Essays in Economic Theory and Industrial Organization

Inauguraldissertation

zur Erlangung des Grades eines Doktors
der Wirtschaftswissenschaften

durch die

Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von

Louis Luther

aus Korbach

2025

Dekan:	Prof. Dr. Martin Böse
Erstreferent:	Prof. Dr. Benny Moldovanu
Zweitreferent:	Prof. Dr. Stephan Lauermann

Tag der mündlichen Prüfung:	10. April 2025

Acknowledgments

I would like to express my gratitude to my supervisors, Benny Moldovanu and Stephan Lauermann, whose support, insightful feedback, and immense knowledge have been indispensable throughout this journey. You encouraged me to follow my own path while consistently offering valuable guidance and holding me to the highest standards.

The support and assistance from fellow students and colleagues have played a crucial role in refining my research. I am thankful to the participants of the Bonn Micro Theory Seminar and Stephan Lauermann's lab meeting group. In particular, the many stimulating discussions with Philipp Krähmer, Laurenz Marstaller, Simon Block, and Ege Bilgin have greatly enriched this thesis.

I am deeply indebted to the faculty and staff at the University of Bonn, particularly the Bonn Graduate School of Economics and the Institute for Microeconomics, for creating an environment that fostered both intellectual rigor and personal growth. I would like to thank Silke Kinzig, Andrea Reykers, and Mirjam Beyer for their administrative support. I have benefited from the funding and infrastructure provided by the Bonn Graduate School of Economics and the briq Institute on Behavior & Inequality. I am also grateful to the University of California, Berkeley, for their hospitality.

Finally, I am eternally thankful to my family and friends. To my parents, Gerlind and Jürgen, and my brother, Yannic, for their unconditional love and belief in me. This accomplishment would not have been possible without you.

Contents

Acknowledgments **iii**

Introduction **1**

References 2

1 The Welfare Implications of Product Recommendations **3**

1.1 Introduction 3

1.1.1 Related Literature 6

1.2 The Model 9

1.3 Consumer Search 10

1.4 Pricing Equilibrium 12

1.4.1 Comparative Statics of the Equilibrium Price 15

1.5 Welfare Analysis 17

1.5.1 Consumer Surplus 17

1.5.2 Equilibrium Demand and Firm Profits 19

1.6 Extensions and Variations 21

1.6.1 Micro-foundations of Recommendations 21

1.6.2 Heterogeneous match probability 24

1.6.3 Full Market Coverage 26

1.7 Discussion 27

1.A Appendix 29

1.A.1 Auxiliary Lemmas 29

1.A.2 Proof of Proposition 1.2 30

1.A.3 Proof of Proposition 1.3 and 1.4 33

1.A.4 Proof of Lemma 1.1 34

1.A.5 Proof of Proposition 1.5 35

1.A.6 Proof of Proposition 1.6 36

1.A.7 Proof of Proposition 1.7 37

References 38

2 Competition and Consumer Search with Costly Product Returns **41**

2.1 Introduction 41

2.1.1 Related Literature . . . 44
2.2 The Model . . . 46
2.3 Optimal Consumer Search . . . 47
2.3.1 Asymptotics of Optimal Consumer Search . . . 50
2.3.2 Comparative Statics of Optimal Consumer Search . . . 50
2.4 Price Equilibrium and Welfare . . . 53
2.4.1 Comparative Statics of the Equilibrium Price . . . 55
2.4.2 Welfare Analysis . . . 56
2.5 Extensions and Variations . . . 58
2.5.1 Sequential Learning . . . 58
2.5.2 Product Return Handling and Production Costs . . . 60
2.5.3 What if consumers can keep numerous products? . . . 61
2.6 Discussion . . . 61
2.A Appendix . . . 63
2.A.1 Proof of Remark 2.1 . . . 63
2.A.2 Proof of Proposition 2.1 . . . 63
2.A.3 Proof of Corollary 2.1 . . . 69
2.A.4 Proof of Proposition 2.2 . . . 71
2.A.5 Proof of Proposition 2.3 . . . 72
2.A.6 Proof of Lemma 2.3 . . . 72
2.A.7 Proof of Proposition 2.4 . . . 74
2.A.8 Proof of Proposition 2.5 . . . 74
2.A.9 Proof of Proposition 2.7 . . . 75
2.A.10 Proof of Proposition 2.8 . . . 76
References . . . 78

3 Information Design in Selection Problems **81**
3.1 Introduction . . . 81
3.2 The Model . . . 84
3.3 Decomposing the Problem . . . 86
3.3.1 Decomposition Approach . . . 86
3.3.2 Undominated (p, σ) . . . 88
3.3.3 Characterization of Feasibility . . . 89
3.3.4 Selection and Action Persuasion . . . 90
3.4 Optimal Persuasion and Trade-off . . . 91
3.5 Comparative Statics . . . 94
3.6 Applications . . . 95
3.6.1 Lobbying . . . 95
3.6.2 Advertising . . . 96

3.7 Extensions and Limitations 98
3.7.1 Asymmetric Priors 98
3.7.2 Continuous State Space 99
3.7.3 Asymmetric Selection 101
3.8 Related Literature 101
3.9 Conclusion 103
3.A Appendix 104
3.A.1 Proof of Lemma 3.1 104
3.A.2 Proof of Lemma 3.2 105
3.A.3 Proof of Lemma 3.3 106
3.A.4 Proof of Theorem 3.1 107
3.A.5 Proof of Proposition 3.4 110
3.A.6 Proof of Proposition 3.6 111
3.A.7 Proof of Proposition 3.7 117
3.A.8 Proof of Proposition 3.8 117
References 119

Introduction

This thesis consists of three essays in economic theory and industrial organization.

The first chapter, *The Welfare Implications of Product Recommendations*, studies the implications of informative product recommendations on consumer search, market prices and welfare. Consumers, who are heterogeneous in search costs, search sequentially for firms to discover prices and valuations. An industry-profit-maximizing intermediary recommends a firm to each consumer based on noisy information about consumers' preferences. The findings reveal that recommendations not only enhance the efficiency with which existing consumers find products but also attract high search costs consumers who would otherwise refrain from participating in the search process. These high search costs consumers, who rely solely on recommendations, are less price-sensitive, leading to higher market prices. While high search costs consumers benefit from reduced search efforts, low search costs consumers are harmed by increasing market prices.

The primary contribution of the first chapter is to extend the literature on competitive markets with consumer search, building on the foundational works of Anderson and Renault (1999); Wolinsky (1986), by demonstrating that an intermediary can substantially reshape market structure. While prior research has established that improved consumer information encourages existing consumers to search for fewer products, often leading to higher market prices, my work, to the best of my knowledge, is the first to identify that the participation of previously inactive, less price-sensitive consumers is the key driver behind rising market prices.

The second chapter, *Competition and Consumer Search with Costly Product Returns*, incorporates product returns into a model of price competition and sequential consumer search. Consumers search for firms to discover prices and observable valuations by incurring positive search costs and can return unsatisfactory products by incurring positive and fixed product return costs. The optimal search rule is stationary: Consumers buy a product if and only if the observable valuation is above a threshold and return a product if and only if the total valuation is below some threshold. Importantly, product returns not only incentivize consumers to buy products but also to search for other products instead of buying. Equilibrium prices decrease with rising search costs but show a non-monotonic relationship with return costs, which leads to ambiguous welfare predictions.

This chapter contributes to the understanding of consumer product returns by incorporating product returns into a consumer search model. While most previous research has focused on monopolies and optimal return policies, I examine a model with multiple firms where refund policies are mandated rather than chosen by sellers. A key distinction is that the option to return products incentivizes consumers to continue searching rather than committing to a purchase, a mechanism absent in monopoly settings. This leads to a reversal in findings: when pre-purchase information is less costly to acquire, monopolies experience fewer returns due to reduced uncertainty, whereas in competitive markets, lower information costs increase returns as consumers are encouraged to search further rather than commit to a purchase.

The third chapter, *Information Design in Selection Problems*, studies information design when there is a receiver who selects one out of many alternatives and takes an action, and a sender who transmits information about the viability of alternatives to persuade the receiver to select a favorable alternative and take a favorable action. The main theorem characterizes which distributions of posterior beliefs about the favorable alternative conditional on selection can be induced by some signal structure. This theorem facilitates a decomposition of the multi-dimensional information design problem into a selection persuasion problem and an action persuasion problem. I derive properties of optimal persuasion for special cases, and analyze applications to advertising and lobbying.

The analysis of the third chapter contributes to the literature on Bayesian Persuasion and Information Design, building on the seminal research of Kamenica and Gentzkow (2011); Rayo and Segal (2010), by demonstrating that, in the specific context of selection problems, a multi-dimensional Bayesian persuasion problem can be reduced to a corresponding uni-dimensional problem. This reduction allows the application of existing solution methods to address the problem effectively.

References

Simon P Anderson and Regis Renault. Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics*, pages 719–735, 1999. (document)

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011. (document)

Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of political Economy*, 118(5):949–987, 2010. (document)

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

Chapter 1

The Welfare Implications of Product Recommendations

1.1 Introduction

There has never been a time when consumer data has been so widely available as today. Among other factors, consumer data is valuable because it enables platforms to offer personalized recommendations or advertisements. An Ipsos (2018) study finds that in France, 70% of websites in the sample personalize the ranking of offers. Personalized recommendations are helpful to consumers, since recommendations provide a starting point for consumers to search for products they might want to buy. Precise recommendations significantly shorten the consumers' time and effort required to find suitable products. Consequently, personalized product recommendations have become standard practice on most e-commerce platforms. For instance, CDEI (2020) reports that 35% of purchases on Amazon come from recommendations.

It may seem that consumers benefit from personalized product recommendations, as they enable faster and more effortless discovery of suitable products. However, platforms or intermediaries often use these recommendations to target a broader consumer base, and numerous studies indicate that personalized recommendations significantly alter consumers' search behavior.[1] These changes can have a profound impact on the market structure and demand, ultimately affecting market prices and raising concerns about the broader welfare implications of personalized product recommendations.

This paper aims to study the question of how personalized recommendations of a platform or an intermediary affect consumer search, market prices, and welfare. In particular, do market prices increase or decrease due to precise product recommendations? The main finding of this article is that product recommendation not only helps existing consumer to

[1]For instance, Zhou, Lin, Xiao, and Fang (2024) report that more targeted search results induce consumers to (i) search less, (ii) buy the featured products more often, and (iii) make fewer unplanned purchases.

find suitable products faster but also induces consumers who are unable to find products without recommendations to participate in the search market. These consumers only search for the recommended product, and as a result require a lower net utility threshold to buy a product. Consequently, demand becomes less price-sensitive, and firms set higher prices in equilibrium.

I incorporate product recommendations into a model of price competition and sequential consumer search. In particular, I employ a consumer search model with a continuum of firms and consumers who differ in their search costs and sequentially search for firms until they buy a product or decide to terminate the search by exercising an outside option. Heterogeneous search costs have first been studied by Moraga-González, Sándor, and Wildenbeest (2017) who include heterogeneous search costs into the canonical consumer search framework of Anderson and Renault (1999) and Wolinsky (1986). I depart from the framework of Moraga-González et al. (2017) in two major ways:

First, consumers can encounter firms that are not a match. If a firm is not a match to a consumer, the consumer's match valuation is equal to 0, while a consumer's match valuation is independently drawn from some CDF F if a firm is match. I assume that consumers are initially not only uninformed about valuations and prices but also about which firms are matches. Ex-ante, each firm is a match for a consumer with some prior match probability $\mu \in (0, 1)$.

Second, an intermediary (for instance, an e-commerce platform) observes a noisy signal about consumer-firm matches and recommends a single firm to each consumer with the objective of maximizing industry profits. Consumers observe the intermediary's recommendation, and update their belief about which firms are matches. I show that under certain symmetry conditions on the intermediary's information and recommendation choice, the recommendation problem results in a match probability $\eta \in [\mu, 1]$ for the recommended firm, and an unchanged and independent match probability for other firms. I refer to η as the *recommendation precision* which captures the informativeness of the intermediary's information and recommendation choice.

Proposition 1.1 applies the optimal search rule for independent alternatives derived by Weitzman (1979) to this article's setting. The optimal search rule of a consumer depends on her search costs. High search costs consumers are *inactive* since they do not search at all as any potential benefits of searching are exceeded by the costs of searching for products. Consumers with intermediate search costs search for the recommended firm but not for any other firms as their search costs lie in between the benefits of searching for the recommended firm and searching for any other firm. I refer to these consumers as *passive search consumers* since they only follow the recommendation and do not actively search for other firms. Lastly, consumers with sufficiently small search costs are referred to as *active consumer*s. These consumers sample the recommended firm first but continue to search for other firms if the net benefit from buying from the recommended firm does not exceed the continuation value of searching for other firms.

The key impact of product recommendations is that they induce previously inactive consumers to adopt a passive search strategy. The demand from passive search consumers is inherently less price-sensitive than the demand from active search consumers. Since product recommendations increase the total mass of passive search consumers, they also increase the relative mass of passive search consumers in the market to active search consumers which implies that total demand is less price-sensitive as a result of informative product recommendations, and that market prices increase. This effect is further exacerbated in equilibrium: Since consumers correctly anticipate increasing market prices as a result of product recommendations in any equilibrium, the continuation payoff of searching will no longer outweigh the search costs for some active search consumers. These consumers will instead adopt the passive search strategy which increases the share of passive search consumers in the market, and thus also increases the equilibrium price even further.

Rising market prices naturally raise concerns about the exact welfare ramifications of product recommendations. Informative product recommendations offer the innate advantage of increasing the match probability of the first firm the consumer searches for. This is crucial to passive search consumers since they do not visit any other firms but it is also advantageous to active search consumers since they will search for fewer firms on average before finding a satisfactory product, and thus save on search expenditures. However, depending on the search costs this effect might be offset by higher market prices. Corollary 1.1 and Corollary 1.2 show that product recommendations are harmful to low search costs consumers for which the advantage of searching for fewer firms is small, and beneficial to high search costs consumers.

While product recommendations induce some inactive consumers to participate in the search market, participating consumers are more likely to terminate search and exercise an outside option instead of buying from some firm in the market. Thus, product recommendations increase the extensive margin of demand, but decrease the intensive margin of demand.[2] The resulting effect of product recommendations on equilibrium demand is ambiguous. This could potentially be concerning for firms since market prices could be inefficiently large as a result of product recommendations such that increased market prices might not outweigh decreasing demand.[3] However, Proposition 1.5 shows that firms still profit from recommendations under weak assumptions.

[2] I define the extensive margin of demand (in equilibrium) as the mass of consumers searching for at least one firm, and the intensive margin as the probability that a consumer buys some product conditional on searching for at least some firm. Equilibrium demand is equal to the product of the intensive and extensive margin.

[3] The pricing problem of an individual firm does not internalize that more consumers participate in the search market when prices are low. When sufficiently many consumers rely only on the recommendation to find products and competitive pressure is therefore low, the market price is larger than the price firms would set if firms could commit and credibly collude.

