - \max_{k \le t_2^s}\{a_{ik}(p_k^m)\}\right)v_L.$$

Proof. Part (i). Note that the U-type firm does not face the exploration-exploitation trade-off. Given Lemma A.5, in equilibrium, there exists a unique price path $\overline{p}_t$ such that any price higher than $\overline{p}_t$ will be rejected by the customers. Therefore, the U-type firm optimizes its revenue by choosing between prices $\overline{p}_t$ and v_L.

Part (ii). We first show the existence and uniqueness of t_2^s. We argue that the exploration stage must be continuous, i.e., if the P-type firm stops exploring in any period, then it would not be optimal to restart exploration later. To show this, suppose the exploration stops at period t_s after $\gamma \le t_s$ periods of exploration. Then, similar to the argument in the proof of Proposition A.3, we must have

$$\begin{aligned}
&(T - t_s)v_L \min\left(I - \alpha\lambda(1-(1-\lambda)^\gamma), \alpha(1-\lambda)^\gamma + (1-\alpha)\right) \\
\ge&\overline{p}_{t_s}\alpha(1-\lambda)^\gamma\lambda + v_H\alpha\lambda^2(1-\lambda)^\gamma(T - t_s - 1) \\
&+ \min\left(I - \alpha\lambda(1-(1-\lambda)^{\gamma+1}), \alpha(1-\lambda)^{\gamma+1} + (1-\alpha)v_L(T - t_s - 1)\right).
\end{aligned}$$

Since the right-hand side of the inequality is decreasing in t_s, if the P-type firm stops exploration in period t_s, it would never be optimal to explore in any period $t > t_s$. Furthermore, because it is optimal for the P-type firm to exploit in period T, we have $t_s \le T - 1$. This argument shows the existence and uniqueness of t_2^s.

For the existence and uniqueness of t_1^s, we first show that price thresholds $\overline{p}_t$ are non-decreasing in t for any $t < t_2^s$. For $t = t_2^s - 1$, since the firm stops exploring before t_2^s, we have $\overline{p}_{t_2^s-1} = v_H - (T - t_2^s + 1)Pr(\omega = 1|\mathcal{H}_{t_2^s-1}^c)\xi_{t_2^s-1}^c(\mathcal{H}_{t_2^s-1}^c|\omega = 1)\lambda(v_H - v_L)$. Note that $\mathcal{H}_{t_2^s-1}^c$ is the information of a customer with valuation v_H in period $t_2^s - 1$. Additionally, let $\mathcal{H}_{t_2^s-1}^{c\prime}$ be the information set of an H-type customer who had valuation v_H in period $t_2^s - 2$ but did not purchase and has valuation v_L in Period $t_2^s - 1$. Then, $\overline{p}_{t_2^s-2}$ satisfies the following equation.

$$v_H - \overline{p}_{t_2^s-2} \geq \mathbb{E}\left[\lambda(v_H - p_{t_2^s-1}) + (1-\lambda)(T - t_2^s + 1)Pr(\omega = 1|\mathcal{H}^{c\prime}_{t_2^s-1})\xi^c_{t_2^s-1}(\mathcal{H}^{c\prime}_{t_2^s-1}|\omega = 1)\lambda(v_H - v_L)\right]. \quad (58)$$

Note that the term on the left-hand side of the inequality is the utility the customer would receive if she purchases the product at a price higher than v_L. In this case, for the rest of the periods, she would receive a zero utility. The right-hand side corresponds to the utility the customer would receive from forgoing a purchase. Thus, price $\overline{p}_{t_2^s-2}$ should be set in a way that the customer has the incentive to reveal her type.

This, combined with that $\overline{p}_{t_2^s-1} = v_H - (T - t_2^s + 1)Pr(\omega = 1|\mathcal{H}^c_{t_2^s-1})\xi^c_{t_2^s-1}(\mathcal{H}^c_{t_2^s-1}|\omega = 1)\lambda(v_H - v_L)$ concludes that $\overline{p}_{t_2^s-2} \leq \overline{p}_{t_2^s-1}$. One can inductively extend this result to any $t < t_2^s$. We then define t_1^s as the maximum value such that for $t < t_1^s$, we have $\overline{p}_t < v_L$. In other words, before period t_1^s, price $\overline{p}_t$ does not exist such that H-type customers would purchase the product at this price. Hence, the P-type firm must set price v_L for all customers to receive a positive revenue. This completes the proof of Part (ii). □

Appendix B Chapter 3

B.1 Proofs from Chapter 3

B.1.1 Proof of Proposition 3.1

Part (i) We initially assume $p = 0$ to guarantee the price would never interrupt the customer learning. Then, we define $\overline{p}$ to be the lowest customer belief that may arise over time.

Suppose $p = 0$. Since Q has a bounded support on $[0, 1]$, $Q_t^c \geq 0$. Therefore, $Z_t = 1$ for all t. We show that the support of customer beliefs over time, i.e., $Supp(\tilde{Q}_t)$, form a set of non-empty contracting intervals $[\underline{Q}_t, \overline{Q}_t)$.

Suppose $Supp(\tilde{Q}_t) = [\underline{Q}_t, \overline{Q}_t)$. If $r_t = i$, then the lower bound of Customer $t+1$'s belief is

$$\underline{Q}_{t+1} = \max\{\underline{Q}_t, \lambda_i + Q_t^c\}, \tag{59}$$

and its upper bound is given by

$$\overline{Q}_{t+1} = \min\{\overline{Q}_t, \lambda_{i+1} + Q_t^c\}. \tag{60}$$

Since $\tilde{Q}$ has a continuous and strictly increasing distribution, $Supp(\tilde{Q}_{t+1}) = [\underline{Q}_{t+1}, \overline{Q}_{t+1})$. Hence, $Supp(\tilde{Q}_{t+1}) \subseteq Supp(\tilde{Q}_t)$. Also, note that, when $\alpha = 1$, the customer rates $r_t = i$ if $Q - Q_t^c \in [\lambda_i, \lambda_{i+1})$. Thus, $Q \in Supp(\tilde{Q}_t)$.

We claim $Supp(\tilde{Q}_\infty)$ exists and $\underline{Q}_\infty = \overline{Q}_\infty$. If $\underline{Q}_t = \overline{Q}_t$ for some t, then the proof is complete. Suppose $\underline{Q}_t < \overline{Q}_t$. In this case, $Q_t^c \in (\underline{Q}_t, \overline{Q}_t)$. If $r_t = i > 0$, from (59), $\underline{Q}_{t+1} > \underline{Q}_t$. Alternatively, if $r_t = i < 0$, from (60), we have $\overline{Q}_{t+1} < \overline{Q}_t$. Therefore, $\text{Supp}(\tilde{Q}_{t+1}) \subset \text{Supp}(\tilde{Q}_t)$. Since $Supp(\tilde{Q}) = [0, 1]$, we can inductively conclude the support of the beliefs remain intervals that contract over time. As such, there exists a limiting interval for the support of the beliefs. In other words, $Supp(\tilde{Q}_\infty) = [\underline{Q}_\infty, \overline{Q}_\infty]$. If $\underline{Q}_\infty = \overline{Q}_\infty$, the proof is complete because we argued earlier that both Q and Q_t^c must belong to the support of the beliefs. Suppose $\overline{Q}_\infty > \underline{Q}_\infty$. Recall that $\overline{Q}_\infty = \min\{\overline{Q}_\infty, \lambda_{i+1} + Q^c_\infty\}$ and $\underline{Q}_\infty = \max\{\underline{Q}_\infty, \lambda_i + Q^c_\infty\}$. If $r_\infty > 0$, then $\lambda_i \geq 0$. Consequently, $\lambda_i + Q^c_\infty > \underline{Q}_\infty$, which is a contradiction with $\underline{Q}_\infty = \max\{\underline{Q}_\infty, \lambda_i + Q^c_\infty\}$. Similarly, one can derive a contradiction if $r_\infty < 0$. Hence, $\overline{Q}_\infty > \underline{Q}_\infty$ cannot hold and we must have $\overline{Q}_\infty = \underline{Q}_\infty$, which proves Part (i). **Part (ii)** define $\overline{p} = \min_t\{Q^c_t\}$. Then, if $p \leq \overline{p}$, all customers continue purchasing and learning dynamics are the same as in the case with $p = 0$. However, when $p > \overline{p}$, customers stop purchasing, which interrupts learning. □

B.1.2 Proof of Proposition 3.2

First, we prove the convergence of the customer beliefs for all α. Suppose the intrinsic quality is Q. Let $\kappa(r; Q, \tilde{Q})$ be the probability that the customer rates the product at r given the intrinsic quality Q and customer belief $\tilde{Q}$. This probability can be written as

$$\kappa(r = i; Q, \tilde{Q}_t) = \mathbb{P}_\theta[\lambda_i \leq Q - \alpha Q^c_t - (1-\alpha)(p_t - \theta_t) < \lambda_{i+1}]. \tag{61}$$

Given r_{t-1}, Customer t can eliminate any Q' such that $Q' - \alpha Q^c_{t-1} + (1-\alpha)\overline{\theta} - (1-\alpha)p < \lambda_{r_{t-1}}$, or $Q' - \alpha Q^c_{t-1} - (1-\alpha)\overline{\theta} - (1-\alpha)p \geq \lambda_{r_{t-1}+1}$. Therefore,

$$Supp(\tilde{Q}_t|\mathcal{I}^c_t) = \\ \big[\alpha Q^c_{t-1} - (1-\alpha)\overline{\theta} + (1-\alpha)p + \lambda_{r_{t-1}}, \alpha Q^c_{t-1} + (1-\alpha)\overline{\theta} + (1-\alpha)p + \lambda_{r_{t-1}+1}\big) \bigcap Supp(\tilde{Q}_{t-1}|\mathcal{I}^c_{t-1}). \tag{62}$$

Consequently, $Supp(\tilde{Q}_t|\mathcal{I}^c_t) \subseteq Supp(\tilde{Q}_{t-1}|\mathcal{I}^c_{t-1})$. Additionally, $Supp(\tilde{Q}_\infty) = \bigcap_{t=1}^\infty Supp(\tilde{Q}_t|\mathcal{I}^c_t)$. This implies that $\kappa(r_t; Q', \tilde{Q}_t) > 0$ for all t and $Q' \in Supp(\tilde{Q}_\infty)$. Let $R_\infty = \{r : \kappa(r; Q, \tilde{Q}_\infty) > 0\}$.

In a general social learning context, Theorem 3 of Frick et al. (2022) proves the convergence of the beliefs with mild assumptions on the continuity of the actions under the potentially misspecified beliefs. In our context, their assumption on belief continuity translates to the continuity of $\kappa(r; Q, \tilde{Q})$, $\kappa(r; Q', \tilde{Q})/\kappa(r; Q, \tilde{Q})$, and $\kappa(r; Q', \tilde{Q})$ in $\tilde{Q}$ for intrinsic quality Q, and arbitrary values $Q' \in Supp(\tilde{Q}_\infty)$ and $r \in R_\infty$. In this context, an arbitrary deterministic function $g(\cdot)$ is said to be continuous in random variable $\tilde{Q}$ if for any sequence of random variables $\{\tilde{Q}_n\}$ that converges in distribution to $\tilde{Q}$, the sequence $\{g(\tilde{Q}_n)\}$ also converges in distribution to $g(\tilde{Q})$.

Note that, as argued above, the denominator for the second term, i.e., $\kappa(r; Q, \tilde{Q})$ is positive. Hence, it is sufficient to show the continuity of $\kappa(r; Q, \tilde{Q})$ in $\tilde{Q}$, which follows from the definition of $\kappa(r; Q, \tilde{Q})$ and that $Q_t^c = \mathbb{E}(\tilde{Q}_t)$. Hence, we can use Theorem 3 of Frick et al. (2022) to show the convergence of the customer beliefs. Next, we prove Parts (i)-(iii) of the proposition.