From a consumer protection point of view, personalized product recommendations imply a trade-off. Precise product recommendations save search costs since consumers find suitable products faster, but they also increase market prices. This result indicates that the effects of personalized product recommendations on consumer welfare can be viewed similarly to the effects of price discrimination. If sellers set personalized prices, there is also a trade-off: Consumers in low-value market segments benefit from low prices while high prices harm consumers in high-value market segments, and the effect on consumer surplus is generally ambiguous. A similar trade-off arises with product recommendations. Consumers with low search costs are harmed since the value of a recommendation is small and market prices rise, but high search costs consumers profit since the value of a product recommendation is large when search costs are large. Although the usage of personalized prices has long been heavily regulated in the EU and US, the regulation of personalized recommendations based on consumer data has only recently gained policymakers' attention.[4]

The welfare effects of recommendations are different from the welfare effects of price discrimination in a crucial detail. If we assume that low-income consumers are more likely to have small search costs due to lower opportunity costs and more likely to have a low willingness to pay in the context of price discrimination, then product recommendations are beneficial to high-income consumers and harmful to low-income consumers which is diametric to the welfare effect of price discrimination. Therefore, for a policy-maker interested in redistributive policies, this article might serve as a cautionary tale on the welfare consequences of personalized product recommendations.

The paper is structured as follows: In the remainder of the Introduction, I discuss the related literature. In Section 1.2, the model is formalized. In Section 1.3, I apply the rule of Weitzman (1979) to derive the optimal consumer search rule with product recommendations. Section 1.4 solves the firms' pricing problem, and discusses properties of the equilibrium price. In Section 1.5, I analyze the effects of product recommendations on consumers' and firms' welfare, and in Section 1.6, some extensions are discussed. All proofs omitted in the main text can be found in the appendix.

1.1.1 Related Literature

This article builds upon the literature on consumer search, where consumers search to discover both price and match valuations. The canonical contributions in this literature are Wolinsky (1986) and Anderson and Renault (1999), who apply the optimal sequential search rule of Weitzman (1979) to a model of price competition with heterogeneous sellers. Product heterogeneity avoids the well-known Diamond (1971) Paradox and gives rise to

[4] Di Noia, Tintarev, Fatourou, and Schedl (2022); Schwemer (2021) discuss recent attempts to regulate recommender systems. De Streel and Jacques (2019) provide a comprehensive overview of EU law regarding personalized pricing.

a model of monopolistic competition.

In this paper's model, each consumer receives a product recommendation, which is indicative of the consumers' valuation for the recommended product. If the intermediary has no additional information about consumers' tastes, the product recommendations are uninformative, and as a result, this paper's model reduces to the model of Wolinsky (1986) with consumers who are heterogeneous in search costs. This framework has been studied by Moraga-González et al. (2017) who provide conditions for the existence of a symmetric and pure equilibrium. However, the log-concavity conditions from Moraga-González et al. (2017) are not sufficient to guarantee existence in my work.

My work is also closely related to the small literature on the effects of prominence on consumer search because it is optimal for a consumer to sample the recommended firm first. Armstrong, Vickers, and Zhou (2009) fix a prominent firm which is sampled first by all consumers. In the market equilibrium, the prominent firm sets a lower price than all of its competitors since it faces more elastic demand as the proportion of returning consumers to all consumers sampling a firm is lower for the prominent firm. Unlike Armstrong, Vickers, and Zhou (2009), prominence is endogenously determined in my work via the recommendation of an intermediary. The recommendation is informative about a consumer's valuation for the prominent firm such that a consumer has an incentive to sample the recommended firm first, even if all firms charge the same price. If the environment is symmetric (firms are ex-ante identical), each firm has the same probability of being prominent since the intermediary recommends different firms to different consumers. Consequently, contrary to Armstrong, Vickers, and Zhou (2009), an equilibrium in which all firms charge the same price still exists.

More closely related are contributions that consider endogenous prominence. Armstrong and Zhou (2011) consider models in which firms can pay for prominence. Garcia and Shelegia (2018) consider a duopoly model in which consumers can observe the purchase decision of predecessors. Since valuations among consumers are assumed to be positively correlated, the purchase decisions of other consumers are informative about the consumers' own tastes. Thus, the purchase decisions of predecessors serve a similar role as the product recommendation of an intermediary. In Garcia and Shelegia (2018), as well as in this article, consumers optimally start their search by following their predecessors or the intermediary's recommendation, respectively. In the model of Garcia and Shelegia (2018), this emulation effect decreases the equilibrium price since firms have an incentive to set low prices in order to increase the share of consumers buying their product, which in turn also increases the share of consumers starting their search by sampling this firm in the next period.

I assume that the intermediary's recommendation is solely based on match probabilities rather than on prices, such that this effect is absent in my model. However, similar to Garcia and Shelegia (2018), consumers receive information before they search which induces some consumers to free-ride on the information, and as a result, consumers search for fewer firms on average and are less price-sensitive. This learning effect unambiguously increases the equilibrium price in the setup of Garcia and Shelegia (2018).

The central trade-off in my work is related to a well-known insight: Improved informedness of consumers induces higher market prices. This is a common observation in the literature of privacy, and information acquisition in multi-product models:

Anderson and Renault (2000) employ a sequential search model and show that consumers who are ex-ante informed about match values and only search to discover prices exert a negative externality on uninformed consumers who search to discover both match values and prices. The reason is that the demand of informed consumers is less elastic, which drives up equilibrium prices. A similar effect is present in my model: As the recommendation of the intermediary increases in precision, consumers who adopt a passive search strategy and solely rely on the intermediary's recommendation are less price-sensitive and exert a negative externality on other consumers.

Rhodes and Zhou (2024) study a sequential search model in which consumers can perfectly reveal their match valuations to a platform recommending the firm with the highest match valuation. Similar to my work, consumers are endogenously separated into different groups. In Rhodes and Zhou (2024), consumers with small privacy costs share data and consumers with large privacy costs do not. Like the passive search consumers in my model, consumers who share data are less price-sensitive due to recommendations and induce a negative externality by driving up the equilibrium price. Minaev (2021) utilizes a structural model and data from the hotel booking platform Expedia to estimate that personalized rankings decrease consumers' search expenditure by 1.1% but decrease consumers' utility by 3.6% due to increased prices.

Other works which find trade-offs between recommendations and prices include Ichihashi (2020) and Hidir and Vellodi (2021) who study the data sharing incentives of consumers facing a monopolist selling several varieties of a product. Data sharing consumers can be provided with the best-fitting variety of the product but pay higher prices.

To the best of my knowledge, my work is the first to identify that, as a result of product recommendations, increased participation by consumers with high search costs is a key factor influencing market prices.

I assume that consumers differ in the set of firms they can or want to buy from. The same assumption is typical in models of oligopoly pricing following Varian (1980) and Rosenthal (1980). Unlike this literature, I assume that at first consumers do not initially know which firms are a match, and the match values for firms' products in the consideration sets are random and independent. Furthermore, consumers receive an informative recommendation about which firms are matches, and consumers search for firms to perfectly reveal the match values. I assume that consumer-firm matches are symmetric and independent, which is also a common assumption in the oligopoly pricing literature (for instance, Ireland (1993); McAfee (1994); Spiegler (2006); Szech (2011)).[5]

1.2 The Model

I analyze a market with a continuum of firms on the supply side. Each firm $i \in [0,1]$ chooses a price p_i for its product which is produced at constant marginal costs normalized to 0.

On the demand side, there is a unit mass of consumers. Each consumer is initially uninformed about match valuations and prices $(u_i, p_i)_{i\in[0,1]}$, but can discover valuation and price of a firm by incurring search costs s. Recall is free but not necessary. Consumers are heterogeneous in search costs where $G(s)$ denotes the mass of consumers with search costs smaller than $s \in \mathbb{R}$. For most of the article, I assume that search costs are distributed uniformly on $[0, \bar{s}]$.

Furthermore, the match valuation u_i of a consumer for the product of firm $i \in [0,1]$ is assumed to be equal to 0 with probability $\mu \in (0,1)$ and otherwise drawn according to an at-least twice differentiable CDF F with support on $[0, \bar{v}]$. If $u_i > 0$, we say that firm i is a match, and it is not a match if $u_i = 0$. Each match valuation is drawn independently from other valuations, and from the search costs. We impose two regularity conditions on the CDF of match valuations:

Assumption 1.1. *(i) $1 - F$ is log-concave, (ii) $\Pi^m(p) := p \cdot (1 - F(p))$ is concave on $[0, \bar{v}]$.*[6]

There exists a single intermediary, for instance a platform. The intermediary has noisy information about which firms match to consumers and can recommend a single firm to a consumer. For instance, the intermediary might observe a set of consumer characteristics that are correlated with which firms are matches. For now, we take a reduced-form modeling approach: I assume that the recommended firm is a match with probability $\eta \in [\mu, 1]$, other firms match probabilities are not affected and every firm is recommended to a single consumer.

I show that this reduced-form approach is the result of an intermediary's strategic choice of a recommendation strategy under symmetry conditions in Section 1.6. The *recommendation precision* η captures how informative the intermediary's recommendations are. If the recommendation precision is equal to the prior match probability ($\eta = \mu$), then the recommended firm is no different from a randomly selected firm. On the other hand, recommendations are perfectly informative in the sense that consumers can be sure that the recommended firm is always a match if $\eta = 1$.

[5] Armstrong and Vickers (2022) provide an excellent overview of the literature and unifying results.

[6] Π^m is concave if and only if $-p\frac{f'(p)}{f(p)} \leqslant 2$ which is a common assumption in the literature on price discrimination (see for instance Aguirre, Cowan, and Vickers (2010) or Preuss (2023))

The timing of the game is as follows. First firms set prices simultaneously. Then, the recommendations are realized, and the consumers observe which firm is recommended. Finally, each consumer sequentially searches for firms.[7] The game ends if a consumer buys a product or the outside option is exercised.

If a consumer with search costs s buys a product that she values at u at a price p after searching for $m \in \mathbb{N}$ firms, her realized payoff is $u - p - m \cdot s$. If a consumer exercises her outside option after searching for $m \in \mathbb{N}$ firms, her realized payoff is $-m \cdot s$. Each consumer maximizes her expected payoff. Firms maximize expected profits. The profit of firm i if mass α_i of consumers buy the product of firm i is $\alpha_i p_i$.

The solution concept is symmetric and pure Perfect Bayesian Nash Equilibrium with two common restrictions: First, a positive mass of consumers sample at least one firm.[8] Second, a consumer does not revise her belief about prices of firms $j \neq i$, if she encounters firm i, for which the actual price p_i does not match the expected price p.

1.3 Consumer Search

First, we derive the optimal consumer search rule of a consumer with search costs $s \in [0, \bar{s}]$. Since we are focusing on pure and symmetric equilibria, suppose that each consumer expects that firms charge some price $p \in \mathbb{R}_+$.

Weitzman (1979) shows that if valuations of alternatives are independently distributed, and the valuation of alternative i is drawn from some CDF H_i, then optimal consumer search can be characterized by a simple reservation value based index. The reservation value $w_i(s)$ of alternative i solves

$$s = \int_{w_i(s)}^{\bar{v}} (u - w_i(s)) dH_i(u).$$

The search rule of Weitzman (1979) states that it is optimal to search for alternatives in decreasing order of reservation values, and to stop searching and take an alternative if $u_i - p_i \geqslant w_j(s) - p$ for all alternatives j that have not yet been sampled. Furthermore, if all remaining net reservation values $w_j(s) - p$ are negative, the consumer takes the alternative with the highest net utility $u_i - p_i$ if it is positive and otherwise exercises the outside option.

In this paper's model, the valuation for a firm's product is drawn from F if the firm is a match and otherwise it is 0. The recommended firm is a match with probability η while any other firm is a match with probability μ. Therefore, the valuation for the

[7] After sampling a firm, the consumer decides whether to sample another firm, buy from any of the sampled firms, or exercise the outside option which yields a payoff of 0. The search stage mirrors Weitzman (1979) and Wolinsky (1986). I refrain from introducing rigorous definitions of searching strategies.

[8] This excludes equilibria in which every firm sets high prices (for instance, $p > \bar{v}$) and the consumers never samples any firms because they expect prices to be too high.

recommended firm is drawn from $\eta F + (1-\eta)\delta_0$ where δ_0 denotes the CDF associated with the Dirac measure on 0, and the valuation for other firms' products are drawn from $\mu F + (1-\mu)\delta_0$. Letting

$$\begin{aligned} W(x) &:= \int_x^{\bar{v}} (u-x)dF(u) \\ &= \int_x^{\bar{v}} (1-F(u))du \end{aligned}$$

denote the incremental benefit of a match when the price is $x \in \mathbb{R}_+$, the reservation value for the recommended firm $r_+(s)$ solves $\frac{s}{\eta} = W(r_+(s))$, and the reservation value for any other firm satisfies $\frac{s}{\mu} = W(r(s))$. Note that the net reservation value for the recommended firm is positive if and only if

$$r_+(s) \geqslant p \quad \Longleftrightarrow \quad s \leqslant s^*_{pas} := \eta W(p)$$

and the net reservation value for other firms is positive if and only if

$$r(s) \geqslant p \quad \Longleftrightarrow \quad s \leqslant s^*_{act} := \mu W(p).$$

Therefore, the optimal consumer search rule of Weitzman (1979) applied to the setting of this paper yields the following result.

Proposition 1.1. *Suppose that consumers expect prices $p_i = p$ for all firms $i \in [0,1]$. The optimal consumer search rule satisfies:*

(i) Consumers with search costs $s \leqslant \mu W(p)$ search for the recommended product first. They buy if $v_i - p_i \geqslant r(s) - p$. Otherwise, they search for other firms until they find some firm for which $v_i - p_i \geqslant r(s) - p$.

(ii) Consumers with search costs $s \in (\mu W(p), \eta W(p)]$ search for the recommended product and buy if $v_i \geqslant p_i$. Otherwise, they exercise the outside option.

(iii) Consumers with search costs $s > \eta W(p)$ exercise the outside option immediately without searching.

The optimal consumer search rule is illustrated in Figure 1.1. Consumers are endogenously separated into three groups: Firstly, *active search* consumers who start with the recommended firm and search for firms until the net payoff of some firm exceeds the net reservation value $r(s) - p$. Secondly, *passive search* consumers who only search for the recommended product and buy this product if the match valuation exceeds the price. Thirdly, *inactive consumers* who do not search at all.

I assume that $\bar{s} > \eta \mathbb{E}_F[v]$ which implies that there always exist inactive search consumers such that product recommendations increase the extensive margin of search. In the existing literature with homogeneous search costs, it is assumed that every consumer

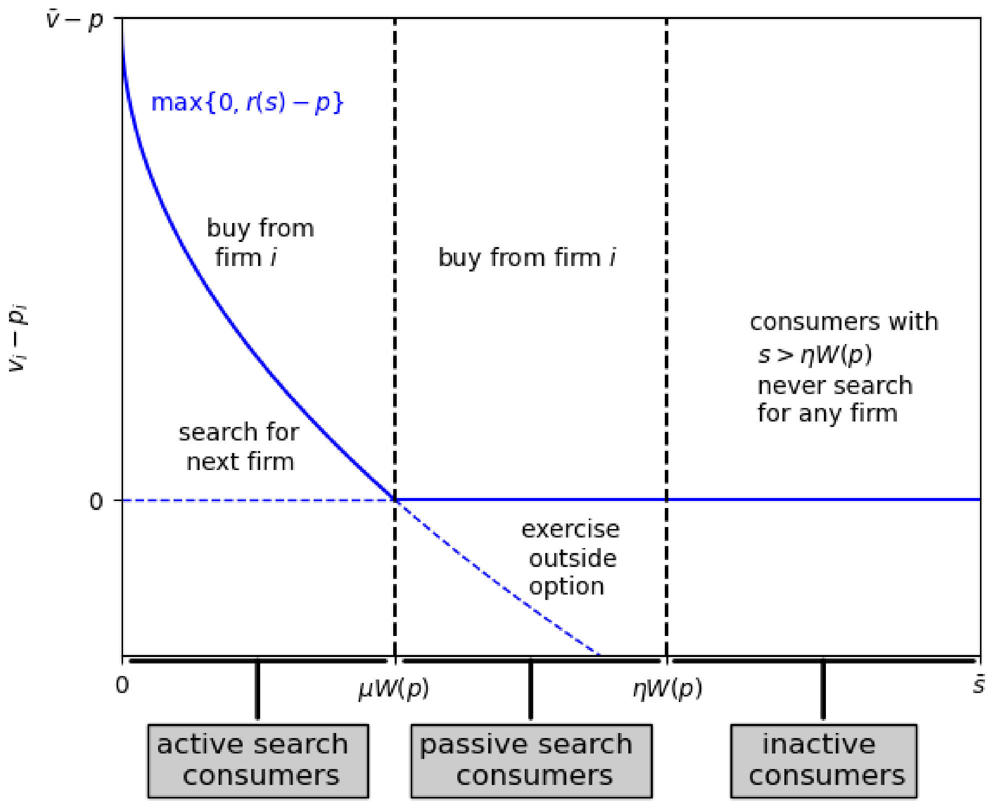

Figure 1.1: Consumers' search and purchasing rule depends on search costs.

is an active search consumer. In the literature with heterogeneous search costs, there exist typically inactive consumers with high search costs as well as active search consumers with low search costs, but no passive search consumers.

1.4 Pricing Equilibrium

It is useful to derive the masses of active search and passive search consumers in the market. Given the optimal consumer search rule (Proposition 1.1), consumers with search costs in the interval $[\mu W(p), \eta W(p)]$ passively follow the recommendation, and consumers with search costs in $[0, \mu W(p)]$ actively search for firms. Thus, let

$$\gamma_{pas}(p, \eta) = G(\eta W(p)) - G(\mu W(p))$$

and

$$\gamma_{act}(p, \eta) = G(\mu W(p))$$

denote the mass of passive search consumers and active search consumers, respectively. One very important property of the uniform distribution is that the relative mass of passive consumers to active consumers is constant in the anticipated price p. That is, for the uniform distribution we obtain

$$\frac{\gamma_{pas}(p,\eta)}{\gamma_{act}(p,\eta)+\gamma_{pas}(p,\eta)} = 1-\frac{\mu}{\eta}.$$

This property ensures a unique equilibrium. To see why, suppose that the relative mass of passive consumers is increasing in the expected price. If consumers expect a large price, a large share of consumers will only search passively for the recommended firm which induces firms to set a large price and confirm the initial expectation. On the other hand, if consumers expect small market price, many consumers will actively search for firms which induces low market prices and again confirms the initial belief. To avoid the issue of multiple equilibria, we assume that search costs are uniformly distributed such that the relative mass of passive consumers to active consumers does not depend on price beliefs.

Since passive search consumers search only for the recommended firm i, the recommended firm is a match with probability η and they buy from the recommended firm if and only if $v_i \geqslant p_i$, the demand induced by passive search consumers is

$$\begin{aligned} D_{pas}(p_i,p) &= \gamma_{pas}(p,\eta)\cdot\eta\cdot(1-F(p_i)) \\ &= \frac{1}{\bar{s}}(\eta-\mu)W(p)\cdot\eta(1-F(p_i)) \end{aligned}$$

where the second lines follows from the uniform distribution.

Active search consumers buy from a firm if and only if $v_i - p_i \geqslant r(s) - p$, and otherwise search for the next firm. Suppose that some firm sets price $p_i \in \mathbb{R}_+$ while all other firms set prices $p_j = p$. The probability that an active search consumer with type $s \leqslant \mu W(p)$ is buying the product of firm i is

$$\begin{aligned} D_{act}(p_i,p;s) = &\underbrace{\rho(1-F(r(s)+p_i-p))}_{\text{buying from recommended product}} \\ &+ \underbrace{\sum_{m=1}^{\infty}\left[1-\rho(1-F(r(s))\right]\left(1-\mu(1-F(r(s))\right)^{m-1}\mu(1-F(r(s)+p_i-p))}_{\text{buying from firm } i \text{ at position } m \text{ in the search order}} \\ &= \frac{1-F(r(s)+p_i-p)}{1-F(r(s))}. \end{aligned}$$

Thus, the aggregated demand coming from active search consumers is then given by

$$\begin{aligned} D_{act}(p_i,p) &= \int_0^{\mu W(p)} D_{act}(p_i,p;s)g(s)ds \\ &= \frac{1}{\bar{s}}\int_0^{\mu W(p)} \frac{1-F(r(s)+p_i-p)}{1-F(r(s))}ds. \end{aligned}$$

Substituting $\tilde{r} = r(s)$ yields the following for the demand induced by active search consumers.

$$D_{act}(p_i, p) = \frac{\mu}{\bar{s}} \int_p^{\bar{v}} (1 - F(\tilde{r} + p_i - p)) d\tilde{r}.$$

Letting $D(p_i, p) = D_{act}(p_i, p) + D_{pas}(p_i, p)$ denote the total demand, we can state the firm's profit maximization problem as follows

$$\max_{p_i} D(p_i, p) p_i = \frac{p_i}{\bar{s}} \cdot \left((\eta - \mu)\eta W(p)(1 - F(p_i)) + \mu \int_p^{\bar{v}} (1 - F(\tilde{r} + p_i - p)) d\tilde{r} \right).$$

The first-order condition of profit maximization yields $p_i = \frac{D(p_i,p)}{-D(p_i,p)/\partial p_i}$. In every symmetric equilibrium, it must be satisfied that firm i has no incentive to deviate from the anticipated price, i.e. $p_i = p$. Thus, a candidate equilibrium price p^* satisfies

$$p^* = \frac{D(p^*, p^*)}{-\frac{\partial D(p^*,p^*)}{\partial p_i}} \equiv \frac{\eta(\eta - \mu)W(p^*)(1 - F(p^*)) + \mu W(p^*)}{\eta(\eta - \mu)W(p^*)f(p^*) + \mu(1 - F(p^*))}. \tag{1.1}$$

Proposition 1.2. *There exists a unique equilibrium price p^* defined by* (1.1).

The proof of Proposition 1.2 is contained in the appendix. I show that the equilibrium candidate condition $p = D(p,p)/(-\frac{\partial D(p,p)}{\partial p_i})$ has a unique solution p^* and $p^* \in [p^a, p^m]$ where p^m is the monopoly price and p^a is the equilibrium price without recommendations. Next, I verify that the firm's profit function $D(p_i, p)p_i$ is concave in p_i which implies that the unique equilibrium candidate is indeed the unique equilibrium. Usually, the literature relies on log-concavity of the demand function in p_i which implies that the profit function is quasi-concave such that the first-order condition must constitute a maximum. However, this approach fails in this paper's setting. Although the demand from passive search consumers is log-concave by assumption and an aggregation theorem of Prékopa (1971) implies that the (aggregated) demand from active search consumers is log-concave, as pointed out by Moraga-González et al. (2017), total demand is not log-concave in general, since the sum of log-concave functions is not necessarily log-concave. Instead, I rely on Assumption 1.1 which implies that the profit functions for active search as well as passive search consumers is concave which is preserved under summation.

For an interpretation, note that with uninformative recommendations ($\eta = \mu$), this paper's model reduces to sequential search with consumers who are heterogeneous in their search costs. As Moraga-González et al. (2017) show, if search costs are uniformly distributed, there exists a unique equilibrium price p^a satisfying $p^a = \frac{W(p^a)}{1-F(p^a)}$, which is the same condition as (1.1) for $\eta = \mu$. When recommendations are informative, some consumers will only search for the recommended product. These consumers essentially act as if the market was a monopoly with the recommended firm as the sole supplier. Therefore, as the recommendation precision η increases, the market price shifts towards the monopoly price p^m satisfying $p^m = \frac{1-F(p^m)}{f(p^m)}$. As a limit case, if the match probability without recommendations μ is small ($\mu \to 0$), almost all consumers rely on the intermediary's recommendation to find and buy a product, and as a result the equilibrium price equals the monopoly price.

1.4.1 Comparative Statics of the Equilibrium Price

How do informative product recommendations affect market prices? I show that the unique equilibrium price defined by (1.1) is increasing in the recommendation precision η, i.e. informative product recommendations drive market prices up. Intuitively, there are two reasons why this is the case: First, passive search consumers have less elastic demand than active search consumers, and second the relative mass of passive search to active search consumers increases in the recommendation precision.

To see the first point recall that passive search consumers buy (from the recommended firm) if and only if $v_i \geqslant p_i$. Therefore, in any equilibrium the probability that a passive search consumer buys a product is $\eta(1-F(p^*))$ and marginally increasing the price yields a marginal demand reduction of $\eta f(p^*)$. Thus, the equilibrium point price elasticity of a passive search consumer E_{pas} is given by

$$E_{pas} = \frac{f(p^*)}{1-F(p^*)} p^*.$$

On the other hand, an active search consumer of type $s \leqslant \mu W(p^*)$ buys from a sampled firm's product if and only if $v_i \geqslant r(s) + p_i - p^*$. Thus, the equilibrium point price elasticity $E_{act}(s)$ of an active search consumer with search costs s is given by

$$E_{act}(s) = \frac{f(r(s))}{1-F(r(s))} p^*.$$

Log-concavity of $1-F$ is satisfied if and only if the hazard rate $f(\cdot)/(1-F(\cdot))$ is increasing, and by definition of $r(s)$ we have $r(s) \geqslant p^*$ for all $s \leqslant \mu W(p^*)$. Hence, $E_{act}(s) \geqslant E_{pas}$, i.e. demand from active search consumers is more elastic than demand from passive search consumers.

To see the second point, note that when product recommendations become more informative, some inactive consumers find it worthwhile to search for the recommended product. This directly increases the relative share of passive search consumers to active search consumers. In equilibrium, consumers correctly anticipate that the prices increase and therefore the value of searching beyond the recommended firm of some active search consumers will be no longer positive. These active search consumers switch to the passive search strategy which additionally increases the relative share of the passive search consumers and multiplies the effect on the equilibrium price.

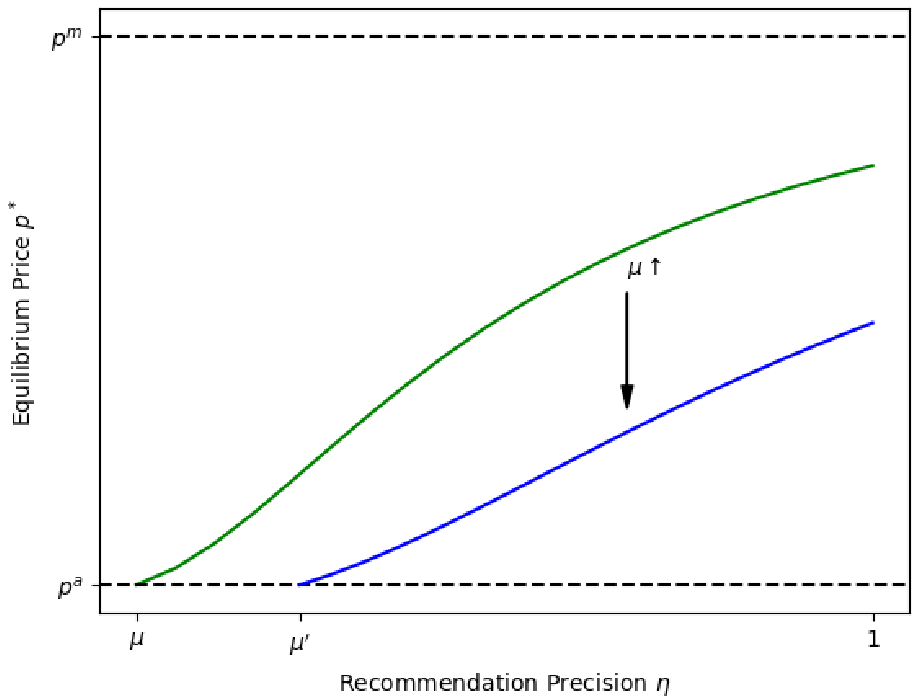

Figure 1.2: The equilibrium price p^* for two different match probabilities μ and μ' as a function of the recommendation precision η.

Proposition 1.3. *The equilibrium price p^* is increasing in the recommendation precision η.*

Figure 1.2 plots the equilibrium price p^* as function of the recommendation precision ρ. When product recommendations are uninformative ($\eta = \mu$), the equilibrium price p^* equals p^a, and increasing the recommendation precision shifts the equilibrium price towards the monopoly price p^m.

Figure 1.2 also indicates that the equilibrium price is decreasing in the prior match probability μ. This seems to be unintuitive at first, since firms can condition on a match and as result the price without recommendations p^a as well as the monopoly price p^m are independent of μ. However, note that if the prior match probability is small, only very few consumers with small search costs find it optimal to search beyond the recommended firm. In other words, the relative share of passive search consumers to active search consumers is decreasing in the prior match probability which is diametric to the effect of a larger recommendation precision.

Proposition 1.4. *The equilibrium price p^* is decreasing in the prior match probability μ.*

1.5 Welfare Analysis

1.5.1 Consumer Surplus

Given that the equilibrium price is increasing in the recommendation precision η, it is immediate that product recommendation might be harmful to at least some consumers.

To clarify, let $p^* \equiv p^*(\eta, \mu)$ denote the equilibrium price when the recommendation precision is given by η and the match probability is μ. As a benchmark, without recommendation ($\eta = \mu$) we know that consumers are active if and only if $s \leqslant s^*_{act}(\mu, \mu) := \mu W(p^*(\mu, \mu))$ and the consumer surplus of type s is simply given by the difference of the reservation value and the equilibrium price, $r(s) - p^*(\mu, \mu)$. Consumers with search costs above $s^*_{act}(\eta, \mu)$ do not search and consequently their consumer surplus is equal to 0.

Now suppose, we have informative recommendations ($\eta > \mu$). Active search consumers with search costs $s \leqslant s^*_{act}(\eta, \mu) = \mu W(p^*(\eta, \mu))$ search first for the recommended firm. Either they buy from the recommended firm, or if they do not buy from the recommended firm, they search for other firms and their continuation payoff is again equal to the net reservation value $r(s) - p^*(\eta, \mu)$. Therefore, if $s \leqslant s^*_{act}(\eta, \mu)$ the consumer surplus is given by

$$\begin{aligned} U_{act}(s; (\eta, \mu)) &= -s + \underbrace{\eta \cdot \int_{r(s)}^{\bar{v}} (u - p^*(\eta, \mu)) dF(u)}_{\text{buying from recommended firm}} + [\eta F(r(s)) + (1 - \eta)] \left(\underbrace{r(s) - p^*(\eta, \mu)}_{\text{buying from other firm}} \right) \\ &= r(s) - p^*(\eta, \mu) + (\eta - \mu) W(r(s)). \end{aligned}$$

Hence, the consumer surplus of an active consumer is given by the net reservation value of searching for firm $r(s) - p^*(\eta, \mu)$ plus $(\eta - \mu) W(r(s))$, which reflects the benefit of receiving a recommendation.