Part (i) We show that if the converged belief $\tilde{Q}_\infty$ exists and α is sufficiently small, then $\mathbb{P}(\tilde{Q}_\infty = Q) = 1$.

Suppose $Supp(\tilde{Q}_\infty)$ contains more than one element, namely Q and Q' such that $Q \neq Q'$. In other words, Q and Q' cannot be separated from each other by customers through time. Note that if the heterogeneity is rich enough, then there exists some rating r such that $\kappa(r; Q, \tilde{Q}_\infty) \neq \kappa(r; Q', \tilde{Q}_\infty)$, in which case Q and Q' can be separated over time. We show that when α is small, such r exists.

If $\kappa(r; Q, \tilde{Q}_\infty) \neq \kappa(r; Q', \tilde{Q}_\infty)$ for some r, there is nothing to prove. Also, since $F_\theta(\cdot)$ is strictly monotonic, then $\kappa(r; Q, \tilde{Q}_\infty) \neq \kappa(r; Q', \tilde{Q}_\infty)$ for $Q \neq Q'$ except when $\kappa(r; Q, \tilde{Q}_\infty) = \kappa(r; Q', \tilde{Q}_\infty) = 1$. If $\kappa(r; Q, \tilde{Q}_\infty) = \kappa(r; Q', \tilde{Q}_\infty) = 1$, only "one" of the ratings can arise. We show this cannot be the case by finding conditions such that at least two ratings arise.

We consider the cases of $k = 1$ and $k \geq 2$ separately. When $k = 1$, $R = \{-1, 1\}$, and customers rate the product positively if and only if $u^r(\alpha, \theta, Q^c) \geq 0$. A sufficient condition for both ratings to arise is

$$\begin{aligned} \min_\theta\{\max_{Q,Q^c} Q - \alpha Q^c + (1-\alpha)(\theta - p)\} &= 1 - 0 + (1-\alpha)(-\overline{\theta} - p) \leq 0, \\ \text{and } \max_\theta\{\min_{Q,Q^c} Q - \alpha Q^c + (1-\alpha)(\theta - p)\} &= 0 - \alpha + (1-\alpha)(\overline{\theta} - p) \geq 0, \end{aligned} \tag{63}$$

which holds if $\alpha \leq \max\{1 - \frac{1}{1+\overline{\theta}-p}, 1 - \frac{1}{\overline{\theta}+p}\}$. In this case, define $\overline{\alpha} \equiv \max\{1 - \frac{1}{1+\overline{\theta}-p}, 1 - \frac{1}{\overline{\theta}+p}\}$, which completes the proof for the case with $k = 1$.

When $k \geq 2$, if $\max_\theta u^r(\alpha, \theta, Q^c) - \min_\theta u^r(\alpha, \theta, Q^c) > \lambda_{i+1} - \lambda_i = \frac{1}{k-1}$, then at least two ratings will arise. From the definition of $u^r(\alpha, \theta, Q^c)$, we have $\max_\theta u^r(\alpha, \theta, Q^c) - \min_\theta u^r(\alpha, \theta, Q^c) = 2(1-\alpha)\overline{\theta}$. Consequently, it is sufficient to have $\alpha \leq 1 - \frac{1}{2\overline{\theta}(k-1)}$. In this case, denote $\overline{\alpha} \equiv 1 - \frac{1}{2\overline{\theta}(k-1)}$ to prove the case of $k \geq 2$.

To complete the proof for Part (i), define

$$\overline{\alpha}(k,\overline{\theta}) = \begin{cases} \max\{1-\frac{1}{1+\overline{\theta}-p}, 1-\frac{1}{\overline{\theta}+p}\} & \text{if } k=1 \\ 1-\frac{1}{2\overline{\theta}(k-1)} & \text{if } k\geq 2. \end{cases} \tag{64}$$

Part (ii) Note that Q has a continuous support. Hence, when $\alpha \geq \overline{\alpha}$, from the argument of Part (i), there exists some Q' close to Q that cannot be separated from Q. Specifically, this occurs when for some i, $\min_\theta u^r(\alpha,\theta,Q^c) \geq \lambda_i$ and $\max_\theta u^r(\alpha,\theta,Q^c) < \lambda_{i+1}$ for both Q and Q', i.e., when $|Q-Q'| \leq 2(1-\alpha)\overline{\theta}$. Any Q' outside this neighborhood of Q can be separated from the intrinsic quality. This implies that the converged belief has support with the length of at most $2(1-\alpha)\overline{\theta}$. This also shows that $\lim_{\alpha\to 1} Q_\infty(\alpha) = Q$.

Part (iii) Follows from the definition of $\overline{\alpha}(k,\overline{\theta})$ in Equation (64). □

B.1.3 Proof of Proposition 3.3

The proof follows from the proof of Proposition 3.2, where it is shown that, when $k \geq 2$, customer beliefs converge to the intrinsic quality if $\alpha \leq 1-\frac{1}{2\overline{\theta}(k-1)}$. We can then rewrite the inequality with respect to k and obtain $k \geq 1+\frac{1}{2\overline{\theta}(1-\alpha)}$. □

B.1.4 Proof of Lemma 3.1

We prove the result by induction. For $t=1$, the rating frequencies are of the form $\pi_1^i = e_i$. It is then straightforward to see the result.

Suppose these properties hold for Period $t-1$. Consider any two rating frequencies of the form $\pi_t^i = \pi + e_i$ and $\pi_t^j = \pi + e_j$ at time t for $j > i$ and arbitrary frequency π. We first show that for any rating history leading to π_t^i, there exists a rating history leading to π_t^j that induces a higher expected customer belief.

Note that π_t^i can be constructed from two rating frequencies in Period $t-1$: $\pi_{t-1} = \pi$ and $\pi_{t-1}^i = \pi + e_i - e_m$, where m represents the index of any non-zero element in π_t^i other than i. Similarly, π_t^j can only arise from $\pi_{t-1} = \pi$ and $\pi_{t-1}^j = \pi + e_j - e_m$. Since π_t^i and π_t^j differ only in i and j, only the following two cases are feasible: (i) π_t^i and π_t^j share the same history path until period $t-1$, i.e, $\pi_{t-1}^i = \pi_{t-1}^j = \pi$; or (ii) they are different before period $t-1$, i.e., $\pi_{t-1}^i \neq \pi_{t-1}^j$.

Case (i). Rating frequencies π_t^i and π_t^j arise from π by $r_t = i$ and $r_t = j$, respectively. Thus, when $\alpha = 1$,

$$Supp(\tilde{Q}|\pi_t = \pi + e_i \ \& \ \pi_{t-1} = \pi) = \left[\lambda_i + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi\right], \lambda_{i+1} + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi\right]\right) \cap Supp(\tilde{Q}|\pi_{t-1} = \pi),$$

and

$$Supp(\tilde{Q}|\pi_t = \pi + e_j \ \& \ \pi_{t-1} = \pi) = \left[\lambda_j + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi\right], \lambda_{j+1} + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi\right]\right) \cap Supp(\tilde{Q}|\pi_{t-1} = \pi).$$

Since $\lambda_j \geq \lambda_{i+1}$, the equations above imply that

$$\mathbb{E}\left[\tilde{Q}|\pi_t = \pi + e_j \ \& \ \pi_{t-1} = \pi\right] > \mathbb{E}\left[\tilde{Q}|\pi_t = \pi + e_i \ \& \ \pi_{t-1} = \pi\right].$$

Case (ii). In this case, for each π_t^i that arises from $\pi_{t-1}^i = \pi + e_i - e_m$ for each m, there exists π_t^j that arises from $\pi_{t-1}^j = \pi + e_j - e_m$. We have

$$\begin{aligned} &Supp(\tilde{Q}|\pi_t = \pi + e_i \ \& \ \pi_{t-1} = \pi + e_i - e_m) \\ &= \left[\lambda_m + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi + e_i - e_m\right], \lambda_{m+1} + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi + e_i - e_m\right]\right) \cap Supp(\tilde{Q}|\pi_{t-1} = \pi + e_i - e_m), \end{aligned} \tag{65}$$

and

$$\begin{aligned} &Supp(\tilde{Q}|\pi_t = \pi + e_j \ \& \ \pi_{t-1} = \pi + e_j - e_m) \\ &= \left[\lambda_m + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi + e_j - e_m\right], \lambda_{m+1} + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi + e_j - e_m\right]\right) \cap Supp(\tilde{Q}|\pi_{t-1} = \pi + e_j - e_m). \end{aligned} \tag{66}$$

From the induction hypothesis, the separation and monotonicity properties hold for period $t-1$. Therefore, we have $Supp(\tilde{Q}|\pi_{t-1} = \pi + e_i - e_m) \cap Supp(\tilde{Q}|\pi_{t-1} = \pi + e_j - e_m) = \emptyset$ and $\mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi + e_i - e_m\right] < \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi + e_j - e_m\right]$. Combining this observation with Equations (65) and (66), we establish the separation and monotonicity properties for Period t.

To see the partition property, we need to show that $\bigcup_{i=1}^{n} Supp\left(\tilde{Q}|\pi_t^i\right) = \mathcal{Q}$. To prove this, we first argue that $Supp\left(\tilde{Q}|\pi_t^i\right)$ contains intrinsic quality Q. From the induction hypothesis, we have $Q \in Supp\left(\tilde{Q}|\pi^i_{t-1}\right)$. Additionally, by definition of the customer rating behavior, the intrinsic quality must belong to

$$\left[\lambda_i + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi\right], \lambda_{i+1} + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi\right]\right).$$

Therefore, Q also belongs to their intersection, i.e.,

$$Supp(\tilde{Q}|\pi_t = \pi+e_i \ \& \ \pi_{t-1} = \pi) = \left[\lambda_i + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi\right], \lambda_{i+1} + \mathbb{E}\left[\tilde{Q}|\pi_{t-1} = \pi\right]\right) \cap Supp(\tilde{Q}|\pi_{t-1} = \pi).$$

Since this holds for any arbitrary $Q \in \mathcal{Q}$, the result follows. □

B.1.5 Proof of Lemma 3.2

Part (i) Follows from the proof of Lemma 3.1.

Part (ii) In an aggregate system, the customer refines her belief by considering all scenarios leading to the current rating frequency, π_t. Thus, $Supp\left(\tilde{Q}|\pi_t\right)$ is the union of all belief supports conditional on different complete rating paths, i.e., $Supp\left(\tilde{Q}|\pi_t \ \& \ (r_1, ..., r_t)\right)$. Since the platform observes the entire rating history, we have

$$\begin{aligned} Supp\left(\tilde{Q}^p|r = (r_1, ..., r_t)\right) = \bigcap_{\tau<t}\bigcup_{r} Supp\left(\tilde{Q}|\pi_\tau \ \& \ r = (r_1, ..., r_\tau)\right) = \\ \bigcap_{\tau<t}[\underline{Q}_\tau, \overline{Q}_\tau) = [\max_{\tau<t}\{\underline{Q}_\tau\}, \min_{\tau<t}\{\overline{Q}_\tau\}), \end{aligned} \tag{67}$$

which completes the proof. □

B.1.6 Proof of Proposition 3.4

We first prove the following claim.

Claim 1. For $t \geq 2$, only ratings $r_t = -1$ and $r_t = 1$ may occur.

The claim is trivial for the case with $k = 1$. Suppose $k \geq 2$. When $\alpha = 1$, the reference utility $u^r(1, \theta, Q^c) = Q - Q^c \in [-1, 1]$. This implies that the ratings $-k$ and k never occur. Therefore, for all i such that $\pi_t(i) > 0$,

$$Supp(\tilde{Q}_t|\pi_t) \subseteq \bigcup_{i \in R \setminus \{-k,k\}} \Big([\lambda_i + Q_t^c, \lambda_{i+1} + Q_t^c) \cap Supp(\tilde{Q}_{t-1}|\pi_t - e_i) \Big). \tag{68}$$

For each first-period rating r_1, we have $Supp(\tilde{Q}_2) = [\lambda_{r_1} + E(\tilde{Q}), \lambda_{r_1+1} + E(\tilde{Q}))$. We observe that $|Supp(\tilde{Q}_2)| = \frac{1}{k-1}$, which yields that $|Q - Q_2^c| \leq \frac{1}{k-1}$. Note that from Equation (68), $Supp(\tilde{Q}_t) \subseteq Supp(\tilde{Q}_2)$. Hence, $|Q - Q_t^c| \leq \frac{1}{k-1}$ and r_t is either $+1$ or -1 for all $t \geq 2$. This proves Claim 1.