Consumers with search costs $s \in [s^*_{act}(\eta, \mu), s^*_{pas}(\eta, \mu)]$ where $s^*_{pas}(\eta, \mu) = \eta W(p^*(\eta, \mu))$ adopt the passive search strategy, i.e. they only search for the recommended firm and buy from the recommended firm if and only if $v_i \geqslant p^*(\eta, \mu)$. Therefore, the consumer surplus of passive search consumers is given by

$$\begin{aligned} U_{pas}(s; (\eta, \mu)) &= -s + \eta \cdot \int_{p^*(\eta, \mu)}^{\bar{v}} (v_i - p^*(\eta, \mu)) f(v_i) dv_i \\ &= -s + \eta W(p^*(\eta, \mu)). \end{aligned}$$

Now we define the consumer surplus of type $s \in [0, \bar{s}]$ as

$$U(s;(\eta,\mu)) = \begin{cases} U_{act}(s;(\eta,\mu)) & \text{if } s \in [0,\ s^*_{act}(\eta,\mu)] \\ U_{pas}(s;(\eta,\mu)) & \text{if } s \in [s^*_{act}(\eta,\mu),\ s^*_{pas}(\eta,\mu)] \\ 0 & \text{if } s \in [s^*_{pas}(\eta,\mu), \bar{s}] \end{cases}$$

The first result shows that indeed low search consumers are harmed by product recommendations.

Corollary 1.1. *There exists search costs s' such that recommendations are harmful to all consumers with search costs $s \leqslant s'$. That is, $U(s;(\eta,\mu))$ is decreasing in η for all $s \leqslant s'$.*

Proof. Note that

$$\begin{aligned} \lim_{s\to 0} \frac{\partial U(s;(\eta,\mu))}{\partial \eta} &= \lim_{s\to 0}\left(W(r(s)) - \frac{\partial p^*(\eta,\mu)}{\partial \eta}\right) \\ &= W(\bar{v}) - \frac{\partial p^*(\eta,\mu)}{\partial \eta} \\ &= -\frac{\partial p^*(\eta,\mu)}{\partial \eta} \end{aligned}$$

which is negative by Proposition 1.3. Therefore, there exists some search costs s' such that $\frac{\partial U(s;(\eta,\mu))}{\partial \eta} < 0$ for all $s \leqslant s'$. □

We can see that for active search consumers, there exists the price effect $\frac{\partial p^*(\eta,\mu)}{\partial \eta} > 0$ of product recommendations which decreases consumer surplus, as well as an effect described by $W(r(s))$ which increases consumer surplus by enabling active search consumers to find a match faster, and thereby reduce the expected number of sampled firms. However, for consumers with small search costs the magnitude of this effect is small and converges to 0 as search costs converge to 0, while the price effect harms all consumers uniformly.

Since low search costs consumers are harmed by product recommendations, it is natural to ask if high search costs consumers profit from product recommendations. This is the case if the positive effects of recommendations due to the increased match probability outweighs the rising equilibrium price. We impose the following condition that ensures that the slope of the equilibrium price in the recommendation precision η is limited.

Assumption 1.2. Let $\phi(x) := 1 - F(x) - f(x)\frac{W(x)}{1-F(x)}$. *Suppose* $\phi(p^*) \leqslant \frac{\mu}{\eta} \cdot \frac{1}{2\eta-\mu}$.

To interpret, note that $\phi(x) \geqslant 0$ for all x is a necessary and sufficient condition for log-concavity of $W(x) = \int_x^1 (1-F(z))dz$. Thus, Assumption 1.2 imposes that W is not too log-concave at p^*, in the sense that $\phi(p^*)$ is bounded above by the term $\frac{\mu}{\eta} \cdot \frac{1}{2\eta-\mu} \in [\frac{\mu}{2-\mu}, \frac{1}{\mu}]$. Assumption 1.2 is satisfied by many commonly studied distributions including the uniform distribution provided that the prior match probability μ is not too small.

Lemma 1.1. *Suppose that Assumption 1.2 is satisfied. Then,*

$$\frac{\partial p^*}{\partial \eta} < \frac{1}{\eta} \cdot \frac{W(p^*)}{1 - F(p^*)}.$$

This result immediately implies that the increase in market prices is sufficiently small such that higher recommendation precision benefits every passive search consumer and more consumers adopt the passive search strategy. In particular, consumers with search costs $s = s^*_{pas}(\eta, \mu)$, who are indifferent between adopting the passive search strategy and not searching at all, adopt the passive search strategy when the recommendation precision η increases. Additionally, consumers anticipate higher prices in equilibrium such that consumers with search costs $s = s^*_{act}(\eta, \mu)$, who are indifferent between adopting the active search strategy and adopting the passive search strategy, adopt the passive search strategy.

Corollary 1.2. *Suppose that Assumption 1.2 is satisfied. Then, (i) $U_{pas}(s; (\eta, \mu))$ is increasing in η and (ii) the mass of passive search consumers is increasing in η.*

Proof. (i) It is straightforward to show that $U_{pas}(s; (\eta, \mu)) = -s + \eta W(p^*(\eta, \mu))$ is increasing in η if and only if

$$\frac{\partial p^*}{\partial \eta} \leqslant \frac{1}{\eta} \cdot \frac{W(p^*)}{-W'(p^*)} = \frac{1}{\eta} \cdot \frac{W(p^*)}{1 - F(p^*)}$$

which holds by Lemma 1.1.

(ii) Every consumer with search costs in the interval $[s^*_{pas}(\eta, \mu), s^*_{act}(\eta, \mu)]$ is a passive search consumer. Since the equilibrium price p^* is increasing in η by Proposition 1.3, $s^*_{pas}(\eta, \mu) = \mu W(p^*(\eta, \mu))$ is decreasing in η and $s^*_{act}(\eta, \mu) = \eta W(p^*(\eta, \mu))$ is increasing in η by (i). Henceforth, the mass of passive search consumers expands at the expense of the mass of active search consumers and inactive consumers. □

Corollary 1.1 and 1.2 are illustrated in Figure 1.3 which plots the consumer surplus $U(s; (\eta, \mu))$ for fixed μ and two different recommendation precision. When recommendations are uninformative ($\eta = \mu$, green) we can see that low search costs consumers benefit compared to informative recommendations ($\eta > \mu$, blue). For uninformative recommendations, no consumers adopt the passive search strategy. When recommendations are informative, some consumers who actively searched with uninformative recommendations as well as some consumers who used to be inactive adopt the passive search strategy.

1.5.2 Equilibrium Demand and Firm Profits

Are more products sold with informative or uninformative recommendations? In light of Corollary 1.2, the answer might seem straightforward. More precise recommendations induce some inactive consumers to passively search and thereby increase market participation which should lead to more sold products. However, some active search consumers

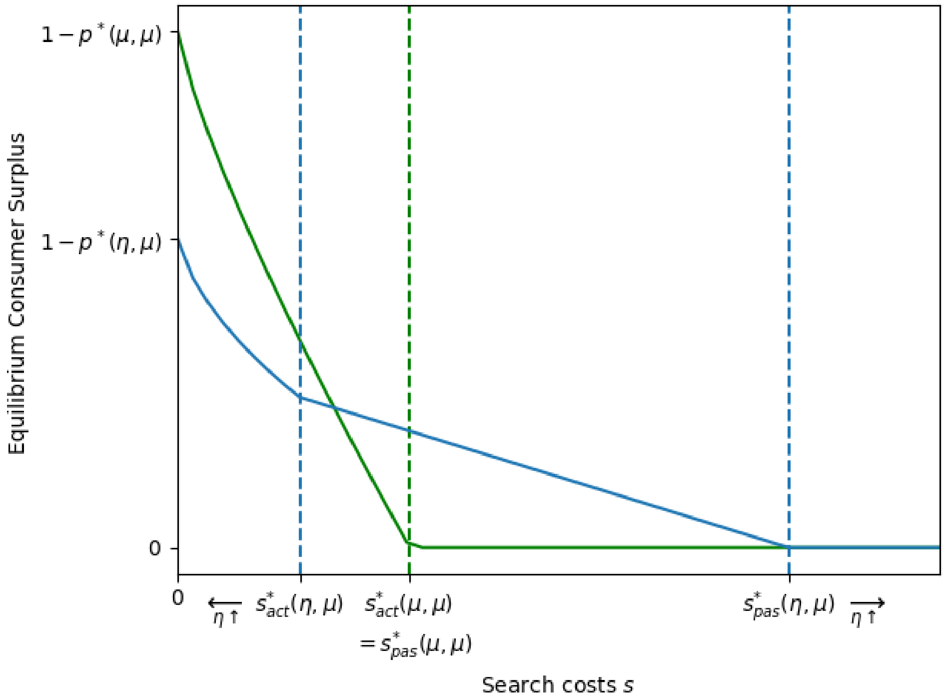

Figure 1.3: Consumer Surplus as function of search costs for uninformative recommendations ($\eta = \mu$, green) and informative recommendations ($\eta > \mu$, blue).

start to adopt the passive search strategy since they anticipate larger prices in equilibrium when recommendations are more precise. Active search consumers search until they find a satisfactory product, while passive search consumers buy a product if and only if the match valuation of the recommended firm is above the price. The extensive margin (how many consumers participate in the search market) is given by

$$\lambda_{ext}(\eta) = \frac{1}{\bar{s}}\eta W(p^*(\eta,\mu))$$

while the intensive margin (how many products a participating consumer buys on average) is given by

$$\lambda_{int}(\eta) = (1 - \frac{\mu}{\eta}) \cdot \eta \cdot (1 - F(p^*(\eta,\mu)) + \frac{\mu}{\eta}.$$

Although the extensive margin is increasing, the intensive margin might be decreasing since the relative mass $\frac{\mu}{\eta}$ of active search consumers to participating consumers decreases in η. However, any passive search consumer is more likely to buy from the recommended product.[9] Corollary 1.3 shows that recommendations decrease the intensive margin at

[9] That is, $\eta \cdot (1 - F(p^*(\eta,\mu))$ is increasing in η. The partial derivative w.r.t. η is positive if and only if

$$\frac{\partial p^*}{\partial \eta} \leqslant \frac{1}{\eta}\frac{1-F(p^*)}{f(p^*)}.$$

By Lemma 1.1, we know that $\frac{\partial p^*}{\partial \eta} \leqslant \frac{1}{\eta}\frac{W(p^*)}{1-F(p^*)} \leqslant \frac{1}{\eta}\frac{1-F(p^*)}{f(p^*)}$ where the final inequality stems from Lemma 1.2.

least for sufficiently imprecise recommendations. This is not surprising since for imprecise recommendations, there are almost no passive search consumers in the market.

Corollary 1.3. *Suppose that Assumption 1.2 is satisfied. The extensive margin $\lambda_{ext}(\eta)$ is increasing in the recommendation precision η and there exists η' such that the intensive margin is decreasing in $\eta \in [\mu, \eta']$.*

Proof. The extensive margin is increasing if and only if $\eta W(p^*(\eta, \mu))$ is increasing in η which is implied by Lemma 1.1. The intensive margin is decreasing in η if and only if

$$\lambda'_{int}(\eta) = -\frac{\mu}{\eta^2} + (1 - F(p^*(\eta, \mu)) - (\eta - \mu)\frac{\partial p^*}{\partial \eta} f(p^*(\eta, \mu)) < 0.$$

Note that

$$\lim_{\eta \to \mu} \lambda'_{int}(\eta) = -\frac{1}{\mu} + (1 - F(p^a)) < 0.$$

Therefore, there exists η' such that $\lambda'_{int}(\eta) < 0$ for all $\eta \leqslant \eta'$. □

Although numerical examples show that the equilibrium demand $D^*(\eta) := \lambda_{ext}(\eta) \cdot \lambda_{int}(\eta)$ is typically increasing in the recommendation precision, there exists no straightforward statistical properties that ensure that this is indeed valid in general. Since informative product recommendation might lead to a reduced equilibrium demand it might also appear plausible that the equilibrium profits of firms are negatively impacted by product recommendations. However, Proposition 1.5 shows that this is not a concern as higher market prices outweigh any potential demand reduction.

Proposition 1.5. *Suppose that Assumption 1.2 is satisfied. The equilibrium profits of firms are increasing in the recommendation precision η.*

1.6 Extensions and Variations

1.6.1 Micro-foundations of Recommendations

Up until now, I have taken a reduced-form approach to recommendations: The intermediary has been a myopic player who recommended firms in such a way that searching for the recommended firm resulted in a match with probability $\eta \geqslant \mu$ for every consumer and every firm has ex-ante the same chance to be recommended.

For exposure, I assume that there exists a single consumer whose search costs are drawn uniformly from $[0, \bar{s}]$ and each firm is a match with probability μ and the valuation of matches is drawn independently from F. As usual, there is an equivalent representation with a unit mass of consumers. I assume that the intermediary has information about whether a firm is a match to the consumer. We denote $\theta_i = 1$ if firm i is a match and $\theta_i = 0$ if it is not a match, so that the state space Θ is given by $\Theta = \{0, 1\}^{[0,1]}$.

For exposure and to avoid issues of non-measurability, I assume that first, nature draws $m \in \mathbb{N}$ firms uniformly and independently from the unit interval of firms, and the intermediary only receives information about the m drawn firms and recommends one of the m firms. Let I denote the realized set of m firms about which the intermediary receives information. Since $(\theta_i)_{i\in[0,1]}$ are assumed to satisfy independence and with slight abuse of notation, the prior probability μ of matches in the set I satisfies

$$\mu((\theta_i)_{i\in I}) = \mu^{(\sum_{i\in I}\theta_i)} \cdot (1-\mu)^{(m-\sum_{i\in I}\theta_i)}$$

by definition.[10]

The intermediary has access to a Blackwell-signal $\pi : \Theta^I \to \Delta(S)$ where S is a finite signal space. I assume that the intermediary maximizes industry-profits.[11]

Letting $\theta_I = (\theta_i)_{i\in I}$, it is well known that any Blackwell-signal induces a distribution σ over posterior beliefs which satisfy $\sum_{\xi\in\text{supp}\sigma}\sigma(\xi)\xi(\theta_I) = \mu(\theta_I)$. We now define a recommendation strategy κ as a function that maps posterior beliefs $\xi \in \text{supp}\ \sigma$ to a probability distribution over recommended firms. That is, $\kappa : \text{supp}\ \sigma \to \Delta(I)$ and $\kappa_i(\xi)$ denotes the probability that firm i is recommended if the intermediary has posterior belief ξ.

The timing of the game including the strategic intermediary is now as follows: First, the intermediary chooses a recommendation strategy κ. Next, firms set prices simultaneously. Then, the recommended firm is drawn according to the recommendation strategy and the realized signal of the intermediary. The consumer observes which firm is recommended and searches for firms sequentially.

I impose the following symmetry assumptions on the distribution σ of posterior beliefs and the recommendation strategy κ:

1. The distribution σ over posterior beliefs is symmetric: That is, for any permutation ρ on I and any $\xi \in \text{supp}\ \sigma$ there exists $\xi^\rho \in \text{supp}\ \sigma$ such that $\sigma(\xi) = \sigma(\xi^\rho)$ and $\xi(\theta_I) = \xi^\rho(\theta_{\rho(I)})$ for all $\theta_I \in \Theta$.
2. The intermediary is restricted to symmetric recommendation strategies: A recommendation κ is symmetric if for any permutation ρ on I, and any $\xi \in \text{supp}\ \sigma$ it holds that $\kappa_i(\xi) = \kappa_{\rho(i)}(\xi^\rho)$.

Symmetry of σ states that we can permute firms and the information structure does not change. This is for example satisfied if the intermediary receives an identical and independent signal $\pi_i : \{0,1\} \to \Delta(S)$ for any firm $i \in I$. The symmetry of the recommendation

[10] A continuum of probability measures $(P_i)_{i\in[0,1]}$ are said to be independent if any finite selection $(P_i)_{i\in I}$ satisfy independence.

[11] This might be the case if for example the intermediary is a platform which obtains a share of the profits of the sellers operating on the platform. Of course, maximizing objective functions like consumer surplus or sale probability might be reasonable as well. Since we have solved the sub-game following any intermediary's decision (if symmetry is satisfied) other objective functions of the intermediary can be easily accommodated.

strategy κ might be interpreted as fairness condition: If the intermediary decides to recommend firm i with probability α and firm j with probability β when she has posterior belief ξ, then symmetry implies that with some posterior belief ξ' which is equal to ξ except that firms i and j are interchanged, firm i is recommended with probability β and firm j is recommended with probability α.

Letting $\xi_i = \sum_{\theta_I \in \Theta^I, \theta_i = 1} \xi(\theta_I)$ denote the probability that some firm $i \in I$ is a match when the posterior belief is ξ, we define the recommendation precision $\eta \equiv \eta(\kappa, \sigma)$ as the probability that some firm is a match conditional on that the firm is recommended. That is,

$$\eta(\kappa, \sigma) = \frac{\sum_{\xi \in \text{supp } \sigma} \sigma(\xi)\kappa_i(\xi)\xi_i}{\sum_{\xi \in \text{supp } \sigma} \sigma(\xi)\kappa_i(\xi)}$$

for some firm $i \in I$. The next Proposition shows that the crucial assumptions about the recommendation process are satisfied. In particular, ex-ante before the intermediary receives information about the firms, every subset of firms of equal size has the same chance to be recommended,[12] and the recommendation precision does not depend on the identity of the recommended firm.[13] These results ensure that firms are ex-ante symmetric such that the analysis of Section 1.3 and 1.4 is valid.

Proposition 1.6. *Suppose that the symmetry conditions* 1. *and* 2. *are satisfied. Then, (i) firm i is recommended with probability* $1/m$ *conditional on $i \in I$.(ii) Ex-ante (before I is realized) the probability that some firm in the interval* $[a, b] \subset [0, 1]$ *is recommended is equal to $b - a$. (iii) The recommendation precision $\eta \equiv \eta(\kappa, \sigma)$ is independent of $i \in I$.*

Hence, under the symmetry conditions 1. and 2. the equilibrium of the consumer search and pricing subgame only depends on the induced recommendation precision ρ and not directly on the specific distribution of posterior beliefs σ and recommendation strategy κ. If the intermediary is industry-profit maximizing, Proposition 1.5 immediately implies that the intermediary maximizes the recommendation precision ρ, as more precise recommendations induce higher industry profits. This raises the question which recommendation strategy maximizes the recommendation precision.

The answer is straightforward: Maximization of the recommendation precision requires that the intermediary recommends a firm which she believes has the largest probability to be a match for the consumer. The following Proposition formalizes this intuition.

Proposition 1.7. *κ^* maximizes $\eta(\kappa, \sigma)$ over all symmetric recommendation strategies if and only if κ^* satisfies*

$$\sum_{i \in \arg\max \xi_i} \kappa_i^*(\xi) = 1$$

[12] In the equivalent interpretation with a continuum of consumers, this implies that every firm is recommended to exactly one consumer.

[13] Also note that the match probability of firms outside of I is unaffected, as I is a measure zero set. Since consumers do not observe I but only the recommended firm, this implies that the match probability of any firm that is not recommended is unaffected and equal to μ.

for all $\xi \in$ *supp* σ. *Furthermore, the equilibrium recommendation precision* $\eta(\kappa^*, \sigma)$ *is given by*

$$\eta(\kappa^*, \sigma) = \sum_{\xi \in supp\ \sigma} \sigma(\xi) \max_{i \in I} \xi_i.$$

In light of Corollary 1.1 which states that recommendations are harmful to low search costs consumers, a policy-maker might be interested in limiting the recommendation precision. In practice, a policy-maker can influence the intermediary's recommendation problem by restricting the intermediary's access or usage to information. Both options can be modeled as a shift in the distribution of posterior beliefs σ. We say that a distribution of posterior beliefs σ is more informative than σ' if σ is a mean-preserving spread of σ'. Since Proposition 1.7 establishes that it is optimal for the intermediary to recommend the "best" firm, it seems natural that a more informative distribution of posterior beliefs induces a larger equilibrium recommendation precision. Corollary 1.4 verifies this intuition.

Corollary 1.4. *Suppose that* σ *and* σ' *satisfy symmetry. If* σ *is more informative than* σ'*, then* $\eta(\kappa^*, \sigma) \geqslant \eta(\kappa^*, \sigma')$.

Proof. It is well known that σ is a mean-preserving spread of σ' if and only if

$$\sum_{x \in \text{supp}\ \sigma} \sigma(x)c(x) \geqslant \sum_{x \in \text{supp}\sigma'} \sigma'(x)c(x)$$

for every convex $c(\cdot)$. Since $\xi \mapsto \max_{i \in I} \xi_i$ is convex, the result immediately follows from Proposition 1.7. □

1.6.2 Heterogeneous match probability

Consumers might not only differ in search costs but also in their prior match probability. For instance, some consumers might have a niche taste looking for a particular kind of product, while others might find almost any product suitable. In this Section, I demonstrate how heterogeneous prior match probabilities can be incorporated in the model.

Suppose that consumers differ in search costs and match probability, and a mass $G(s, \mu)$ has search costs smaller than s and match probability smaller than μ. Suppose that the intermediary recommends a firm to every consumer and this firm is a match with probability $\eta(\mu) \geqslant \mu$.[14] Letting the reservation value $\tilde{r}(s/\mu)$ solve $s/\mu = W(\tilde{r}(s/\mu))$ and slightly adjusting the approach from Section 1.3 yields an analogue to Proposition 1.1.

[14] $\eta(\mu) \geqslant \mu$ is a minimal requirement on the recommendation precision since an intermediary can achieve $\eta(\mu) = \mu$ by recommending a random firm. In order to determine the pricing equilibrium and for comparative statics, it is also sensible to assume that $\eta(\mu)$ is differentiable and increasing.

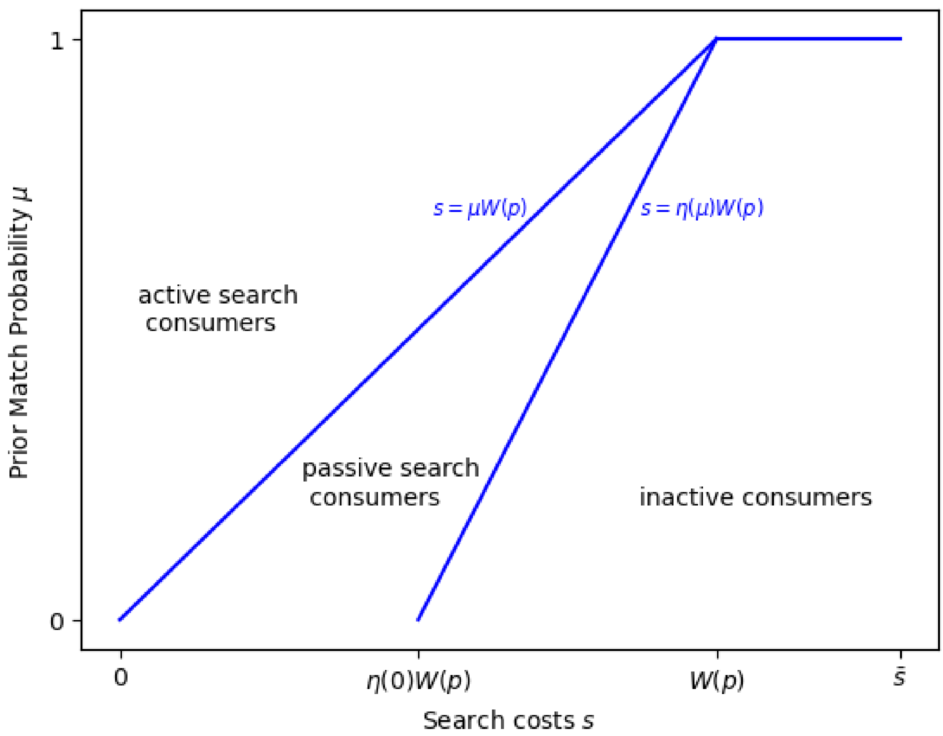

Figure 1.4: Separation of consumers into active search, passive search and inactive consumers for heterogeneous match probability and search costs.

Proposition 1.8. *Suppose that consumers expect prices $p_i = p$. The optimal consumer search rule satisfies:*

(i) Consumers of type (s, μ) for which $s \leqslant \mu W(p)$ search for the recommended product first. They buy if $v_i - p_i \geqslant \tilde{r}(\frac{s}{\mu}) - p$. Otherwise, they search for other firms until finding some firm for which $v_i - p_i \geqslant \tilde{r}(\frac{s}{\mu}) - p$.

(ii) Consumers of type (s, μ) for which $\mu W(p) < s \leqslant \eta(\mu)W(p)$ search for the recommended product and buy if $v_i \geqslant p_i$. Otherwise, they exercise the outside option.

(iii) Consumers of type (s, μ) for which $s \geqslant \eta(\mu)W(p)$ exercise the outside option immediately without searching.

Figure 1.4 illustrates the separation of consumers into active search consumers, passive search consumers and inactive consumers. Different to Section 1.3, the choice of which search rule to adapt does not depend only on the search costs but also on the prior match probability. If the search costs s are smaller than the incremental benefit $\eta(\mu)W(p)$ of searching for the recommended firm, consumers search for the recommended firm. If additionally, the search costs are smaller than the incremental benefit $\mu W(p)$ of searching for a firm that is not recommended, consumers will search beyond the recommended firm if the recommended product is not satisfactory.

Thus, the pricing problem of firms is also determined by the relative mass of passive search consumers who act as if they were in a monopoly to active search consumers who potentially visit many firms. It is reasonable to suspect that the comparative statics and welfare implications of product recommendations are similar to those derived in this article. However, inferring comparative statics and verifying that the first-order condition of profit-maximization constitutes a unique equilibrium is not straightforward and not within the scope of this article.

1.6.3 Full Market Coverage

This article has focused on the case in which the market is partially covered. That is, regardless of recommendation precision there are consumers with sufficiently small search costs searching actively for firms as well as consumers with sufficiently large search costs who are inactive and never buy a product. In this case, the implications of recommendations are mainly driven by the effect that precise recommendations induce inactive consumers to search for the recommended product and, as a consequence of rising market price, also active search consumers to adopt the passive search strategy.

Now, we study what happens when the market is fully covered even without product recommendations. Suppose that the largest search costs $\bar{s}$ satisfy $r(\bar{s}) \geqslant p^m$ or equivalently $\bar{s} \leqslant \mu W(p^m)$. As prices above the monopoly price p^m cannot be optimal, this implies that for every consumer the reservation value of searching exceeds the anticipated price in any equilibrium. Hence, every consumer adopts the active search strategy regardless of recommendation precision η. Following the derivation from Section 1.4, the demand function $\tilde{D}(p_i, p)$ is now entirely driven by active search consumers and can be described by

$$\tilde{D}(p_i, p) = \frac{1}{\bar{s}} \int_0^{\bar{s}} \frac{1 - F(r(s) + p_i - p)}{1 - F(r(s))} ds = \frac{\mu}{\bar{s}} \int_{r(\bar{s})}^{\bar{v}} (1 - F(\tilde{r} + p_i - p)) d\tilde{r}.$$

As a result, the equilibrium price $\tilde{p}$ when the market is fully covered is characterized by

$$\tilde{p} = \frac{D(\tilde{p}, \tilde{p})}{\frac{\tilde{D}}{\partial p_i}(\tilde{p}, \tilde{p})} = \frac{\mu}{1 - F(r(\bar{s}))}.$$

When the market is fully covered, the equilibrium price is independent of the recommendation precision η. There are two reasons why this is the case: Firstly, in contrast to Section 1.3, precise product recommendations do not induce any high search costs consumers to participate in the search market since any consumer already adopts an active search strategy. High search costs consumers induce a negative externality since they search for fewer firms and drive up the equilibrium price. Second, product recommendations do not change the search and purchase strategy of active search consumers (see Proposition 1.1) since the intermediary recommends only a single firm, and thus does not influence the continuation payoff of searching beyond the recommended firm.

Perhaps surprisingly, this immediately implies that the welfare implications of product recommendations are almost completely reversed. Without full market coverage, firm profits increase in the recommendation precision (Proposition 1.5), while low search costs consumers are harmed by product recommendations and high search costs consumers profit from product recommendations (Corollary 1.1 and 1.2). With full market coverage, firms are unaffected by product recommendations since the equilibrium profits $\tilde{p} \cdot D(\tilde{p}, \tilde{p})$ are independent of the precision of product recommendations, and the surplus of a consumer with search costs $s \in [0, \bar{s}]$ is given by

$$\tilde{U}(s) = r(s) - \tilde{p} + (\eta - \mu)W(r(s))$$

which is increasing in the recommendation precision η. Intuitively, all consumers profit from product recommendations since the first firm has a larger match probability, and thus consumer search for fewer firms on average and save on search costs, while market prices are not affected.

The analysis for fully covered search markets clarifies that the crucial economic mechanism of product recommendations in partially covered markets is that it incentivizes high search costs consumers, who were previously inactive, to participate in the market which causes the demand to become less price-elastic, and thereby induces market prices to rise.

1.7 Discussion

This article has examined the effects of personalized product recommendations on consumer search, market price, and welfare. I have extended the literature of competition and sequential consumer search with heterogeneous search costs by introducing an intermediary with information about consumers' preferences recommending firms.

In this article, I impose a very particular assumption on the structure of match valuations, and the information consumers receive about match valuations due to recommendations: Consumers only have positive and random match valuations if a firm is a match and the information a consumer receives is only informative about matches and not about the match valuations.

This assumption is different to how recommendations are modeled in most of the related literature: For instance, Rhodes and Zhou (2024) assume that consumers sharing their data receive a recommendation which perfectly reveals the product with the highest match valuations. In this case, consumers have no incentives to sample beyond the recommended firm which increases the equilibrium price. Therefore, the economic mechanism driving the results in most of the literature is that existing consumers search less and are less price-sensitive as a result of informative recommendations.

If the recommendations are only informative about the match probability of the recommended firm, the demand of active search consumers is unaffected by recommendations. Therefore, my modeling assumptions isolate and highlight another economic mechanism which is, to the best of my knowledge, not discussed in the literature: Product recommendations incentivize new, less price-sensitive, consumers to participate which induces firms to set higher prices.

In particular, one of the main findings is that product recommendations induce a separation of consumers along their search costs. Low search costs consumers first search for the recommended firm, but also search beyond the recommended firm if the match valuation for the product of the recommended product is not satisfactory. Consumers with intermediate search costs also search first for the recommended firm, but do not search beyond the recommended firm if the match valuation is unsatisfactory. Finally, high search costs consumers do not participate in the market. Product recommendations enable inactive consumers to search for the recommended product. This exerts a negative externality on all participating consumers as demand of high search costs consumers is more price-insensitive than of low search costs. As a consequence, product recommendations lead to higher market prices and harm low search costs consumers.

There are several avenues for future research. First, although recommendations harm low search costs consumers and benefit high search costs consumers, the effect of recommendations on total consumer surplus remains ambiguous. Cowan (2016) establishes conditions which guarantee an overall positive or negative effect on consumer surplus in the context of third-degree price discrimination. It would be interesting to find analogue conditions for personalized product recommendation.