Next, we prove the following claim.

Claim 2. Suppose $\tilde{Q}_{T-1} \sim U(\underline{Q}_{T-1}, \overline{Q}_{T-1})$ and $|r_t| = 1$ for $\forall t > T$. If $Supp(\tilde{Q}_t|\pi_t) = [\underline{Q}_t, \overline{Q}_t) \subseteq [\underline{Q}_{T-1}, \overline{Q}_{T-1})$, then $\forall Q \in [\underline{Q}_{T-1}, \overline{Q}_{T-1})$, we have $Q \in Supp(\tilde{Q}_t) = [\underline{Q}_t, \overline{Q}_t)$ for some $\underline{Q}_t$ and $\overline{Q}_t$. Furthermore, $\lim_{t\to\infty}(\overline{Q}_t - \underline{Q}_t) = 0$.

From Claim 1, $\forall t \geq 2$, only two rating paths may generate rating frequency π_t: $\{\pi_{t-1}^{(1)} = \pi_t - e_1 \ \&\ r_{t-1} = 1\}$ or $\{\pi_{t-1}^{(-1)} = \pi_t - e_{-1} \ \&\ r_{t-1} = -1\}$.

Additionally, from Lemma 3.2, we know that $Supp(\tilde{Q}_{t-1}|\pi_{t-1}^{(1)}) = [\underline{Q}_{t-1}^{(1)}, \overline{Q}_{t-1}^{(1)})$ and $Supp(\tilde{Q}_{t-1}|\pi_{t-1}^{(-1)}) = [\underline{Q}_{t-1}^{(-1)}, \overline{Q}_{t-1}^{(-1)})$ for some $\underline{Q}_{t-1}^{(1)}, \overline{Q}_{t-1}^{(1)}, \underline{Q}_{t-1}^{(-1)}$, and $\overline{Q}_{t-1}^{(-1)}$. Moreover, from the partition property established in Lemma 3.1, $\overline{Q}_{t-1}^{(1)} = \underline{Q}_{t-1}^{(-1)}$. As such, from Equation (68),

$$Supp(\tilde{Q}_t|r_{t-1} = 1 \ \&\ \pi_{t-1}^{(1)}) = \left[\mathbb{E}\left[\tilde{Q}_{t-1}|\pi_{t-1}^{(1)}\right], \lambda_2 + \mathbb{E}\left[\tilde{Q}_{t-1}|\pi_{t-1}^{(1)}\right]\right) \cap Supp(\tilde{Q}_{t-1}|\pi_{t-1}^{(1)}), \tag{69}$$

and

$$Supp(\tilde{Q}_t|r_{t-1} = -1 \ \&\ \pi_{t-1}^{(-1)}) = \left[\lambda_{-1} + \mathbb{E}\left[\tilde{Q}_{t-1}|\pi_{t-1}^{(-1)}\right], \mathbb{E}\left[\tilde{Q}_{t-1}|\pi_{t-1}^{(-1)}\right]\right) \cap Supp(\tilde{Q}_{t-1}|\pi_{t-1}^{(-1)}). \tag{70}$$

Let $Supp(\tilde{Q}_t|\pi_t) = [\underline{Q}_t, \overline{Q}_t)$. Then, Equation (68) also implies that

$$\underline{Q}_{t-1}^{(1)} < \mathbb{E}\left[\tilde{Q}_{t-1}|\pi_{t-1}^{(1)}\right] = \underline{Q}_t < \overline{Q}_t = \mathbb{E}\left[\tilde{Q}_{t-1}|\pi_{t-1}^{(-1)}\right] < \overline{Q}_{t-1}^{(-1)}. \tag{71}$$

Without loss of generality, reparameterize time to $\tau = t - T + 1$. With this reparameterize, $supp(\tilde{Q}_1) = [\underline{Q}_T, \overline{Q}_T)$. For technical convenience, we further normalize interval $[\underline{Q}_T, \overline{Q}_T)$ to $[-1, 1)$ by changing rating thresholds, intrinsic quality, and the initial belief distribution. This normalization will help us simplify the subsequent analysis. Thus, $supp(\tilde{Q}_1) = [\underline{Q}_T, \overline{Q}_T) = [-1, 1)$. From the partition property established in Lemma 3.2, there exist a set of cutoffs $Cut_\tau = \{-1 = x_\tau^{N+1}, x_\tau^N, ..., x_\tau^2, x_\tau^1, x_\tau^0 = 1\}$ such that $Supp(\tilde{Q}_\tau | \pi_\tau^n) = [x_\tau^{N+2-n}, x_\tau^{N+1-n})$, where π_τ^n are the feasible rating frequencies defined in Lemma 3.1 and $N+1$ is the number of feasible rating frequencies by time τ. We also observe that when $\alpha = 1$ and beliefs are uniformly distributed, the distribution of beliefs on $Supp(\tilde{Q}|\pi_\tau^n)$ is symmetric and uniformly distributed upon observing π_τ^n. This observation, combined with Equation (68), implies that $x_{\tau+1}^n = (x_\tau^n + x_\tau^{n-1})/2$ for any n.

With this background, we also observe that $N = \tau$. Furthermore, due to symmetry of the distributions, when τ is odd, we can rewrite the cutoffs as $Cut_\tau = \{-1, ..., -x_\tau^{(\tau-1)/2}, 0, x_\tau^{(\tau-1)/2}, ..., x_\tau^1, 1\}$. When τ is even, we can rewrite the cutoffs as $Cut_\tau = \{-1, -x_\tau^1, ..., -x_\tau^{\tau/2}, x_\tau^{\tau/2}, ..., x_\tau^1, 1\}$.

For the rest of the proof, we focus on the positive cutoffs, i.e., $x_\tau^n > 0$. The argument for the negative cutoffs is similar due to symmetry. Define $x^n = \{x_\tau^n : \tau \geq 2n\}$ (note that for x_τ^n to be one of the cutoffs, we must have $\tau \geq 2n$). In the following, we characterize x_τ^n.

For $n = 1$ and $\tau = 2$, we have $x_2^1 = \frac{0+1}{2}$. Further, $\forall \tau > 2$, we have $x_\tau^1 = \frac{x_{\tau-1}^1 + 1}{2}$. Hence,

$$x_\tau^1 = 1 - \frac{1}{2^{\tau-1}}. \tag{72}$$

For $n = 2$ and $\tau = 4$, Then, $x_4^2 = \mathbb{E}\left[\tilde{Q}_3 | Supp(\tilde{Q}_3) = [0, x_3^1)\right] = \frac{0 + x_3^1}{2}$. In addition, $x_{\tau+1}^2 = \frac{x_\tau^2 + x_\tau^1}{2}$. Therefore, for $n = 2$ and $\tau \geq 4$,

$$x_\tau^2 = \sum_{k=3}^{\tau-1} \frac{x_k^1}{2^{\tau-k}} = 1 - \frac{\tau+1}{2^{\tau-1}}. \tag{73}$$

Repeating these steps, we find that

$$x_\tau^n = \sum_{k=2n-1}^{\tau-1} \frac{x_k^{n-1}}{2^{\tau-k}}, \text{ for } \tau \geq 2n. \tag{74}$$

We argue that for $\tau \geq 2n$, we have

$$x_\tau^{n-1} - x_\tau^n = \frac{\Pi_{k=1}^{n-1}(\tau + 1 - k)}{(n-1)!\, 2^{\tau-1}}. \tag{75}$$

We prove this by induction on $n-1$, i.e, if the result holds for $x_\tau^{n-2} - x_\tau^{n-1}$, then it holds for $x_\tau^{n-1} - x_\tau^n$.

We begin with $n = 2$ as the basis of the induction. Then,

$$x_\tau^1 - x_\tau^2 = 1 - \frac{1}{2^{\tau-1}} - (1 - \frac{\tau+1}{2^{\tau-1}}) \quad = \frac{\tau}{2^{\tau-1}}. \tag{76}$$

Next, suppose Equation (75) holds for $x_\tau^{n-2} - x_\tau^{n-1}$ for all τ. Then, for Period $\tau - 1$,

$$x_{\tau-1}^{n-2} - x_{\tau-1}^{n-1} = \frac{\Pi_{k=1}^{n-2}(\tau - k)}{(n-2)!\,2^{\tau-2}}. \tag{77}$$

Alternatively,

$$x_{\tau-1}^{n-2} = x_{\tau-1}^{n-1} + \frac{\Pi_{k=1}^{n-2}(\tau - k)}{(n-2)!\,2^{\tau-2}}. \tag{78}$$

Recall that for all τ and n, $x_\tau^n = \frac{x_{\tau-1}^n + x_{\tau-1}^{n-1}}{2}$. Therefore,

$$\begin{aligned} x_\tau^{n-1} - x_\tau^n &= \frac{x_{\tau-1}^{n-1} + x_{\tau-1}^{n-2}}{2} - \frac{x_{\tau-1}^n + x_{\tau-1}^{n-1}}{2} \\ &= \frac{x_{\tau-1}^{n-1} + \frac{\Pi_{k=1}^{n-2}(\tau-k)}{(n-2)!2^{\tau-2}}}{2} - \frac{x_{\tau-1}^n}{2} \\ &= \frac{x_{\tau-1}^{n-1} - x_{\tau-1}^n}{2} + \frac{\Pi_{k=1}^{n-2}(\tau - k)}{(n-2)!\,2^{\tau-1}}. \end{aligned} \tag{79}$$

By iteratively applying (78) for $\tau = 2n - 1$, we have

$$\begin{aligned} x_{2n}^{n-1} - x_{2n}^n &= \frac{x_{2n-1}^{n-2}}{2} \\ &= \frac{1}{2} \cdot \left(x_{2n-1}^{n-3} - \frac{\Pi_{k=1}^{n-3}(2n-k)}{(n-3)!\,2^{2n-2}} \right) \\ &= \frac{1}{2} \cdot \left(x_{2n-1}^{1} - \sum_{j=1}^{n-3} \frac{\Pi_{k=1}^{j}(2n-k)}{(j)!\,2^{2n-2}} \right) \\ &= \frac{1}{2} \cdot \left(1 - \frac{1}{2^{2n-2}} - \sum_{j=1}^{n-3} \frac{\Pi_{k=1}^{j}(2n-k)}{(j)!\,2^{2n-2}} \right) \\ &= \frac{1}{2^{2n-1}} \cdot \left(2^{2n-2} - 1 - \sum_{j=1}^{n-3} \frac{\Pi_{k=1}^{j}(2n-k)}{(j)!} \right) \\ &= \frac{1}{2^{2n-1}} \cdot \left(2^{2n-2} - 1 - \sum_{j=1}^{n-1} \frac{\Pi_{k=1}^{j}(2n-k)}{(j)!} + \frac{\Pi_{k=1}^{n-2}(2n-k)}{(n-2)!} + \frac{\Pi_{k=1}^{n-1}(2n-k)}{(n-1)!} \right) \\ &= \frac{1}{2^{2n-1}} \cdot \left(2^{2n-2} - 1 - \sum_{j=1}^{n-1} \frac{\Pi_{k=1}^{j}(2n-k)}{(j)!} + \frac{\Pi_{k=1}^{n-1}(2n+1-k)}{(n-1)!} \right). \end{aligned} \tag{80}$$

Note that since $\frac{\Pi_{k=1}^{j}(2n-k)}{(j)!} = \binom{2n-1}{j}$,

$$\sum_{j=1}^{n-1} \frac{\Pi_{k=1}^{j}(2n-k)}{(j)!} = 2^{2n-1} - 1 - \sum_{j=0}^{n-1} \binom{2n-1}{j} = 2^{2n-1} - 1 - 2^{2n-2} = 2^{2n-2} - 1. \tag{81}$$

Plugging this equation into Equation (80), we find that $x_{2n}^{n-1} - x_{2n}^{n} = \frac{\Pi_{k=1}^{n-1}(2n+1-k)}{(n-1)!2^{2n-1}}$, which shows the induction result for $\tau = 2n$.

$$\begin{aligned} x_{2n+1}^{n-1} - x_{2n+1}^{n} &= \frac{x_{2n}^{n-1} - x_{2n}^{n}}{2} + \frac{\Pi_{k=1}^{n-2}(2n+1-k)}{(n-2)!\,2^{2n}} \\ &= \frac{\Pi_{k=1}^{n-1}(2n+1-k)}{(n-1)!\,2^{2n}} + \frac{\Pi_{k=1}^{n-2}(2n+1-k)}{(n-2)!\,2^{2n}} \\ &= \frac{\Pi_{k=1}^{n-1}(2n+2-k)}{(n-1)!\,2^{2n}}, \end{aligned} \tag{82}$$

where the first equality followed from applying Equation (79). By iteratively repeating these steps we find that

$$x_{\tau}^{n-1} - x_{\tau}^{n} = \frac{\Pi_{k=1}^{n-1}(\tau+1-k)}{(n-1)!\,2^{\tau-1}}, \tag{83}$$

for $\tau \geq 2n$, which is what we desired to show.

Note that $x_{\tau}^{n-1} - x_{\tau}^{n}$ is the length of each partition at Period τ. Further, $\forall n$,

$$\lim_{\tau\to\infty}(x_{\tau}^{n-1} - x_{\tau}^{n}) = \lim_{\tau\to\infty} \frac{\Pi_{k=1}^{n-1}(\tau+1-k)}{(n-1)!\,2^{\tau-1}} = 0. \tag{84}$$

As such, as $t \to \infty$, the size of each element of the partition approaches 0. That is to say, $\lim_{t\to\infty}(\overline{Q}_t - \underline{Q}_t) = 0$, which completes the proof for Claim 2.

Combining Claims 1 and 2, and the observations that all feasible rating frequencies partition the entire support, and that the intrinsic quality Q belongs to a certain element of the partition at Period t, we conclude that the $\lim_{t\to\infty} Q_t^c = Q$. □

B.1.7 Proof of Lemma 3.3

We prove the above lemma by showing that for any $Q \in [0,1]$, we can separate Q from any other potential quality value. We prove cases of $k = 1$ and $k \geq 2$, separately.

Case (i) $k \geq 2$. Consider $Q_1, Q_2 \in [0,1]$ such that $Q_1 > Q_2$. We show that we can separate Q_1 from Q_2 using strict separation in distribution defined in Definition 2 when α is small. Let $\overline{r}$ be the highest rating that satisfies $\inf_{\tilde{q} \in \mathcal{F}} \kappa(\overline{r}; Q_1, \tilde{q}) > 0$.

We use $S = \{\overline{r}\}$ to establish the condition required for the separation in distribution. Recall that Q^c depends on the distribution of the customer beliefs, i.e., $\tilde{Q}$. For notational convenience, we drop this dependence from the argument of Q^c. With this note in mind, we have

$$\inf_{\tilde{q} \in \mathcal{F}} \kappa(\overline{r}; Q_1, \tilde{q}) = \inf_{\tilde{q} \in \mathcal{F}} \mathbb{P}_\theta(\theta \geq \frac{\lambda_{\overline{r}} + \alpha Q^c - Q_1}{1-\alpha} + p) \geq \mathbb{P}_\theta(\theta \geq \frac{\lambda_{\overline{r}} + \alpha - Q_1}{1-\alpha} + p), \tag{85}$$

where the inequality followed from the fact that $Q^c \leq 1$. Similarly,

$$\sup_{\tilde{q} \in \mathcal{F}} \kappa(\overline{r}; Q_2, \tilde{q}) = \sup_{\tilde{q} \in \mathcal{F}} \mathbb{P}_\theta(\theta \geq \frac{\lambda_{\overline{r}} + \alpha Q^c - Q_2}{1-\alpha} + p) \leq \mathbb{P}_\theta(\theta \geq \frac{\lambda_{\overline{r}} - Q_2}{1-\alpha} + p), \tag{86}$$

where the inequality followed from the fact that $Q^c \geq 0$.

Consequently, a sufficient condition for the strict separation in distribution to hold is

$$\mathbb{P}_\theta(\theta \geq \frac{\lambda_{\overline{r}} + \alpha - Q_1}{1-\alpha} + p) > \mathbb{P}_\theta(\theta \geq \frac{\lambda_{\overline{r}} - Q_2}{1-\alpha} + p). \tag{87}$$

Similar to the Proof of Proposition 3.2, we can show that $\kappa(\overline{r}; Q_1, \tilde{Q}) < 1$ for any arbitrary belief $\tilde{Q}$, when $\alpha \leq 1 - \frac{1}{2\overline{\theta}(k-1)}$. Additionally, $F_\theta(\cdot)$ is continuous and strictly increasing. Hence, condition in Equation (87) holds if an only if $Q_1 - Q_2 > \alpha$. Intuitively, any two quality values with a distance greater than α can be strictly separated in distribution.

Let Q_2 be the intrinsic quality. Following the discussion above, any $Q_1 \notin (Q_2 - \alpha, Q_2 + \alpha)$ can be separated from the intrinsic quality. Hence, the support of the customer belief is $Supp(\tilde{Q}) \subseteq [\max(Q_2 - \alpha, 0), \min(Q_2 + \alpha, 1)]$. Therefore, $Q^c \in [\max(Q_2 - \alpha, 0), \min(Q_2 + \alpha, 1)]$. Let $\mathcal{F}_1$ be the space of all distribution functions with support $[\max(Q_2 - \alpha, 0), \min(Q_2 + \alpha, 1)]$. Additionally, let $\overline{r}_1$ be the highest rating that satisfies $\inf_{\tilde{q} \in \mathcal{F}_1} \kappa(\overline{r}_1; Q_1, \tilde{q}) > 0$. We use

$S_1 = \{\bar{r}_1\}$ to establish the condition required for the separation in distribution between Q_1 and Q_2 with a more refined domain. We have

$$\inf_{\tilde{q}\in\mathcal{F}_1} \kappa(\bar{r}_1; Q_1, \tilde{q}) = \inf_{\tilde{q}\in\mathcal{F}_1} \mathbb{P}_\theta(\theta \geq \frac{\lambda_{\bar{r}_1} + \alpha Q^c - Q_1}{1-\alpha} + p) \geq \mathbb{P}_\theta(\theta \geq \frac{\lambda_{\bar{r}_1} + \alpha \min\{1, Q_2 + \alpha\} - Q_1}{1-\alpha} + p), \tag{88}$$

where the last inequality followed from $Q^c \leq \min\{1, Q + \alpha\}$. Similarly,

$$\sup_{\tilde{q}\in\mathcal{F}_1} \kappa(\bar{r}_1; Q_2, \tilde{q}) = \sup_{\tilde{q}\in\mathcal{F}_1} \mathbb{P}_\theta(\theta \geq \frac{\lambda_{\bar{r}_1} + \alpha Q^c - Q_2}{1-\alpha} + p) \leq \mathbb{P}_\theta(\theta \geq \frac{\lambda_{\bar{r}_1} + \alpha \max\{0, Q_2 - \alpha\} - Q_2}{1-\alpha} + p), \tag{89}$$

where the last inequality followed from $Q^c \geq \max\{0, Q_2 - \alpha\}$.

Therefore, a sufficient condition this time for the strict separation in distribution to hold is

$$\mathbb{P}_\theta(\theta \geq \frac{\lambda_{\bar{r}_1} + \alpha \min\{1, Q_2 + \alpha\} - Q_1}{1-\alpha} + p) > \mathbb{P}_\theta(\theta \geq \frac{\lambda_{\bar{r}_1} + \alpha \max\{0, Q_2 - \alpha\} - Q_2}{1-\alpha} + p). \tag{90}$$

Similar to the argument for Condition (87), one can argue that Condition (90) is equivalent to

$$Q_1 - Q_2 > \alpha(\min\{1, Q_2 + \alpha\} - \max\{0, Q_2 - \alpha\}). \tag{91}$$

Note that since $\min\{1, Q_2 + \alpha\} - \max\{0, Q_2 - \alpha\} \leq 2\alpha$, a sufficient condition for (90) to hold is $Q_1 - Q_2 \geq \alpha(2\alpha)$. Intuitively, any two quality values with a distance greater than $\alpha(2\alpha)$ can be strictly separated in distribution. Suppose $\alpha < 1/2$; then, $\alpha(2\alpha) < \alpha$. In other words, in the second iteration of the argument, the distance between Q_1 and Q_2 required for strictly separating the two values in distribution decreases. By repeating this belief refinement process n times, we can show that Q_1 and Q_2 can be strictly separated in distribution if $Q_1 - Q_2 > \alpha(2\alpha)^{n-1}$. As $n \to \infty$, we have $\alpha(2\alpha)^{n-1} \to 0$. Hence, Q_1 and Q_2 can be strictly separated in distribution if $Q_1 > Q_2$. One can similarly argue the case when $Q_1 < Q_2$. Recall that we assumed $\alpha < 1/2$. Additionally, to ensure $\kappa(\bar{r}; Q_1, \tilde{Q}) < 1$, we required $\alpha \leq 1 - \frac{1}{2\bar{\theta}(k-1)}$. Hence, one can separate any two possible values for the intrinsic quality if $\alpha < \bar{\alpha}_I^1 \equiv \min\{\frac{1}{2}, 1 - \frac{1}{2\bar{\theta}(k-1)}\}$.

Case (ii) $k = 1$. When $k = 1$, we use $S = \{\bar{r}\} = \{1\}$ to strictly separate any two quality values Q_1 and Q_2. Similar to the proof of Proposition 3.2, to guarantee $\kappa(r = 1; Q_1, \tilde{Q})$ and

$\kappa(r=1; Q_2, \tilde{Q})$ to be non-negative and less than 1, we must have $\alpha < \max\{1-\frac{1}{1+\overline{\theta}-p}, 1-\frac{1}{\overline{\theta}+p}\}$. We can then repeat a similar argument as in Case (i) to prove any two quality values can be strictly separated in distribution if $\alpha < \overline{\alpha}_I^2 \equiv \min\{\frac{1}{2}, \max\{1-\frac{1}{1+\overline{\theta}-p}, 1-\frac{1}{\overline{\theta}+p}\}\}$.

The result of the Lemma then follows by defining $\overline{\alpha}_I \equiv \overline{\alpha}_I^2$ when $k=1$ and $\overline{\alpha}_I \equiv \overline{\alpha}_I^1$ when $k \geq 2$. □

B.1.8 Proof of Proposition 3.5

Part (i) Immediately follows from Lemma 3.3.

Part (ii) We add argument (α, Q) to $Q_t^c(\alpha, Q)$ to emphasize its dependence on the disconfirmation bias and the intrinsic quality. If $\lim_{t\to\infty} Q_t^c(\alpha, Q) = Q_\infty^c(\alpha, Q)$ exists, then we show that $Q_\infty^c(\alpha, Q)$ is strictly increasing in the intrinsic quality Q. We prove this by contradiction. Consider $Q_1 < Q_2$. By contradiction, assume $Q_\infty^c(\alpha, Q_2) \leq Q_\infty^c(\alpha, Q_1)$. Then,

$$\begin{aligned} u^r(\alpha, Q_1) &= Q_1 - \alpha Q_\infty^c(\alpha, Q_1) - (1-\alpha)p + (1-\alpha)\theta \\ &< Q_2 - \alpha Q_\infty^c(\alpha, Q_2) - (1-\alpha)p + (1-\alpha)\theta = u^r(\alpha, Q_2). \end{aligned} \tag{92}$$

Furthermore, similar to the Proof of Proposition 3.2, we find that, when $\alpha \leq \overline{\alpha}(k, \overline{\theta})$, at least two ratings are realizable under Q_1 and Q_2 (note that Proposition 3.2 is for the complete history setting while in this section we study the aggregate system. However, the customer ratings only depend on the reference utility and not the customer information set. So, we can use the same argument as in the proof of Proposition 3.2).