Second, as noted by Moraga-González et al. (2017), in order to show existence and uniqueness of the equilibrium in a sequential search model with heterogeneous search costs, it is necessary to assume that the supply side consists of infinitely many sellers. However, Rhodes and Zhou (2024) show by example that an intriguing effect can occur with finitely many sellers: If the informativeness of the recommendation increases in the number of sellers, for instance because the intermediary recommends the best product and has independent information about each seller, then market prices might be increasing in the number of sellers because more consumers rely on the recommendation and this might even outweigh larger competition. In theory, this mechanism should also be present in a version of this article's model with finitely many sellers.

Lastly, this article has imposed the simplifying assumption that recommendations are only informative about consumer-firm matches and are not informative about match valuations or prices. In practice, recommendation algorithms might be much more nuanced. For instance, Donnelly, Kanodia, and Morozov (2024) find that personalized recommendations might be an appealing tool to shift demand from best-selling products to niche products which offer a higher profitability. I view it as a fruitful direction for future research to study unrestricted recommendation algorithms.

1.A Appendix

1.A.1 Auxiliary Lemmas

The following auxiliary Lemmas are repeatedly used.

Lemma 1.2. *$W(p) = \int_p^{\bar{v}}(1-F(z))dz$ is log-concave in p on $[0, \bar{v}]$. In particular, $\frac{W(p)}{1-F(p)} \leqslant \frac{1-F(p)}{f(p)}$ for all $p \in [0, \bar{v}]$.*

Proof. First, note that W is at least twice differentiable, and $W'(p) = -(1 - F(p))$ and $W''(p) = f(p)$. Thus, W is log-concave if and only $W'(p)^2 \geqslant W(p)W''(p)$ for all $p \in [0, \bar{v}]$ which is equivalent to $\frac{W(p)}{1-F(p)} \leqslant \frac{1-F(p)}{f(p)}$. This inequality is satisfied if and only if

$$\int_p^{\bar{v}} (\frac{1-F(z)}{1-F(p)} - \frac{f(z)}{f(p)})dz \leqslant 0.$$

This is satisfied if the integrand is positive for all $z \geqslant p$ which holds since log-concavity of $1 - F$ implies

$$\frac{1-F(z)}{f(z)} \leqslant \frac{1-F(p)}{f(p)}$$

for all $z \geqslant p$. □

Lemma 1.3. *Let $\Pi_z^m(p) = (1 - F(p + z))p$ and let p^m solve $p^m = \frac{1-F(p^m)}{f(p^m)}$. If $p > p^m$, then $\Pi_z^m(p^m) > \Pi_z^m(p)$ for any $z \geqslant 0$.*

Proof. Note that

$$\begin{aligned}
\frac{\partial \Pi_z^m(p)}{\partial p} &= (1 - F(p + z)) - f(p + z)p \\
&= f(p + z)(\frac{1 - F(p + z)}{f(p + z)} - p) \\
&\leqslant f(p + z)(\frac{1 - F(p)}{f(p)} - p) \\
&< 0
\end{aligned}$$

where the first inequality follows from log-concavity of $1 - F$ (the reverse hazard rate is decreasing), and the second inequality follows from $p > p^m$. Thus, $(1 - F(p + z))p < (1 - F(p^m + z))p^m$ for any $z \geqslant 0$. □

Lemma 1.4. *$D(p_i, p)$ is differentiable in p_i on $(0, \bar{v})$ and twice differentiable everywhere on $(0, \bar{v})$ except at $p_i = p$.*

Proof. First, note that the problem with naively differentiating $D(p_i, p)$ is that

$$\frac{\partial D(p_i, p)}{\partial p_i} = -(\eta - \mu)\eta \frac{\partial F(p_i))}{\partial p_i} - \mu \int_p^{\bar{v}} \frac{\partial F(\tilde{r} + p_i - p))}{\partial p_i} d\tilde{r}$$

where $\frac{\partial F(\tilde{r}+p_i-p))}{\partial p_i}$ is not differentiable at $\tilde{r} = \bar{v} - p_i + p$. Therefore, we first rewrite the demand function as follows

$$\begin{aligned}
D(p_i, p) &= (\eta - \mu)\eta W(p)(1 - F(p_i)) + \mu \int_p^{\bar{v}} (1 - F(\tilde{r} + p_i - p))d\tilde{r} \\
&= (\eta - \mu)\eta W(p)(1 - F(p_i)) + \mu \int_{p_i}^{\bar{v}+p_i-p} (1 - F(z))dz \\
&= \begin{cases} (\eta - \mu)\eta W(p)(1 - F(p_i)) + \mu \int_{p_i}^{\bar{v}+p_i-p} (1 - F(z))dz & \text{if } p_i < p \\ (\eta - \mu)\eta W(p)(1 - F(p_i)) + \mu \int_{p_i}^{\bar{v}} (1 - F(z))dz & \text{if } p_i > p \end{cases}.
\end{aligned}$$

Now,

$$\frac{\partial D(p_i, p)}{\partial p_i} = \begin{cases} -(\eta - \mu)\eta W(p) f(p_i) - \mu(F(\bar{v} + p_i - p)) - F(p_i)) & \text{if } p_i < p \\ -(\eta - \mu)\eta W(p) f(p_i) - \mu(1 - F(p_i)) & \text{if } p_i > p \end{cases}.$$

Thus, $D(p_i, p)$ is differentiable everywhere since left and right derivatives agree for $p_i = p$.

Further, it can easily be seen that $D(p_i, p)$ is twice differentiable everywhere except at $p_i = p$ (unless $\lim_{v \to \bar{v}} f(v) = 0$ in which case it is also twice differentiable at $p_i = p$). □

1.A.2 Proof of Proposition 1.2

Outline First, we show that the optimal price cannot be strictly larger than the monopoly price p^m (Lemma 1.5). This implies that we can exclude any prices above p^m as the equilibrium price. Next, Lemma 1.6 establishes that the equilibrium candidate price p^* is well-defined, i.e. the first-order condition of profit-maximization paired with the symmetry condition $p_i = p$ admits a unique solution. Finally, Lemma 1.7 shows that a firm's profit function is concave which implies that the solution to the first-order condition solves the profit-maximization problem.

Lemma 1.5. *$p_i D(p_i, p)$ is maximized for some $p_i \leqslant p^m$ where p^m solves $p^m = \frac{1-F(p^m)}{f(p^m)}$ for any $p \in [0, \bar{v}]$.*

Proof. Suppose $p_i > p^m$. We will show that a firm can increase profits by charging the monopoly price, i.e. $p^m D(p^m, p) > p_i D(p_i, p)$ for all $p_i > p^m$. We have

$$\begin{aligned}
p^m D(p^m, p) &= (\mu - \eta)\eta W(p)(1 - F(p^m))p^m + \mu \int_p^{\bar{v}} (1 - F(\tilde{r} + p^m - p))p^m d\tilde{r} \\
&> (\mu - \eta)\eta W(p)(1 - F(p_i))p_i + \mu \int_p^{\bar{v}} (1 - F(\tilde{r} + p_i - p))p_i d\tilde{r} \\
&= p_i D(p_i, p).
\end{aligned}$$

where the inequality follows immediately from Lemma 1.3. □

Lemma 1.6. *There exists a unique solution p^* to $p = \frac{D(p,p)}{-\frac{\partial D}{\partial p_i}(p,p)}$ and $p^* \in [p^a, p^m]$ where $p^a = \frac{W(p^a)}{1-F(p^a)}$ is the price without recommendations and $p^m = \frac{1-F(p^m)}{f(p^m)}$ is the monopoly price.*

Proof. Let
$$\varepsilon(p,\eta) = \frac{D(p,p)}{-\frac{\partial D}{\partial p_i}(p,p)} = \frac{\eta(\eta-\mu)(1-F(p))+\mu}{\eta(\eta-\mu)f(p)+\mu\frac{1-F(p)}{W(p)}}.$$

We are going to show that (i) $\varepsilon(p^a,\eta) \geqslant p^a$, (ii) $\varepsilon(p^m,\eta) \leqslant p^m$ and (iii) $\partial\varepsilon(p,\eta)/\partial p < 0$ for any $\eta \in [\mu, 1]$ and $p \leqslant p^m$. By the intermediate value theorem, (i) and (ii) imply that for every $\eta \in [\mu, 1]$ there exists a solution p^*, and (iii) implies that this solution is unique.

Let $A(p,\eta) = \eta(\eta-\mu)(1-F(p))+\mu$ denote the numerator of ε and $B(p,\eta) = \eta(\eta-\mu)f(p)+\mu\frac{1-F(p)}{W(p)}$ denote the denominator. Then, (i) is equivalent to

$$\begin{aligned}
0 &\leqslant A(p^a,\eta) - B(p^a,\eta)p^a \\
&= \eta(\eta-\mu)(1-F(p^a)) + \mu - p^a \cdot (\eta(\eta-\mu)f(p^a) + \mu\frac{1-F(p^a)}{W(p^a)}) \\
&= \eta(\eta-\mu)(1-F(p^a) - p^a f(p^a)) + \mu\left(1 - p^a \cdot \frac{1-F(p^a)}{W(p^a)}\right) \\
&= \eta(\eta-\mu)f(p^A)\left(\frac{1-F(p^a)}{f(p^a)} - \frac{W(p^a)}{1-F(p^a)}\right) + 0
\end{aligned}$$

which is satisfied since $\frac{1-F(p)}{f(p)} \geqslant \frac{W(p)}{1-F(p)}$ for all p by Lemma 1.2.

Next, (ii) is equivalent to

$$\begin{aligned}
0 &\geqslant A(p^m,\eta) - p^m B(p^m,\eta) \\
&= \eta(\eta-\mu)(1-F(p^m)) + \mu - p^m \cdot \left(\eta(\eta-\mu)f(p^m) + \mu\frac{1-F(p^m)}{W(p^m)}\right) \\
&= \eta(\eta-\mu)(1-F(p^m) - p^m f(p^m)) + \mu\left(1 - p^m \cdot \frac{1-F(p^m)}{W(p^m)}\right) \\
&= 0 + \mu\left(1 - \frac{1-F(p^m)}{f(p^m)} \cdot \frac{1-F(p^m)}{W(p^m)}\right)
\end{aligned}$$

which is satisfied since $(1-F(p))^2 \geqslant f(p)W(p)$ for all p by Lemma 1.2.

Lastly, (iii) is satisfied if and only if

$$\begin{aligned}
0 &< \frac{\partial A(p,\eta)}{\partial p} B(p,\eta) - A(p,\eta)\frac{\partial B(p,\eta)}{\partial p} \\
&= (-\eta(\eta-\mu)f(p))\left(\eta(\eta-\mu)f(p) + \mu\frac{1-F(p)}{W(p)}\right) \\
&\quad - (\eta(\eta-\mu)(1-F(p)) + \mu)\left(\eta(\eta-\mu)f'(p) + \mu\frac{d}{dp}(\frac{1-F(p)}{W(p)})\right) \\
&= \eta^2(\eta-\mu)^2\left(-f(p)^2 - (1-F(p))f'(p)\right) \\
&\quad + \mu^2\left(-\frac{d}{dp}(\frac{1-F(p)}{W(p)})\right) \\
&\quad + \eta(\eta-\mu)\mu\left(-f(p)\frac{1-F(p)}{W(p)} - (1-F(p))\frac{d}{dp}(\frac{1-F(p)}{W(p)}) - f'(p)\right).
\end{aligned}$$

which holds since first

$$f(p)^2 + (1-F(p))f'(p) > 0$$

by log-concavity of $1-F$, and second

$$\frac{d}{dp}(\frac{1-F(p)}{W(p)}) = \frac{-f(p)W(p) + (1-F(p))^2}{W(p)^2} > 0$$

by Lemma 1.2, and lastly

$$\begin{aligned}
f(p)\frac{1-F(p)}{W(p)} + (1-F(p))\frac{d}{dp}(\frac{1-F(p)}{W(p)}) + f'(p) &> \frac{(1-F(p))f(p)}{W(p)} + f'(p) \\
&> \frac{f(p)^2}{1-F(p)} + f'(p) \\
&> 0.
\end{aligned}$$

where the first inequality follows from $\frac{d}{dp}(\frac{1-F(p)}{W(p)}) > 0$, the second inequality follows from Lemma 1.2, and the third inequality follows form log-concavity of $1-F(p)$. □

Lemma 1.7. *The profit function $p_i D(p_i, p)$ is concave in p_i for all p.*

Since $p_i D(p_i, p)$ is twice differentiable almost everywhere, as established by Lemma 1.4, it is concave in p_i if and only if

$$0 > p_i \frac{\partial^2 D(p_i,p)}{\partial p_i^2} + 2\frac{\partial D(p_i,p)}{\partial p_i}.$$

If $p_i > p$, then this inequality amounts to

$$0 > -(\eta-\mu)\eta(p_i f'(p_i) + 2f(p_i)) - \mu(2(1-F(p_i)) + p_i f(p_i))$$

which is satisfied since $p_i \frac{f'(p_i)}{f(p_i)} > -2$ by Assumption 1.1 and $p_i < p^m$ implies $\frac{1-F(p_i)}{f(p_i)} >$

$\frac{1-F(p^m)}{f(p^m)} = p^m > p_i$.

And if $p_i > p$, then the inequality amounts to

$$\begin{aligned} 0 &> -(\eta-\mu)\eta(p_i f'(p_i) + 2f(p_i)) - \mu(2(F(\bar{v}+p_i-p) - F(p_i)) + p_i(f(\bar{v}+p_i-p) - f(p_i)) \\ &= -(\eta-\mu)\eta(p_i f'(p_i) + 2f(p_i)) + \mu \cdot \int_{p_i}^{\bar{v}+p_i-p} (-2f(z) - p_i f'(z))dz \end{aligned}$$

which is satisfied since $p_i f'(p_i) + 2f(p_i) > 0$ by Assumption 1.1 and this also implies $-2f(z) - p_i f'(z) < 0$ for all $z > p_i$ since either $f'(z) \geqslant 0$ in which case the inequality is trivially satisfied or $f'(z) < 0$ in which case $-2f(z) - p_i f'(z) < -2f(z) - zf'(z) < 0$.