Combining (92) and that at least two ratings realize under Q_1 and Q_2, we observe that the distribution of the ratings generated under intrinsic quality Q_2 strictly stochastically dominates the distribution of ratings generated under intrinsic quality Q_1. Hence, $Q_\infty^c(\alpha, Q_2) > Q_\infty^c(\alpha, Q_1)$, which is a contradiction. Therefore, $Q_\infty^c(\alpha, Q)$ is strictly increasing in the intrinsic quality Q. Furthermore, since customers are Bayesian rational, they can identify any systematic bias. As such, $Q_\infty^c(\alpha, Q) = Q$.

Part (iii) We show that, when $\alpha > \overline{\alpha}(k, \overline{\theta})$, if $Q_\infty^c(\alpha, Q)$ exists, there exist Q_1 and Q_2, such that $Q_1 \neq Q_2$, and $Q_\infty^c(\alpha, Q_1) = Q_\infty^c(\alpha, Q_2)$.

If Q_t^c converges, for any arbitrary $\epsilon > 0$, there exists period T_ϵ such that $|Q_t^c - Q_\infty^c|< \frac{\epsilon}{2}$ for $t \geq T_\epsilon$. Let π_{T_ϵ} denote the rating frequency that Customer T_ϵ observes. First, we prove the following claim.

Claim 1. For any intrinsic quality Q_1, there exists Q_2 such that they can both generate the rating frequency π_{T_ϵ} at period T_ϵ.

We prove the claim by induction on T_ϵ. When $T_\epsilon = 1$, then all customers have the same prior belief. Since the reference utility is Continuous in Q, for a fixed Q^c, we can choose Q_2 arbitrarily close to Q_1 such that both Q_1 and Q_2 induce the same rating in Period 1. Suppose the result holds for $T_\epsilon = T - 1$, we show the result holds for $T_\epsilon = T$. Since there exists Q_2 such that it generates the same rating frequency as for Q_1 until period $T - 1$, in Period T we have $Q_T^c(\alpha, Q_1) = Q_T^c(\alpha, Q_2)$. Once again, since the reference utility is continuous in Q, we can find another quality value that generates the same rating r_T as the rating that is induced by Q_1. This proves Claim 1.

Next, we prove the following claim.

Claim 2. For any intrinsic quality Q_1, there exists neighborhood $\mathcal{N}_{Q_1}$ such that $\kappa(i^*, Q_2, \tilde{Q}_t) = 1$ for some rating i^* and any $t \geq T_\epsilon$ and $Q_2 \in \mathcal{N}_{Q_1}$.

To see this, note that customers rate according to their reference utility. Consider $\alpha > 1 - \frac{1}{2\overline{\theta}(k-1)} + \epsilon$. The minimum value for the reference utility is

$$\begin{aligned}\min_{Q_t^c, \theta} u^r(\alpha, \theta, Q_t^c) &= Q_1 - \alpha \max(Q_t^c) - (1 - \alpha) \max(\theta) + (1 - \alpha)p \\ &= Q_1 - \alpha(Q_\infty^c(\alpha, Q) + \epsilon/2) - (1 - \alpha)\overline{\theta} + (1 - \alpha)p,\end{aligned} \tag{93}$$

where we used the facts that $|Q_t^c - Q_\infty^c|< \frac{\epsilon}{2}$ for $t \geq T_\epsilon$ and that $\theta \leq \overline{\theta}$.

Similarly,

$$\begin{aligned}\max_{Q_t^c, \theta} u^r(\alpha, \theta, Q_t^c) &= Q_1 - \alpha \min(Q_t^c) - (1 - \alpha) \min(\theta) + (1 - \alpha)p \\ &= Q_1 - \alpha(Q_\infty^c(\alpha, Q) - \epsilon/2) + (1 - \alpha)\overline{\theta} + (1 - \alpha)p.\end{aligned} \tag{94}$$

Hence,

$$\max_{Q_t^c, \theta} u^r(\alpha, \theta, Q_t^c) - \min_{Q_t^c, \theta} u^r(\alpha, \theta, Q_t^c) = 2(1 - \alpha)\overline{\theta} + \alpha\epsilon < 1/(k - 1), \tag{95}$$

where the last inequality followed from $\alpha > 1-\frac{1}{2\overline{\theta}(k-1)}+\epsilon$. Hence, for $t > T_\epsilon$, all customers rate the product with intrinsic quality Q_1 the same, namely i^*. Since $u^r(\alpha, \theta, Q_t^c)$ is continuous in Q, the existence of $\mathcal{N}_{Q_1}$ follows.

Combining Claims 1 and 2 and letting $\epsilon \to 0$, we observe that, when $\alpha > \overline{\alpha}(k, \overline{\theta})$, for any intrinsic quality Q_1, there exists another quality value that generates the same rating frequency in all periods which proves Part (iii).

Part (iv) follows from the definition of $\overline{\alpha}_I(k, \overline{\theta})$ provided in Lemma 3.3. □

B.1.9 Proof of Proposition 3.6

First, we show the convergence of the customer beliefs, i.e., the existence of $\hat{Q}_\infty^c < \infty$. For any $t \geq 1$, we have

$$\hat{Q}_{t+1}^c = \frac{1}{\gamma \cdot t + 1} E_0(Q) + \frac{\gamma \cdot t}{\gamma \cdot t + 1} \overline{r}_t = \frac{\gamma \cdot (t-1) + 1}{\gamma \cdot t + 1} \hat{Q}_t^c + \frac{\gamma}{\gamma \cdot t + 1} r_t. \tag{96}$$

The convergence then follows from noting that $\lim_{t\to\infty} \frac{\gamma \cdot (t-1)+1}{\gamma \cdot t+1} = 1^-$, $\lim_{t\to\infty} \frac{\gamma}{\gamma \cdot t+1} = 0$, and $r_t < \infty$.

Next, we show that $\hat{Q}_\infty^c(Q, \alpha)$ "weakly" increases in Q for "any" value of α. We prove this by contradiction.

By contradiction suppose $Q_1 < Q_2$, but $\hat{Q}_\infty^c(Q_1, \alpha) > \hat{Q}_\infty^c(Q_2, \alpha)$. Then,

$$\begin{aligned} u^r(\alpha, \theta, \hat{Q}_\infty^c(Q_1, \alpha)) = Q_1 - \alpha \hat{Q}_\infty^c(Q_1, \alpha) + (1-\alpha)\theta - (1-\alpha)p < \\ Q_2 - \alpha \hat{Q}_\infty^c(Q_2, \alpha) + (1-\alpha)\theta - (1-\alpha)p = u^r(\alpha, \theta, \hat{Q}_\infty^c(Q_2, \alpha)). \end{aligned} \tag{97}$$

As a result, in steady states, the ratings generated for Q_1 would be weakly worse than the ratings generated for Q_2. In other words, Q_1 receives infinitely many equal or worse ratings than Q_2. Furthermore, from Equation (13), we observe that $\hat{Q}_\infty^c(Q, \alpha) = \overline{r}_\infty(Q, \alpha)$. Therefore, $\hat{Q}_\infty^c(Q_1, \alpha) = \overline{r}_\infty(Q_1, \alpha) \leq \overline{r}_\infty(Q_2, \alpha) = \hat{Q}_\infty^c(Q_2, \alpha)$, which is a contradiction. This proves the weakly increasing property of $\hat{Q}_\infty^c(Q, \alpha)$ in Q for any α. This immediately yields the result in Part (ii).

Next, we prove Part (i) by showing strict monotonicity when α is small. From the weak monotonicity property established in the above argument, for any $Q_1 < Q_2$, we have $\hat{Q}_\infty^c(Q_1, \alpha) \leq \hat{Q}_\infty^c(Q_2, \alpha)$. If the inequality is strict, there is nothing to prove. Suppose

$\hat{Q}^c_\infty(Q_1,\alpha) = \hat{Q}^c_\infty(Q_2,\alpha)$. Note that since the beliefs converges to $\bar{r}_\infty(Q,\alpha)$, in steady state, we must have

$$\hat{Q}^c_\infty(Q,\alpha) = \sum_{j=-K}^{K} j \times \kappa(j;Q,\tilde{Q}_\infty) = \sum_{j=-K}^{K} j \times \mathbb{P}_\theta[\lambda_j \leq Q - \alpha Q^c - (1-\alpha)(p-\theta) < \lambda_{j+1}], \tag{98}$$

where the last equality followed from the definition of $\kappa(r=i;Q,\tilde{Q})$. Hence, $\hat{Q}^c_\infty(Q_1,\alpha) = \hat{Q}^c_\infty(Q_2,\alpha)$ can hold only if, in steady state, only "one" rating can realize under both quality values, i.e., $\kappa(r=i;Q_1,\tilde{Q}) = \kappa(r=i;Q_2,\tilde{Q}) = 1$ for some i. We find conditions such that this does not occur.

Similar to the proof of Proposition 3.2, a sufficient condition for this is to have $Q_1 - \alpha\hat{Q}^c_\infty - (1-\alpha)p - (1-\alpha)\overline{\theta} \geq \lambda_{i-1}$, and $Q_2 - \alpha\hat{Q}^c_\infty - (1-\alpha)p + (1-\alpha)\overline{\theta} < \lambda_i$, which holds if $\alpha \leq 1 - \frac{1}{2\overline{\theta}(k-1)}$. This completes the proof of Part (i).

Part (iii) By contradiction, suppose $Q > \hat{Q}^c_\infty(Q,1)$ (the proof for $Q < \hat{Q}^c_\infty(Q,1)$ is similar). If $\alpha = 1$, we have $u^r(\alpha,\theta,\hat{Q}^c_\infty(Q,\alpha)) = Q - \hat{Q}^c_\infty(Q,1) \equiv \Delta > 0$. Hence, there exists t_Δ such that for any $t > t_\Delta$, we have $\Delta > 0$ and customers rate the product positively. In other words, $r_{t+1} \geq 1$. Consequently, $\hat{Q}^c_\infty(Q,1) = \bar{r}_\infty \geq 1$. Since $Q \in [0,1]$, we have $Q - \hat{Q}^c_\infty(Q,1) \leq 0$, which is a contradiction with the initial assumption that $\Delta > 0$. As such, we must have $\Delta = 0$. □

B.1.10 Proof of Corollary 2

First, we prove the following claim.

Claim 1. When $\alpha \leq \overline{\alpha}(k,\overline{\theta})$, the function $\phi_\alpha^{-1}(x)$ is Lipschitz continuous on $\mathbb{R}$.

From the proof of Proposition 3.6, $\hat{Q}^c_\infty(Q,\alpha) = \sum_{j=-K}^{K} j \times \kappa(j;Q,\tilde{Q}_\infty)$. Since $\kappa(j;Q,\tilde{Q}_\infty)$ is continuous in Q for all j, we know that $\hat{Q}^c_\infty(Q,\alpha) = \phi_\alpha(Q)$ is continuous in Q. Further, since $\phi_\alpha(Q)$ is bounded on the compact interval $\mathcal{Q}$, $\phi_\alpha(Q)$ is uniformly continuous on $\mathcal{Q}$. Additionally, from Part (i) of Proposition 3.6, $\phi_\alpha(Q) = \hat{Q}^c_\infty(Q,\alpha)$ is bounded and strictly increasing in Q when $\alpha \leq \overline{\alpha}(k,\overline{\theta})$. Combining these observations and using established results of Lipschitz continuity of inverse functions, we find that $\phi_\alpha^{-1}(r)$ is Lipschitz continuous. This completes the proof of Claim 1.

Next, we show that the customers' beliefs converge when they use Equation (14) to form beliefs. We have

$$\begin{aligned}\hat{Q}^u_{t+1} &= \frac{1}{\gamma\cdot t+1}\mathbb{E}(\tilde{Q}) + \frac{\gamma\cdot t}{\gamma\cdot t+1}\cdot\phi_\alpha^{-1}(\bar{r}_t) \\ &= \frac{\gamma\cdot(t-1)+1}{\gamma\cdot t+1}\hat{Q}^u_t + \frac{1}{\gamma\cdot t+1}\left(\gamma\cdot t\phi_\alpha^{-1}(\bar{r}_t) - \gamma\cdot(t-1)\phi_\alpha^{-1}(\bar{r}_{t-1})\right) \\ &= \frac{\gamma\cdot(t-1)+1}{\gamma\cdot t+1}\hat{Q}^u_t + \frac{\gamma\phi_\alpha^{-1}(\bar{r}_t)}{\gamma\cdot t+1} + \frac{\gamma\cdot(t-1)}{\gamma\cdot t+1}\left(\phi_\alpha^{-1}(\bar{r}_t) - \phi_\alpha^{-1}(\bar{r}_{t-1})\right).\end{aligned} \tag{99}$$

Since $\phi_\alpha^{-1}(\bar{r}_t) < \infty$, then $\lim_{t\to\infty}\frac{\gamma\phi_\alpha^{-1}(\bar{r}_t)}{\gamma\cdot t+1} = 0$. We show that $\lim_{t\to\infty}\frac{\gamma\cdot(t-1)}{\gamma\cdot t+1}\left(\phi_\alpha^{-1}(\bar{r}_t) - \phi_\alpha^{-1}(\bar{r}_{t-1})\right) = 0$ by arguing $|(t-1)\phi_\alpha^{-1}(\bar{r}_t) - (t-1)\phi_\alpha^{-1}(\bar{r}_{t-1})| < \infty$. Since $\phi_\alpha^{-1}(\cdot)$ is Lipschitz continuous, for any $\bar{r}_1$ and $\bar{r}_2$, there exists constant D such that $|\phi_\alpha^{-1}(\bar{r}_1) - \phi_\alpha^{-1}(\bar{r}_2)| \le D|\bar{r}_1 - \bar{r}_2|$. Therefore,

$$\begin{aligned}|(t-1)\phi_\alpha^{-1}(\bar{r}_t) - (t-1)\phi_\alpha^{-1}(\bar{r}_{t-1})| &\le D|(t-1)\bar{r}_t - (t-1)\bar{r}_{t-1}| \\ &= D|t\bar{r}_t - (t-1)\bar{r}_{t-1} - \bar{r}_t| \\ &= D|r_t - \bar{r}_t| \\ &\le D(|r_t| + |\bar{r}_t|) < \infty.\end{aligned} \tag{100}$$

Hence, $\lim_{t\to\infty}\frac{\gamma\cdot(t-1)}{\gamma\cdot t+1}\left(\phi_\alpha^{-1}(\bar{r}_t) - \phi_\alpha^{-1}(\bar{r}_{t-1})\right) = 0$. As such, from Equation (99),

$$\lim_{t\to\infty}|\hat{Q}^u_{t+1} - \hat{Q}^u_t| = \lim_{t\to\infty}\left|\frac{\gamma\cdot(t-1)+1}{\gamma\cdot t+1}\hat{Q}^u_t - \hat{Q}^u_t\right| = 0, \tag{101}$$

which proves the convergence.

Finally, from the first equality in Equation (99), we observe that $\hat{Q}^u_\infty = \phi_\alpha^{-1}(\bar{r}_\infty) = Q$. □

B.1.11 Proof of Proposition 3.7

Part (i) Since $Q^c_t \to Q$ conditional on complete and correct learning, for any $\epsilon > 0$, there exists N such that $|Q^c_t - Q| \le \epsilon$ for $t \ge N$. Let $\mathcal{N}_\epsilon \equiv (\max(Q-\epsilon, 0), \min(Q+\epsilon, 1))$.

Consider arbitrary quality value $Q' \neq Q$. We show it is separation divergent from Q for large k. We show this for $Q' > Q$. The argument for $Q < Q'$ is similar.

Define $\overline{Q}_\epsilon \equiv \text{argmin}_{Q^c\in\mathcal{N}_\epsilon} Q - \alpha Q^c + (1-\alpha)\theta - (1-\alpha)p = \min\{Q+\epsilon, 1\}$ and $\underline{Q}_\epsilon \equiv \text{argmax}_{Q^c\in\mathcal{N}_\epsilon} Q - \alpha Q^c + (1-\alpha)\theta - (1-\alpha)p = \max\{Q-\epsilon, 0\} \ge -1$.

Since $\overline{\theta} \leq \frac{1}{1-\alpha} + Q - p$, then $\min_{Q^c} \min_{\theta} u^r(\alpha, \theta, Q^c) = Q - \alpha \overline{Q}_\epsilon - (1-\alpha)\overline{\theta} - (1-\alpha)p \geq -1$. In other words, there are ratings that would never arise under intrinsic quality Q. Further, if $Q' < Q$, for sufficiently a small $\epsilon > 0$, we have

$$\begin{aligned}\max_{Q^c} \min_{\theta} u^r(\alpha, \theta, \tilde{Q}|Q') &= Q' - \alpha \underline{Q} - (1-\alpha)\overline{\theta} - (1-\alpha)p \\ &< Q - \alpha \overline{Q}_\epsilon - (1-\alpha)\overline{\theta} - (1-\alpha)p = \min_{Q^c} \min_{\theta} u^r(\alpha, \theta, \tilde{Q}|Q).\end{aligned} \tag{102}$$

Hence, for sufficiently a granular rating system, i.e., $k > k_{Q'}$, at least one of the thresholds of the rating system, namely λ, belongs to the interval

$$\left(\max_{Q^c} \min_{\theta} u^r(\alpha, \theta, \tilde{Q}|Q'), \min_{Q^c} \min_{\theta} u^r(\alpha, \theta, \tilde{Q}|Q)\right)$$

. Therefore, there exist ratings that would arise under Q' but never under intrinsic quality Q. Since, in this section, we assume that $\mathcal{Q}$ is finite, we define $\bar{k} = \max_{Q' \in \{\mathcal{Q} \setminus Q\}} k_{Q'}$ to prove Part (i).

Part (ii) Consider intrinsic quality Q and arbitrary quality value $Q' < Q$. Define $\underline{r}_Q = \max\{i | \lambda_i \leq Q - \alpha \overline{Q}_\epsilon - (1-\alpha)\overline{\theta} - (1-\alpha)p\}$ and $\overline{r}_Q = \max\{i | \lambda_i \leq Q - \alpha \underline{Q}_\epsilon - (1-\alpha)\overline{\theta} - (1-\alpha)p\}$. Similarly, define $\underline{r}_{Q'}$ and $\overline{r}_{Q'}$ by replacing Q with Q'. As shown in the proof of Part (i), for sufficiently a granular rating system, we have $\underline{r}_Q > \overline{r}_{Q'}$. Hence, any rating $r \in R_{Q'} \equiv \{\underline{r}_{Q'}, ..., \overline{r}_{Q'}\}$ would separate Q' from Q. We have

$$\begin{aligned}\mathbb{P}(r \notin R_{Q'}) &= 1 - \sum_{\underline{r}_{Q'} \leq r \leq \overline{r}_{Q'}} \kappa(r; Q, \tilde{Q}) \\ &\geq 1 - \sum_{\underline{r}_{Q'} \leq r \leq \overline{r}_{Q'}} \kappa(r; Q, \tilde{Q} | \mathbb{E}(\tilde{Q}) = \overline{Q}_\epsilon) \\ &= 1 - F_\theta\left(\frac{\lambda_{\overline{r}_{Q'}+1} - Q + \alpha \overline{Q}_\epsilon + (1-\alpha)p}{1-\alpha}\right) \\ &\geq 1 - F_\theta\left(\frac{Q' - \alpha \underline{Q}_\epsilon + (1-\alpha)\overline{\theta} - (1-\alpha)p + \frac{1}{k-1} - Q + \alpha \overline{Q}_\epsilon + (1-\alpha)p}{1-\alpha}\right) \\ &= \bar{F}_\theta\left(\overline{\theta} - \frac{Q - Q' - \alpha(\overline{Q}_\epsilon - \underline{Q}_\epsilon) - \frac{1}{k-1}}{1-\alpha}\right) \equiv \Delta_\epsilon,\end{aligned} \tag{103}$$

where the last inequality followed from the definition of $\overline{r}_{Q'}$ and that $\lambda_{\overline{r}_{Q'}+1} = \lambda_{\overline{r}_{Q'}} + 1/(k-1)$.

Since $k \geq \bar{k}$, $\Delta_\epsilon \neq 0$. As such, the expected number of customers needed to separate Q' from Q with ϵ accuracy is bounded by $1/\Delta_\epsilon$, i.e., the mean of a Geometric distribution with parameter Δ_ϵ.

Part (iii). Follows from Part (ii) and definition of $V^L(Q)$. $\square$

B.1.12 Proof of Proposition 3.8

Consider intrinsic quality Q and correct learning. Hence, $Q_t^c \to Q$. Moreover, under assumption $\overline{\theta} > \frac{1}{1-\alpha} + Q - p$, following a similar argument used in the proof of Proposition 3.7, we find that all ratings can arise with non-zero probability. Recall that $q_t(Q)$ is the probability that Customer t believes the the intrinsic quality is Q. Note that because of correct learning, $\lim_{t\to\infty} \frac{\pi_t(i)}{t} = \kappa(i; Q, \tilde{Q}_\infty)$. Therefore, for any $\epsilon > 0$, there exists $N > 0$ such that for $\tau \geq N$,

$$q_\tau(Q) \geq 1 - \epsilon \text{ and } \left| \frac{\pi_\tau(i)}{\tau} - \kappa(i; Q, \tilde{Q}_\tau) \right| \leq \epsilon, \text{ for all } i \text{ and } \tilde{Q}_\tau. \tag{104}$$

Let $\mathcal{F}_\tau^\epsilon(Q)$ be the space of all feasible customer beliefs realizable under intrinsic quality Q at time τ such that $q_\tau(Q) \geq 1 - \epsilon$. Then, define

$$\begin{aligned} \underline{\kappa}^\epsilon(i, Q) &= \min_{\tilde{q} \in \mathcal{F}_\tau^\epsilon(Q)} \kappa(i; Q, \tilde{q}), \text{ and} \\ \overline{\kappa}^\epsilon(i, Q') &= \max_{\tilde{q} \in \mathcal{F}_\tau^\epsilon(Q)} \kappa(i; Q', \tilde{q}). \end{aligned} \tag{105}$$

We consider the following notations:

$\mathcal{H}_\tau$ and $|\mathcal{H}_\tau|$: respectively, the set and size of all feasible history rating paths that are consistent with the observation at time τ, i.e., π_t.

$\mathbb{P}[h_\tau|Q]$: the probability that rating path h_τ realizes when the intrinsic quality is Q by the time τ.

$h_{N+1:\tau}$: The tail of rating path h_τ after time N.

$\tilde{\pi}_\tau(i)$: the difference between $\pi_\tau(i)$ and the maximum number of rating is over all feasible rating paths up to time N that are consistent with set $\mathcal{H}_\tau$.

$\hat{\pi}_\tau(i)$: the difference between $\pi_\tau(i)$ and the minimum number of rating is over all feasible rating paths up to time N.