1.A.3 Proof of Proposition 1.3 and 1.4

Proof. Recall that p^* uniquely satisfies

$$p^* = \varepsilon(p^*, \mu, \eta) := \frac{\eta(\eta-\mu)(1-F(p)) + \mu}{\eta(\eta-\mu)f(p) + \mu\frac{1-F(p)}{W(p)}}.$$

The implicit function theorem implies that p^* is differentiable around (μ, η) and

$$\frac{\partial p^*}{\partial \eta} = \frac{\frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial \eta}}{1 - \frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial p^*}} \quad \text{and} \quad \frac{\partial p^*}{\partial \mu} = \frac{\frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial \mu}}{1 - \frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial p^*}}.$$

The proof of Lemma 1.6 shows that $\frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial p^*} < 0$. Therefore, Proposition 1.3 is satisfied if and only if $\frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial \eta} > 0$ and Proposition 1.4 is satisfied if and only if $\frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial \mu} < 0$. It is straightforward to derive that $\frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial \eta} > 0$ if and only if

$$\frac{1-F(p^*)}{f(p^*)} > \varepsilon(p^*, \mu, \eta) = p^*$$

which is satisfied since $p^* < p^m = \frac{1-F(p^m)}{f(p^m)}$ and hence $\frac{1-F(p^*)}{f(p^*)} > p^*$. Analogously, it is routine to derive that $\frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial \mu} > 0$ if and only if

$$\eta f(p^*)\left(p^* - \frac{1-F(p^*)}{f(p^*)}\right) + \frac{1-F(p^*)}{W(p^*)}(\frac{W(p^*)}{1-F(p^*)} - p^*) < 0$$

which is satisfied since $p^* > \frac{1-F(p^*)}{f(p^*)}$ by the same arguments as above and $p^* > p^a = \frac{W(p^a)}{1-F(p^a)}$ and Lemma 1.2 imply that $p^* < \frac{W(p^*)}{1-F(p^*)}$. □

1.A.4 Proof of Lemma 1.1

Proof. Again, recall that p^* uniquely solves

$$p^* = \varepsilon(p^*, \mu, \eta) := \frac{\eta(\eta-\mu)(1-F(p)) + \mu}{\eta(\eta-\mu)f(p) + \mu\frac{1-F(p)}{W(p)}}.$$

The implicit function theorem implies that p^* is differentiable around (μ, η) and

$$\frac{\partial p^*}{\partial \eta} = \frac{\frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial \eta}}{1 - \frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial p^*}}.$$

The proof of Lemma 1.6 shows that $\frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial p^*} < 0$ and therefore

$$\frac{\partial p^*}{\partial \eta} \leqslant \frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial \eta}.$$

Differentiating ε w.r.t. η and simplifying yields

$$\frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial \eta} = \frac{(2\eta-\mu)(1-F(p^*)) - (2\eta-\mu)f(p^*)p^*}{\eta(\eta-\mu)f(p^*) + \mu(\frac{1-F(p^*)}{W(p^*)})}.$$

Factoring out $\frac{1}{\eta}\frac{W(p^*)}{1-F(p^*)}$ yields the following upper bound on $\frac{\partial p^*}{\partial \eta}$:

$$\begin{aligned} \frac{\partial p^*}{\partial \eta} &\leqslant \frac{\partial \varepsilon(p^*,\mu,\eta)}{\partial \eta} \\ &= \frac{1}{\eta}\frac{W(p^*)}{1-F(p^*)}\left((2\eta-\mu)\frac{1-F(p^*)-f(p^*)p^*}{(\eta-\mu)f(p^*)\frac{W(p^*)}{1-F(p^*)} + \frac{\mu}{\eta}}\right). \end{aligned}$$

Hence, Lemma 1.1 holds if

$$(2\eta-\mu)\frac{1-F(p^*)-f(p^*)p^*}{(\eta-\mu)f(p^*)\frac{W(p^*)}{1-F(p^*)} + \frac{\mu}{\eta}} \leqslant 1$$

which is equivalent to

$$(2\eta-\mu)(1-F(p^*)-f(p^*)p^*) - (\eta-\mu)f(p^*)\frac{W(p^*)}{1-F(p^*)} < \frac{\mu}{\eta}.$$

Note first that $p^* \geqslant \frac{W(p^*)}{1-F(p^*)}$ since log-concavity of W implies that the $W(p)/(1-F(p))$ is decreasing in p, and $p^* \geqslant p^A$ where $p^A = \frac{W(p^A)}{1-F(p^A)}$. Hence,

$$\begin{aligned}
&(2\eta-\mu)(1-F(p^*)-f(p^*)p^*)-(\eta-\mu)f(p^*)\frac{W(p^*)}{1-F(p^*)} \\
&\qquad < (2\eta-\mu)(1-F(p^*)-f(p^*)\frac{W(p^*)}{1-F(p^*)})-(\eta-\mu)f(p^*)\frac{W(p^*)}{1-F(p^*)} \\
&\qquad = (2\eta-\mu)(1-F(p^*))-(3\eta-2\mu)\cdot f(p^*)\frac{W(p^*)}{1-F(p^*)} \\
&\qquad \leqslant (2\eta-\mu))(1-F(p^*)-f(p^*)\frac{W(p^*)}{1-F(p^*)}) \\
&\qquad \leqslant \frac{\mu}{\eta}
\end{aligned}$$

where the second inequality holds since $3\eta-2\mu \geqslant 2\eta-\mu$ and the final inequality holds by Assumption 1.2. □

1.A.5 Proof of Proposition 1.5

Proof. First note that the equilibrium demand is given by

$$\begin{aligned}
D^*(\eta) &= \lambda_{int}(\eta)\cdot\lambda_{ext}(\eta) \\
&= \frac{W(p^*(\eta,\mu))}{\bar{s}}\left((\eta-\mu)\eta\cdot(1-F(p^*(\eta,\mu))+\mu\right).
\end{aligned}$$

Dropping $\frac{1}{\bar{s}}$ (w.l.o.g.) and letting $p^* \equiv p^*(\eta,\mu)$, we have

$$\begin{aligned}
\frac{\partial D^*(\eta)}{\partial\eta} &= -\frac{\partial p^*}{\partial\eta}(1-F(p^*))\left((\eta-\mu)\eta\cdot(1-F(p^*))+\mu\right) \\
&\quad + W(p^*)\left((2\eta-\mu)(1-F(p^*)-(\eta-\mu)\eta\frac{\partial p^*}{\partial\eta}f(p^*)\right) \\
&= -\frac{\partial p^*}{\partial\eta}\left((\eta-\mu)\eta\cdot(1-F(p^*))^2+\mu(1-F(p^*))+(\eta-\mu)\eta\cdot W(p^*)f(p^*)\right) \\
&\quad + (2\eta-\mu)W(p^*)(1-F(p^*) \\
&= -\frac{\partial p^*}{\partial\eta}\left((\eta-\mu)\eta\cdot((1-F(p^*))^2+f(p^*)W(p^*))+\mu(1-F(p^*))\right) \\
&\quad + (2\eta-\mu)W(p^*)(1-F(p^*).
\end{aligned}$$

The equilibrium profits $\Pi^*(\eta) = D^*(\eta)p^*$ are increasing if and only if

$$\frac{\partial p^*}{\partial\eta}\frac{D^*(\eta)}{p^*} \geqslant -\frac{\partial D^*(\eta)}{\partial\eta}.$$

The equilibrium price p^* satisfies by definition (see 1.1)

$$p^* = \frac{\eta(\eta-\mu)W(p^*)(1-F(p^*)) + \mu W(p^*)}{\eta(\eta-\mu)W(p^*)f(p^*) + \mu(1-F(p^*))} = \frac{D^*(\eta)}{\eta(\eta-\mu)W(p^*)f(p^*) + \mu(1-F(p^*))}.$$

Thus, plugging in $\frac{\partial D^*(\eta)}{\partial \eta}$ and the equation above yields $\Pi^*(\eta)$ is increasing if and only if

$$\frac{\partial p^*}{\partial \eta}(\,\eta(\eta-\mu)W(p^*)f(p^*) + \mu(1-F(p^*))\,)$$
$$\geqslant$$
$$\frac{\partial p^*}{\partial \eta}(\,(\eta-\mu)\eta\cdot((1-F(p^*))^2 + f(p^*)W(p^*)) + \mu(1-F(p^*))\,) - (2\eta-\mu)W(p^*)(1-F(p^*).$$

Simplifying yields

$$\frac{2\eta-\mu}{\eta-\mu}\cdot\frac{1}{\eta}\frac{W(p^*)}{1-F(p^*)} \geqslant \frac{\partial p^*}{\partial \eta}$$

which is implied almost immediately by Lemma 1.1. □

1.A.6 Proof of Proposition 1.6

Proof. (i) I show that any firms $i, j \in I$ have the same ex-ante probability to be recommended if symmetry is satisfied. Let ρ denote a permutation on I such that $\rho(i) = j$. The probability that firm $i \in I$ is recommended is given by

$$\begin{aligned}\sum_{\xi\in\text{supp }\sigma} \sigma(\xi)\kappa_i(\xi) &= \sum_{\xi\in\text{supp }\sigma} \sigma(\xi^\rho)\kappa_{\rho(i)}(\xi^\rho)\\ &= \sum_{\xi\in\text{supp }\sigma} \sigma(\xi^\rho)\kappa_j(\xi^\rho)\\ &= \sum_{\xi\in\text{supp }\sigma} \sigma(\xi)\kappa_j(\xi)\end{aligned}$$

where the first line follows from symmetry of σ and κ, the second line applies $\rho(i) = j$ and the last line is implied by symmetry of σ since $\xi \in \text{supp }\sigma$ if and only if $\xi^\rho \in \text{supp }\sigma$. Thus, it immediately follows that every firm in I has a probability of $\mid I \mid^{-1} = m^{-1}$ to be recommended by the intermediary. (ii) follows immediately since I is drawn uniformly from $[0, 1]$, and the recommended firm is drawn uniformly from I.

(iii) I show that

$$\rho := \frac{\sum_{\xi\in\text{supp }\sigma}\sigma(\xi)\kappa_i(\xi)\xi_i}{\sum_{\xi\in\text{supp }\sigma}\sigma(\xi)\kappa_i(\xi)}$$

does not depend on $i \in I$. (i) already clarifies that the numerator is equal to $1/m$, and therefore does not depend on i. Therefore, it remains to show that the denominator does not depend on i. Let ρ denote the permutation on I for which $\rho(i) = j$, $\rho(j) = i$ and $\rho(l) = l$ for all $l \in I$, $l \neq i, j$. It is immediate that given some posterior belief ξ, the probability that firm i is a match ξ_i satisfies $\xi_i = \xi^\rho_{\rho(i)}$. Thus,

$$\begin{aligned}\sum_{\xi \in \text{supp } \sigma} \sigma(\xi)\kappa_i(\xi)\xi_i &= \sum_{\xi \in \text{supp } \sigma} \sigma(\xi^\rho)\kappa_{\rho(i)}(\xi^\rho)\xi^\rho_{\rho(i)} \\ &= \sum_{\xi \in \text{supp } \sigma} \sigma(\xi^\rho)\kappa_j(\xi^\rho)\xi^\rho_j \\ &= \sum_{\xi \in \text{supp } \sigma} \sigma(\xi)\kappa_j(\xi)\xi_j.\end{aligned}$$

where the first line follows from symmetry of σ and κ and $\xi_i = \xi^\rho_{\rho(i)}$, the second line applies $\rho(i) = j$ and the last line is implied by symmetry of σ since $\xi \in$ supp σ if and only if $\xi_\rho \in$ supp σ. Therefore, also the denominator does not depend on the specific firm $i \in I$ such that ρ is well-defined, i.e. does not depend on i. $\square$

1.A.7 Proof of Proposition 1.7

Proof. By Proposition 1.6, we have $\sum_{\xi \in \text{supp } \sigma} \sigma(\xi)\kappa_i(\xi) = 1/m$ for any symmetric recommendation strategy κ. Therefore,

$$\eta(\kappa, \sigma) = m \sum_{\xi \in \text{supp } \sigma} \sigma(\xi)\kappa_i(\xi)\xi_i$$

such that the optimal recommendation strategy maximizes $\sum_{\xi \in \text{supp } \sigma} \sigma(\xi)\kappa_i(\xi)\xi_i$. Since $\eta(\kappa, \sigma)$ is independent of i by Proposition 1.6, we get that

$$\begin{aligned}\eta(\kappa, \sigma) &= m \sum_{\xi \in \text{supp } \sigma} \sigma(\xi)\kappa_i(\xi)\xi_i \\ &= \sum_{i \in I} \sum_{\xi \in \text{supp } \sigma} \sigma(\xi)\kappa_i(\xi)\xi_i \\ &= \sum_{\xi \in \text{supp } \sigma} \sigma(\xi) \left(\sum_{i \in I} \kappa_i(\xi)\xi_i \right).\end{aligned}$$

Thus, κ maximizes $\eta(\kappa, \sigma)$ if and only if κ maximizes $\sum_{i \in I} \kappa_i(\xi)\xi_i$ for any $\xi \in$ supp σ. Since $\sum_{i \in I} \kappa_i(\xi) = 1$ this is obviously maximized if and only if κ satisfies

$$\sum_{i \in \arg\max_j \xi_j} \kappa_i(\xi) = 1.$$

Lastly, any κ^* satisfying the equation above induces a recommendation precision is given by

$$\rho(\kappa^*, \sigma) = \sum_{\xi \in \text{supp } \sigma} \sigma(\xi) \max_{i \in I} \xi_i.$$

$\square$

References

Inaki Aguirre, Simon Cowan, and John Vickers. Monopoly price discrimination and demand curvature. *American Economic Review*, 100(4):1601–1615, 2010. 6

Simon P Anderson and Regis Renault. Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics*, pages 719–735, 1999. 1.1, 1.1.1

Simon P Anderson and Régis Renault. Consumer information and firm pricing: negative externalities from improved information. *International economic review*, 41(3):721–742, 2000. 1.1.1

Mark Armstrong and John Vickers. Patterns of competitive interaction. *Econometrica*, 90(1):153–191, 2022. 5

Mark Armstrong and Jidong Zhou. Paying for prominence. *The Economic Journal*, 121 (556):F368–F395, 2011. 1.1.1

Mark Armstrong, John Vickers, and Jidong Zhou. Prominence and consumer search. *The RAND Journal of Economics*, 40(2):209–233, 2009. 1.1.1

CDEI. Online targeting: Final report and recommendations. *Discussion Paper, Centre for Data Ethics and Innovation*, 2020. 1.1

Simon Cowan. Welfare-increasing third-degree price discrimination. *The RAND Journal of Economics*, 47(2):326–340, 2016. 1.7

Alexandre De Streel and Florian Jacques. Personalised pricing and eu law. 2019. 4

Tommaso Di Noia, Nava Tintarev, Panagiota Fatourou, and Markus Schedl. Recommender systems under european ai regulations. *Communications of the ACM*, 65 (4):69–73, 2022. 4

Peter A Diamond. A model of price adjustment. *Journal of economic theory*, 3(2):156–168, 1971. 1.1.1

Robert Donnelly, Ayush Kanodia, and Ilya Morozov. Welfare effects of personalized rankings. *Marketing Science*, 43(1):92–113, 2024. 1.7

Daniel Garcia and Sandro Shelegia. Consumer search with observational learning. *The RAND Journal of Economics*, 49(1):224–253, 2018. 1.1.1

Sinem Hidir and Nikhil Vellodi. Privacy, personalization, and price discrimination. *Journal of the European Economic Association*, 19(2):1342–1363, 2021. 1.1.1

Shota Ichihashi. Online privacy and information disclosure by consumers. *American Economic Review*, 110(2):569–595, 2020. 1.1.1

Ipsos. Consumer market study on online market segmentationtrough personalised pricing/offers in the european union. *London Economics and Deloitte*, Study for the European Commission, 2018. 1.1

Norman J Ireland. The provision of information in a bertrand oligopoly. *The Journal of Industrial Economics*, pages 61–76, 1993. 1.1.1

R Preston McAfee. Endogenous availability, cartels, and merger in an equilibrium price dispersion. *Journal of Economic Theory*, 62(1):24–47, 1994. 1.1.1

Andrey Minaev. *Steering of Consumer Search by Information Intermediaries: The Value of Consumers' Personal Data*. PhD thesis, The University of North Carolina at Chapel Hill, 2021. 1.1.1

José Luis Moraga-González, Zsolt Sándor, and Matthijs R Wildenbeest. Prices and heterogeneous search costs. *The RAND Journal of Economics*, 48(1):125–146, 2017. 1.1, 1.1.1, 1.4, 1.7

András Prékopa. Logarithmic concave measures with applications to stochastic programming. 1971. 1.4

Marcel Preuss. Search, learning, and tracking. *The RAND Journal of Economics*, 54(1): 54–82, 2023. 6

Andrew Rhodes and Jidong Zhou. Personalization and privacy choice. *Available at SSRN 4795118*, 2024. 1.1.1, 1.7

Robert W Rosenthal. A model in which an increase in the number of sellers leads to a higher price. *Econometrica: Journal of the Econometric Society*, pages 1575–1579, 1980. 1.1.1

Sebastian Felix Schwemer. Recommender systems in the eu: from responsibility to regulation? In *FAccTRec Workshop*, volume 21, 2021. 4

Ran Spiegler. The market for quacks. *The Review of Economic Studies*, 73(4):1113–1131, 2006. 1.1.1

Nora Szech. Welfare in markets where consumers rely on word of mouth. *Report.[155]*, 2011. 1.1.1

Hal R Varian. A model of sales. *The American economic review*, 70(4):651–659, 1980. 1.1.1

Martin L. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3): 641–654, 1979. ISSN 00129682, 14680262. URL `http://www.jstor.org/stable/1910412`. 1.1, 1.1.1, 1.3, 7

Asher Wolinsky. True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics*, 101(3):493–511, 1986. 1.1, 1.1.1, 7

Wei Zhou, Mingfeng Lin, Mo Xiao, and Lu Fang. Higher precision is not always better: Search algorithm and consumer engagement. *Management Science, forthcoming*, 2024. 1

Chapter 2

Competition and Consumer Search with Costly Product Returns

2.1 Introduction

In both online marketplaces and, to a lesser extent, traditional brick-and-mortar stores, consumers often face uncertainty about the value or quality of a product when making a purchase decision. This uncertainty can stem from various factors, such as the inability to physically inspect or test the product before buying it, particularly in online settings. To alleviate these concerns and encourage purchases, sellers frequently offer return policies.