With this background, we have

$$
\begin{aligned}
\log \frac{1-q_\tau(Q)}{q_\tau(Q)} &= \log\left(\frac{\sum_{Q'\in\{\mathcal{Q}\setminus Q\}}\sum_{h_\tau\in\mathcal{H}_\tau}\mathbb{P}[h_\tau|Q']}{\sum_{h_\tau\in\mathcal{H}_\tau}\mathbb{P}[h_\tau|Q]}\right)\\
&= \log\left(\frac{\sum_{Q'\in\{\mathcal{Q}\setminus Q\}}\sum_{h_\tau\in\mathcal{H}_\tau}\mathbb{P}[h_N|Q']\mathbb{P}[h_{N+1:\tau}|h_N,Q']}{\sum_{h_\tau\in\mathcal{H}_\tau}\mathbb{P}[h_N|Q]\mathbb{P}[h_{N+1:\tau}|h_N,Q]}\right)\\
&\leq \log\left(\frac{\sum_{Q'\in\{\mathcal{Q}\setminus Q\}}(\max_{h_N}\mathbb{P}[h_N|Q'])\sum_{h_\tau\in\mathcal{H}_\tau}\mathbb{P}[h_{N+1:\tau}|h_N,Q']}{(\min_{h_N}\mathbb{P}[h_N|Q])\sum_{h_\tau\in\mathcal{H}_\tau}\mathbb{P}[h_{N+1:\tau}|h_N,Q]}\right) \qquad (106)\\
&\leq \log\left(\frac{\sum_{Q'\in\{\mathcal{Q}\setminus Q\}}(\max_{h_N}\mathbb{P}[h_N|Q'])\cdot|\mathcal{H}_\tau|\cdot\Pi_{i\in R}(\overline{\kappa}^\epsilon(i,Q'))^{\tilde{\pi}_\tau(i)}}{(\min_{h_N}\mathbb{P}[h_N|Q])\cdot|\mathcal{H}_\tau|\cdot\Pi_{i\in R}(\underline{\kappa}^\epsilon(i,Q))^{\hat{\pi}_\tau(i)}}\right)\\
&\leq \log\left(\frac{(M-1)(\max_{h_N}\mathbb{P}[h_N|\overline{Q}])\Pi_{i\in R}(\overline{\kappa}^\epsilon(i,\overline{Q}))^{\tilde{\pi}_\tau(i)}}{(\min_{h_N}\mathbb{P}[h_N|Q])\Pi_{i\in R}(\underline{\kappa}^\epsilon(i,Q))^{\hat{\pi}_\tau(i)}}\right),
\end{aligned}
$$

where $\overline{Q} = \arg\max_{Q'\in\{\mathcal{Q}\setminus Q\}}(\max_{h_N}\mathbb{P}[h_N|Q'])\Pi_{i\in R}(\overline{\kappa}^\epsilon(i,Q'))^{\tilde{\pi}_\tau(i)}$. The first inequality holds because we take the maximum among all rating paths in the numerator, and minimum in the denominator. The second inequality followed from the definitions in (105). The last inequality followed by replacing Q' with $\overline{Q}$ and using the fact that $|\mathcal{Q}|= M$.

As argued in the proof of Proposition 3.7, when $\overline{\theta} \geq \frac{1}{1-\alpha}+Q-p$, there exists some quality value Q' such that it is not separation divergent from Q. In other words, $\overline{\kappa}^\epsilon(i,Q') > 0$ for some Q' and all i. Therefore, $\overline{\kappa}^\epsilon(i,\overline{Q}) > 0$ from the definition of $\overline{Q}$. Consequently, for all i,

$$
\frac{(\overline{\kappa}^\epsilon(i,\overline{Q}))^{\tilde{\pi}_\tau(i)}}{(\underline{\kappa}^\epsilon(i,Q))^{\hat{\pi}_\tau(i)}} \leq \left(\frac{\overline{\kappa}^\epsilon(i,\overline{Q})}{\underline{\kappa}^\epsilon(i,Q)}\right)^{\pi_\tau(i)}\left(\frac{1}{\overline{\kappa}^\epsilon(i,\overline{Q})\underline{\kappa}^\epsilon(i,Q)}\right)^N, \qquad (107)
$$

where the inequality follows from knowing $\tilde{\pi}_\tau(i) \geq \pi_\tau(i) - N$ and $\hat{\pi}_\tau(i) \leq \pi_\tau(i) + N$. Using this inequality in (106), we find that

$$
\log\frac{1-q_\tau(Q)}{q_\tau(Q)} \leq \log\left((M-1)\cdot\frac{\max_{h_N}\mathbb{P}[h_N|\overline{Q}]}{\min_{h_N}\mathbb{P}[h_N|Q]}\cdot\Pi_i\left(\frac{\overline{\kappa}^\epsilon(i,\overline{Q})}{\underline{\kappa}^\epsilon(i,Q)}\right)^{\pi_\tau(i)}\left(\frac{1}{\overline{\kappa}^\epsilon(i,\overline{Q})\underline{\kappa}^\epsilon(i,Q)}\right)^N\right) \qquad (108)
$$

Recall that $\lim_{\tau\to\infty}\frac{\pi_\tau(i)}{\tau} = \kappa(i;Q,\tilde{Q}_\infty)$. Hence,

$$\begin{aligned}
\lim_{\tau\to\infty}\frac{1}{\tau}\log\frac{1-q_\tau(Q)}{q_\tau(Q)} &\leq \lim_{\tau\to\infty}\frac{1}{\tau}\left(\log\left(\frac{\max_{h_N}\mathbb{P}[h_N|\overline{Q}]}{\min_{h_N}\mathbb{P}[h_N|Q]}\right)+\log\left(\left(\frac{1}{\overline{\kappa}^\epsilon(i,\overline{Q})\underline{\kappa}^\epsilon(i,Q)}\right)^N\right)\right)\\
&\quad+\lim_{\tau\to\infty}\frac{1}{\tau}\log(M-1)+\lim_{\tau\to\infty}\frac{1}{\tau}\log\left(\Pi_i\left(\frac{\overline{\kappa}^\epsilon(i,\overline{Q})}{\underline{\kappa}^\epsilon(i,Q)}\right)^{\pi_\tau(i)}\right)\\
&=\lim_{\tau\to\infty}\sum_i\frac{\pi_\tau(i)}{\tau}\log\left(\frac{\overline{\kappa}^\epsilon(i,\overline{Q})}{\underline{\kappa}^\epsilon(i,Q)}\right)\\
&=\sum_i\kappa(i;Q,\tilde{Q}_\infty)\log\left(\frac{\overline{\kappa}^\epsilon(i,\overline{Q})}{\underline{\kappa}^\epsilon(i,Q)}\right).
\end{aligned} \tag{109}$$

Since the above inequality holds for any ϵ, then $\lim_{\epsilon\to 0}\overline{\kappa}^\epsilon(i,\overline{Q}) = \kappa(i,\overline{Q},\tilde{Q}_\infty(\overline{Q}))$ and $\lim_{\epsilon\to 0}\underline{\kappa}^\epsilon(i,Q) = \kappa(i,Q,\tilde{Q}_\infty(Q))$. Therefore,

$$\begin{aligned}
\lim_{\tau\to\infty}\frac{1}{\tau}\log\frac{1-q_\tau(Q)}{q_\tau(Q)} &\leq \sum_i\kappa(i,Q,\tilde{Q}_\infty(Q))\log\left(\frac{\kappa(i,\overline{Q},\tilde{Q}_\infty(\overline{Q}))}{\kappa(i,Q,\tilde{Q}_\infty(Q))}\right)\\
&=-\sum_i\kappa(i,Q,\tilde{Q}_\infty(Q))\log\left(\frac{\kappa(i,Q,\tilde{Q}_\infty(Q))}{\kappa(i,\overline{Q},\tilde{Q}_\infty(\overline{Q}))}\right)\\
&=-D(\vec{\kappa}(Q,q_\infty(Q)=1)||\vec{\kappa}(\overline{Q},q_\infty(\overline{Q})=1))\\
&\leq\max_{\overline{Q}\in\{\mathcal{Q}\backslash Q\}}-D(\vec{\kappa}(Q,q_\infty(Q)=1)||\vec{\kappa}(\overline{Q},q_\infty(\overline{Q})=1))\\
&=-\min_{\overline{Q}\in\{\mathcal{Q}\backslash Q\}}D(\vec{\kappa}(Q,q_\infty(Q)=1)||\vec{\kappa}(\overline{Q},q_\infty(\overline{Q})=1)).
\end{aligned} \tag{110}$$

This proves the upper bound in the proposition.

To prove the lower bound, change the definitions of $\underline{\kappa}^\epsilon(i,Q)$ and $\overline{\kappa}^\epsilon(i,Q')$ to

$$\begin{aligned}
\overline{\kappa}^\epsilon(i,Q) &= \max_{\tilde{q}\in\mathcal{F}^\epsilon_\tau(Q)}\kappa(i;Q,\tilde{q}),\ \text{and}\\
\underline{\kappa}^\epsilon(i,Q') &= \min_{\tilde{q}\in\mathcal{F}^\epsilon_\tau(Q)}\kappa(i;Q',\tilde{q}).
\end{aligned} \tag{111}$$

Also, let $\underline{Q} = \arg\min_{Q'\in\{\mathcal{Q}\backslash Q\}}(\min_{h_N}\mathbb{P}[h_N|Q'])\Pi_{i=1}(\underline{\kappa}^\epsilon(i,Q'))^{\tilde{\pi}_\tau(i)}$ such that $\underline{\kappa}^\epsilon(i,Q')>0$ for all i. Following a similar argument used to prove the upper bound, we can show that for all $\epsilon>0$,

$$\lim_{\tau\to\infty}\frac{1}{\tau}\log\frac{1-q_\tau(Q)}{q_\tau(Q)} \geq \sum_i\kappa(i,Q)\log\left(\frac{\underline{\kappa}^\epsilon(i,\underline{Q})}{\overline{\kappa}^\epsilon(i,Q)}\right). \tag{112}$$

Since this inequality holds for any $\epsilon > 0$, we have

$$\sum_i \kappa(i, Q, \tilde{Q}_\infty(Q)) \log \left(\frac{\kappa(i, \underline{Q}, \tilde{Q}_\infty(\underline{Q}))}{\kappa(i, Q, \tilde{Q}_\infty(Q))} \right) \leq \lim_{\tau \to \infty} \frac{1}{\tau} \log \frac{1 - q_\tau(Q)}{q_\tau(Q)}, \tag{113}$$

which implies the lower bound in the proposition. □

Bibliography

Acemoglu D, Dahleh MA, Lobel I, Ozdaglar A (2011) Bayesian learning in social networks. *The Review of Economic Studies* 78(4):1201–1236.

Acemoglu D, Makhdoumi A, Malekian A, Ozdaglar A (2022) Learning from reviews: The selection effect and the speed of learning. *Econometrica* 90(6):2857–2899, URL `http://dx.doi.org/https://doi.org/10.3982/ECTA15847`.

Acquisti A, Varian HR (2005) Conditioning prices on purchase history. *Marketing Science* 24(3):367–381.

Aflaki A, Feldman P, Swinney R (2020) Becoming strategic: Endogenous consumer time preferences and multiperiod pricing. *Operations Research* 68(4):1116–1131.

Aflaki A, Swinney R (2021) Inventory integration with rational consumers. *Operations Research* 69(4):1025–1043.

Allon G, Bassamboo A (2011) Buying from the babbling retailer? the impact of availability information on customer behavior. *Management Science* 57(4):713–726.

Allon G, Bassamboo A, Randhawa RS (2012) Price as a signal of product availability: Is it cheap? *Available at SSRN 3393502* .

Amaldoss W, He C (2019) The charm of behavior-based pricing: When consumers' taste is diverse and the consideration set is limited. *Journal of Marketing Research* 56(5):767–790.

Anderson EW, Sullivan MW (1993) The antecedents and consequences of customer satisfaction for firms. *Marketing science* 12(2):125–143.

Angwin J, Parris Jr T, Mattu S (2016) Breaking the black box: When algorithms decide what you pay. propublica.

Anunrojwong J, Iyer K, Lingenbrink D (2023) Persuading risk-conscious agents: A geometric approach. *Operations Research* .

Arieli I, Babichenko Y, Smorodinsky R, Yamashita T (2019) Optimal persuasion via bi-pooling. *Available at SSRN 3511516* .

Aviv Y, Pazgal A (2008) Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing & service operations management* 10(3):339–359.

Aviv Y, Wei MM, Zhang F (2019) Responsive pricing of fashion products: The effects of demand learning and strategic consumer behavior. *Management Science* 65(7):2982–3000.

Aydin G, Ziya S (2009) Personalized dynamic pricing of limited inventories. *Operations Research* 57(6):1523–1531.

Aydinliyim T, Pangburn MS, Rabinovich E (2017) Inventory disclosure in online retailing. *European Journal of Operational Research* 261(1):195–204.

Ban GY, Keskin NB (2021) Personalized dynamic pricing with machine learning: High-dimensional features and heterogeneous elasticity. *Management Science* 67(9):5549–5568.

Banerjee AV (1992) A simple model of herd behavior. *The quarterly journal of economics* 107(3):797–817.

Besanko D, Winston WL (1990) Optimal price skimming by a monopolist facing rational consumers. *Management science* 36(5):555–567.

Besbes O, Scarsini M (2018) On information distortions in online ratings. *Operations Research* 66(3):597–610.

Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy* 100(5):992–1026.

Bimpikis K, Papanastasiou Y (2019) Inducing exploration in service platforms. *Sharing Economy*, 193–216 (Springer).

Cachon GP, Feldman P (2015) Price commitments with strategic consumers: Why it can be optimal to discount more frequently than optimal. *Manufacturing & Service Operations Management* 17(3):399–410.

Cachon GP, Swinney R (2009) Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Science* 55(3):497–511.

Calvo E, Cui R, Wagner L (2020) Disclosing product availability in online retail. *Manufacturing & Service Operations Management* .

Çelen B, Kariv S (2004) Observational learning under imperfect information. *Games and Economic behavior* 47(1):72–86.

Che YK, Hörner J (2018) Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics* 133(2):871–925.

Chen B, Simchi-Levi D, Wang Y, Zhou Y (2021a) Dynamic pricing and inventory control with fixed ordering cost and incomplete demand information. *Management Science* .

Chen N, Nasiry J (2020) Does loss aversion preclude price variation? *Manufacturing & Service Operations Management* 22(2):383–395.

Chen S, Liu Z, Moinzadeh K, Tan Y (2021b) When to broadcast? inventory disclosure policies for online sales of limited inventory. *Inventory Disclosure Policies for Online Sales of Limited Inventory (July 3, 2021)* .

Chen X, Hu P, Hu Z (2017) Efficient algorithms for the dynamic pricing problem with reference price effect. *Management Science* 63(12):4389–4408.

Chen X, Owen Z, Pixton C, Simchi-Levi D (2022) A statistical learning approach to personalization in revenue management. *Management Science* 68(3):1923–1937.

Chen X, Simchi-Levi D, Wang Y (2021c) Privacy-preserving dynamic personalized pricing with demand learning. *Management Science* .

Chevalier JA, Mayzlin D (2006) The effect of word of mouth on sales: Online book reviews. *Journal of marketing research* 43(3):345–354.

Cohen MC, Lobel I, Paes Leme R (2020) Feature-based dynamic pricing. *Management Science* 66(11):4921–4943.

Conitzer V, Taylor CR, Wagman L (2012) Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science* 31(2):277–292.

Cover TM, Thomas JA (2012) *Elements of information theory* (John Wiley & Sons).

Crapis D, Ifrach B, Maglaras C, Scarsini M (2017) Monopoly pricing in the presence of social learning. *Management Science* 63(11):3586–3608.

Cui R, Shin H (2018) Sharing aggregate inventory information with customers: Strategic cross-selling and shortage reduction. *Management Science* 64(1):381–400.

Cui R, Zhang DJ, Bassamboo A (2019) Learning from inventory availability information: Evidence from field experiments on amazon. *Management Science* 65(3):1216–1235.

Dana Jr JD, Petruzzi NC (2001) Note: The newsvendor model with endogenous demand. *Management Science* 47(11):1488–1497.

Deloitte (2017) 2017 global mobile consumer survey: Us edition.

Drakopoulos K, Jain S, Randhawa R (2021) Persuading customers to buy early: The value of personalized information provisioning. *Management Science* 67(2):828–853.

Dubé JP, Misra S (2019) Personalized pricing and customer welfare. *Available at SSRN* 2992257.

Elmachtoub AN, Gupta V, Hamilton ML (2021) The value of personalized pricing. *Management Science* 67(10):6055–6070.

Elmaghraby W, Keskinocak P (2003) Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management science* 49(10):1287–1309.

European Commission (2018) Consumer market study on online market segmentation through personalised pricing/offers in the european union .

Frick M, Iijima R, Ishii Y (2022) Belief Convergence under Misspecified Learning: A Martingale Approach. *The Review of Economic Studies* 90(2):781–814, ISSN 0034-6527, URL `http://dx.doi.org/10.1093/restud/rdac040`.

Fudenberg D, Villas-Boas JM (2006) Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems* 1:377–436.

Gallego G, Van Ryzin G (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management science* 40(8):999–1020.

Garbarino E, Lee OF (2003) Dynamic pricing in internet retail: effects on consumer trust. *Psychology & Marketing* 20(6):495–513.

Garg N, Johari R (2019) Designing optimal binary rating systems. *The 22nd International Conference on Artificial Intelligence and Statistics*, 1930–1939 (PMLR).

Guan X, Wang Y, Yi Z, Chen YJ (2020) Inducing consumer online reviews via disclosure. *Production and Operations Management* 29(8):1956–1971.

Guo Y, Shmaya E (2019) The interval structure of optimal disclosure. *Econometrica* 87(2):653–675.

Haws KL, Bearden WO (2006) Dynamic pricing and consumer fairness perceptions. *Journal of Consumer Research* 33(3):304–311.

Herrera H, Hörner J (2013) Biased social learning. *Games and Economic Behavior* 80:131–146.

Ho YC, Wu J, Tan Y (2017) Disconfirmation effect on online rating behavior: A structural model. *Information Systems Research* 28(3):626–642.

Hu N, Pavlou PA, Zhang J (2006) Can online reviews reveal a product's true quality? empirical findings and analytical modeling of online word-of-mouth communication. *Proceedings of the 7th ACM conference on Electronic commerce*, 324–330.

Hu Z, Chen X, Hu P (2016) Dynamic pricing with gain-seeking reference price effects. *Operations Research* 64(1):150–157.

Ifrach B, Maglaras C, Scarsini M, Zseleva A (2019) Bayesian social learning from consumer reviews. *Operations Research* 67(5):1209–1221.

Jadbabaie A, Molavi P, Sandroni A, Tahbaz-Salehi A (2012) Non-bayesian social learning. *Games and Economic Behavior* 76(1):210–225.

Jing B (2017) Behavior-based pricing, production efficiency, and quality differentiation. *Management Science* 63(7):2365–2376.

Kamenica E, Gentzkow M (2011) Bayesian persuasion. *American Economic Review* 101(6):2590–2615.

Keskin NB, Li Y, Song JS (2022) Data-driven dynamic pricing and ordering with perishable inventory in a changing environment. *Management Science* 68(3):1938–1958.

Kremer M, Mantin B, Ovchinnikov A (2017) Dynamic pricing in the presence of myopic and strategic consumers: Theory and experiment. *Production and Operations Management* 26(1):116–133.

Li J, Granados N, Netessine S (2014) Are consumers strategic? structural estimation from the air-travel industry. *Management Science* 60(9):2114–2137.

Li X, Hitt LM (2008) Self-selection and information role of online product reviews. *Information systems research* 19(4):456–474.

Lingenbrink D, Iyer K (2019) Optimal signaling mechanisms in unobservable queues. *Operations research* 67(5):1397–1416.

Lobel I, Sadler E (2015) Information diffusion in networks through social learning. *Theoretical Economics* 10(3):807–851.

Maglaras C, Scarsini M, Shin D, Vaccari S (2023) Product ranking in the presence of social learning. *Operations Research* 71(4):1136–1153.

Mossel E, Sly A, Tamuz O (2014) Asymptotic learning on bayesian social networks. *Probability Theory and Related Fields* 158(1):127–157.

OECD (2018) Personalised pricing in the digital era–note by the united states. `https://tinyurl.com/29rfsjw8`.

OECD (2020) Personalised pricing in the digital era. `https://tinyurl.com/3wkcjtab`.

Papanastasiou Y, Bakshi N, Savva N (2014) Scarcity strategies under quasi-bayesian social learning. *History* .

Papanastasiou Y, Bimpikis K, Savva N (2018) Crowdsourcing exploration. *Management Science* 64(4):1727–1746.

Papanastasiou Y, Savva N (2017) Dynamic pricing in the presence of social learning and strategic consumers. *Management Science* 63(4):919–939, URL `http://dx.doi.org/10.1287/mnsc.2015.2378`.

Pazgal A, Soberman D (2008) Behavior-based discrimination: Is it a winning play, and if so, when? *Marketing Science* 27(6):977–994.

PowerReviews (2023) The ever-growing power of reviews (2023 edition). URL `https://www.powerreviews.com/research/power-of-reviews-2023/`.

Reinartz WJ, Haucap J, Wiegand N (2018) When customers are–and aren't–ok with personalized prices. *Harvard Business Review* 2018(May 31).

Shin D, Vaccari S, Zeevi A (2023) Dynamic pricing with online reviews. *Management Science* 69(2):824–845, URL `http://dx.doi.org/10.1287/mnsc.2022.4387`.

Smith L, Sørensen P (2000) Pathological outcomes of observational learning. *Econometrica* 68(2):371–398.

Stenzel A, Wolf C, Schmidt P (2020) Pricing for the stars: Dynamic pricing in the presence of rating systems. *Proceedings of the 21st ACM Conference on Economics and Computation*, 273–274.

Streitfeld D (2000) On the web price tags blur: What you pay could depend on who you are. *The Washington Post* (September 27).

Su X (2007) Intertemporal pricing with strategic customer behavior. *Management Science* 53(5):726–741.

Su X, Zhang F (2008) Strategic customer behavior, commitment, and supply chain performance. *Management Science* 54(10):1759–1773.

Talwar A, Jurca R, Faltings B (2007) Understanding user behavior in online feedback reporting. *Proceedings of the 8th ACM Conference on Electronic Commerce*, 134–142 (New York, NY, USA: Association for Computing Machinery), ISBN 9781595936530, URL http://dx.doi.org/10.1145/1250910.1250931.

Taylor CR (2004) Consumer privacy and the market for customer information. *RAND Journal of Economics* 631–650.

Trustpilot (2020) The critical role of reviews in internet trust. URL https://business.trustpilot.com/guides-reports/build-trusted-brand/the-critical-role-of-reviews-in-internet-trust.

US Council of Economic Advisers (2016) Big data and differential pricing. Executive Office of the President of the United States.

Wallheimer B (2018) Are you ready for personalized pricing? https://tinyurl.com/4p64ru4e.

Yin R, Aviv Y, Pazgal A, Tang CS (2009) Optimal markdown pricing: Implications of inventory display formats in the presence of strategic customers. *Management Science* 55(8):1391–1408.

Yu M, Ahn HS, Kapuscinski R (2015) Rationing capacity in advance selling to signal quality. *Management Science* 61(3):560–577.

Yu M, Debo L, Kapuscinski R (2016) Strategic waiting for consumer-generated quality information: Dynamic pricing of new experience goods. *Management Science* 62(2):410–435, URL http://dx.doi.org/10.1287/mnsc.2014.2134.