In the European Union, return policies are not merely a matter of seller discretion but are mandated by regulations governing distance sales contracts. These regulations ensure that consumers have the right to return products bought online or through other distance sales methods within a specified period, typically 14 days.

At first glance, return policies appear to benefit consumers, offering them the ability to return unsatisfactory products. However, these policies can present significant challenges for sellers and the broader public. For sellers, excessive product returns can be expensive. They must handle the logistics of returns, including processing, restocking, and managing damaged or unsellable items. These activities result in significant operational costs, which may lead sellers to raise prices, ultimately impacting consumers negatively. Moreover, the environmental impact of product returns is a growing concern. The carbon footprint associated with the transportation of returned goods can be substantial. It is estimated that returns in the US alone generate approximately 15 million metric tonnes of carbon emissions annually and contribute 5 billion pounds of landfill waste (Schiffer, 2019).

These concerns raise several questions. First, how do product returns influence consumer behavior in terms of searching for and purchasing products? Second, what are the determining factors for the number of products returned in a market? Specifically, is the rate of product returns higher in online markets, where discovering new items is relatively easy, or in offline markets, where finding products typically requires more effort? Third, how do product returns shape consumer demand and competition among firms?

To answer these questions, I incorporate product returns into a model of price competition and sequential consumer search. In particular, I employ a consumer search model building on Wolinsky (1986) and Anderson and Renault (1999) with a continuum of firms offering horizontally differentiated products. Firms set prices simultaneously at the start of the game and consumers can search for a firm by paying fixed search costs c_s and thereby discover the firm's price and the observable component θ_i of the match value $v_i = \theta_i + \varepsilon_i$ for the product of firm i. Upon purchasing the product of firm i, the consumer also learns the residual match value ε_i. If the consumer is dissatisfied with the match value v_i, she can return the product and obtain a full refund by incurring positive return costs c_r.

I leverage the stationarity of the consumer's decision problem to derive the consumers' optimal search rule. In equilibrium, a consumer purchases a product if and only if the observable valuation exceeds a threshold θ^* and returns the product if and only if the total match value is below a threshold v^*. The threshold values (θ^*, v^*) satisfy two indifference conditions: At match value v^*, a consumer is indifferent between (i) keeping the product and (ii) returning the product and continuing to search following the optimal search rule. At the threshold θ^*, the consumer is indifferent between (i) purchasing the object and continuing optimally in the return stage and (ii) continuing the search at another firm following the optimal search rule.

An important question that arises is how optimal consumer search behavior depends on the consumer's ability to return products, as measured by the costs c_r associated with a product return. It is well-established that product return policies make it more attractive for consumers to purchase a firm's product, as they have the option to return it if unsatisfied. However, a key insight of this article is that market-wide return rights also enhance the appeal of not purchasing a product and continuing the search for alternatives since the expected value of finding and buying a product is larger with the ability to return a product.

When product return costs are lower, consumers are generally less inclined to return products. This is because, first, the costs of returns are reduced, and second, the continuation payoff from ongoing search is higher. However, the effect on the purchasing threshold θ^* is not straightforward. Lower return costs increase both the expected benefit of purchasing a product and the expected benefit of continuing to search. I show that the former effect outweighs the latter, making consumers more likely to purchase a product when return costs are low. Consequently, easier product returns increase both the likelihood of purchasing and the probability of subsequently returning the product.

This result poses the question of how product returns impact the demand for a firm's product, which is defined as the probability that a consumer buys and does not return a product of a firm she encounters. Under common log-concavity assumptions, product returns reduce the demand for a firm's product. The option to return products encourages consumers to continue searching for alternatives, and thereby increases the expected number of firms a consumer searches for, and reduces the probability of purchasing from a firm conditional on searching for this firm.

Given the consumer's optimal search rule, it is straightforward to derive the demand functions and the equilibrium price p^*. Under weak assumptions, demand is log-concave in a firm's own price, and thus the first-order condition of profit maximization constitutes an equilibrium in which every firm charges an identical price.

When search costs increase, consumers are less likely to buy as well as retain a firm's product, intensifying competition among firms and driving down the equilibrium price. Conversely, as the costs of returns rise, consumers become less likely to purchase a product but, contrary to the effect of search costs, more likely to keep it once bought. The precise effect on equilibrium prices depends on specific parameters and distributional assumptions, generally leading to a non-monotonic relationship between return costs and equilibrium price.

The model also predicts how the number of product returns varies with search costs. Are returns more frequent in markets where products are easily accessible, such as online platforms, or in markets where products are harder to find, like traditional brick-and-mortar stores?[1] The answer is nuanced since there is a trade-off: When search costs are low, the continuation payoff from searching is high, increasing the incentive to return a product and search for a better match. At the same time, lower search costs make it more likely that a consumer finds and purchases a well-suited product, reducing the need for returns. I argue that the first effect dominates, leading to more product returns in markets with lower search costs.

This article might be interpreted as a cautionary tale about the hidden consequences of product returns. The ability to return unsatisfactory products in a market significantly alters consumers' search behavior and can lead to increasing market prices which can ultimately impact consumers and even firms negatively.

Although product returns are the primary focus, this paper's model is not limited to product returns. Consumers can, without any loss of generality, return products only after deciding which one they wish to keep permanently. Therefore, if a consumer already owns a product and buys another, they will inevitably need to return one of the two, regardless of which product they ultimately choose to keep. This implies that the product returns model is strategically equivalent to a model without returns, where the consumer incurs search costs c_s to observe a product's price and discover her observable valuation, and then has the option to further investigate the product at costs c_r to observe the

[1] Of course, I acknowledge that there are other, systemic, differences between an e-commerce platform and an offline market.

residual valuation before deciding to buy the product or continue searching for another one. Hence, the costs of returning a product can be reinterpreted as search costs for the residual valuation.

This strategic equivalence implies a broader contribution to the literature of consumer search. Typically, equilibrium prices increase with search costs in consumer search models involving price discovery. However, this article shows that this insight does not necessarily hold when consumers have the option to further learn about their valuation before deciding to buy a product.

The paper is structured as follows: In the remainder of this Section, I discuss the related literature. In Section 2.2, I introduce the basic model. Section 2.3 deals with the search problem of the consumers. Section 2.4 solves the firm's problem, derives the equilibrium price, and discusses the comparative statics and welfare implications. Section 2.5 includes some extensions and variations of the model. Finally, I conclude this article in Section 2.6. The appendix includes a detailed treatment of the consumer search problem, along with all proofs omitted throughout the article.

2.1.1 Related Literature

The literature on consumer search with product returns is quite sparse. To the best of my knowledge, Petrikaitė (2018) and Janssen and Williams (2024) are the only contributions that examine product returns building on the canonical consumer search framework established by Wolinsky (1986) and Anderson and Renault (1999).

Janssen and Williams (2024) explore the interplay between competition in pricing and refund policies when consumers search for firms and can return products. Their findings indicate that, in the absence of regulation, product return fees are inefficiently high, as firms fail to account for the welfare benefits derived from consumers who return the product and continue to search other firms. In contrast, my analysis assumes that sellers must provide a free return policy as required by the EU for online sales. Here, the costs associated with returning a product represent the inconvenience of the return process and are fixed.

The work of Petrikaitė (2018) is closely aligned to my contribution. Petrikaitė (2018) also assumes an additive structure of match values where the consumer observes the first part after searching for a firm and the second part only after purchasing the product, and the consumer can return the product at fixed return costs. Two insights from Petrikaitė (2018) are particularly relevant for my work. First, it is a weakly dominant strategy for consumers to only ever return a product after deciding which product to keep permanently. Second, product return costs can be reinterpreted as search costs. It is well known that in models with price-directed search, equilibrium prices can decrease with higher search costs (see Choi, Dai, and Kim (2018); Haan, Moraga-González, and Petrikaitė (2018);

Shen (2015)). In this sense, it is not surprising that return costs increase the equilibrium price in the price-directed duopoly model of Petrikaitė (2018).

Different to Petrikaitė (2018), I consider a model with undirected search (i.e., prices are not observable prior to search), with infinitely many firms, general distributions of match values, and positive search as well as return costs. Even though prices are not observable, and firms therefore do not have an incentive to lower prices in order to be sampled earlier, the consumers' ability to return products can still increase market prices in my model since consumers become less price sensitive when buying products.

Greminger (2022) and Gibbard (2022) both analyze sequential consumer search with two-stage information acquisition search without prices. Using insights from the multi-armed bandit problem (see Gittins, Glazebrook, and Weber (2011)), they show that the optimal policy can be fully characterized by reservation values. In the appendix, I show that the optimal consumer search policy in my work can also be framed as a reservation value-based policy. In this way, I extend the consumer search problem of Greminger (2022) and Gibbard (2022) to search including prices. As a caveat, I assume infinitely many alternatives (firms) which allows characterizing the optimal search policy without relying on the literature on multi-armed bandit problems.

There is extensive literature that deals with optimal monopolistic refund policies in various settings (e.g. Courty and Hao (2000); Matthews and Persico (2007); Shulman, Coughlan, and Savaskan (2009); Hinnosaar and Kawai (2020); Inderst and Tirosh (2015); Krähmer and Strausz (2015); Ren and Jerath (2022); von Wangenheim (2024)). This literature generally finds that there are fewer product returns when consumer information is improved as there is less uncertainty with improved information.

As an example, Ren and Jerath (2022) consider a monopolistic seller who sets the price and return policy when consumers can acquire information about their valuation before buying the product. They find that there are fewer product returns when it is inexpensive to acquire pre-purchase information. In contrast, my work suggests that low search costs imply a larger number of returned products, as they lead to a higher continuation payoff from returning a product and finding a new one. This mechanism is absent in models with a single seller.

Regarding comparative statics concerning product return costs, my work is also related to Zhou (2022) who analyzes comparative statics of the canonical Wolinsky (1986) model when the distribution of match valuations changes. With free product returns, the consumer can observe the residual match valuation at no costs. Without the option to return products, she cannot obtain any information about the residual match valuation before deciding which product to keep permanently. Generally, as the costs of product returns increase, the effective distribution of match valuations becomes more dispersed. Similar to Zhou (2022), I find that increased dispersion leads consumers to search for longer on average, but it does not necessarily result in a higher or lower equilibrium price.

2.2 The Model

The market consists of a demand and supply side. The supply side is made up of a continuum of firms. Each firm $i \in [0,1]$ supplies a single indivisible product at some price p_i. The prices $(p_i)_{i\in[0,1]}$ are chosen simultaneously at the start of the game. On the demand side, there is a unit-mass of risk-neutral consumers with single-unit demand. The valuation of consumer $j \in [0,1]$ for the product of firm i is denoted by v_{ij}. For the remaining article, we drop the index j, and denote the valuation for firm i's product of a representative consumer as v_i. We assume that the valuation v_i consists of two additive parts. An observable valuation θ_i and a residual valuation ε_i, where $v_i = \theta_i + \varepsilon_i$.

Initially, consumers are uninformed about their valuations $(v_i)_i$ and prices $(p_i)_i$. Consumers can learn the price p_i of the product of firm i and the observable valuation θ_i by searching for firm i. Every time a consumer searches for a firm, she incurs search costs $c_s \in \mathbb{R}^+$. After observing (p_i, θ_i), the consumer chooses between buying the product or searching for another firm.[2] If the consumer buys from firm i, she pays price p_i and observes her total valuation v_i for the firm's product. The consumer can return the product by incurring return costs $c_r \in \mathbb{R}^+$ and thereby obtain a refund of the price p_i.[3] If the consumer has searched for $m \in \mathbb{N}$ firms, has returned $k < m$ products, and keeps the product of firm i, her realized utility is $v_i - p_i - m \cdot c_s - k \cdot c_r$. For the majority of the paper, I impose the simplifying assumption that the consumer must eventually return any product she does not want to keep. In Section 2.5, I show that this assumption is without loss of generality under weak conditions.

The observable valuations $(\theta_i)_i$ are drawn independently and identically from some CDF G, and the residual valuations (ε_i) are drawn identically and independently from each other and from $(\theta_i)_i$ from some CDF H, where H and G are at least twice differentiable and admit finite moments. Furthermore, to guarantee the unique existence of interior solutions to the consumer search problem for any pair of (c_s, c_r) of search and return costs, I also assume that θ_i and ε_i have full support on $\mathbb{R}$, i.e. $g(x) > 0$ and $h(x) > 0$ for all $x \in \mathbb{R}$.

It is useful to derive the distribution of the total valuation v_i. We denote the CDF of v_i conditional on $\theta_i = \theta$ as $F(\cdot \mid \theta)$. Since $v_i < v$ for some $v \in \mathbb{R}$ if and only if $\varepsilon_i < v - \theta_i$, it is straightforward to see that $F(v \mid \theta) = H(v - \theta)$. Therefore, the joint density of (θ_i, v_i) is given by $f(\theta, v) = g(\theta)h(v - \theta)$. The joint reliability (or survival) function $\bar{F}$ of (θ_i, v_i) is defined by $\bar{F}(\theta, v) = \int_\theta^\infty \int_v^\infty f(\tilde{\theta}, \tilde{v}) d\tilde{\theta} d\tilde{v}$. I impose the following log-concavity assumption.

[2] For simplicity, we assume that the outside option of not buying any product is $-\infty$ (i.e., the consumer must buy some firm's product). With infinitely many firms this assumption is without loss of generality, provided that search costs are sufficiently small.

[3] The return costs c_r represent the effort associated with returning the product. In particular, c_r does *not* represent a return fee such that firm i does not obtain c_r if its product is returned.

Assumption 2.1. *The joint reliability function $\bar{F}$ of (θ_i, v_i) is log-concave.*

Assumption 2.1 immediately implies that also the marginal survival functions $\bar{F}_v(\theta) := \bar{F}(\theta, -\infty) = 1 - G(\theta)$ and $\bar{F}_\theta(v) := \bar{F}(-\infty, v)$ are log-concave. As a leading example, I will illustrate results when the observable valuation θ_i and the residual valuation ε_i independently normally distributed. In the appendix, I show that normality of $\theta_i$and ε_i implies Assumption 1.

Remark 2.1. Suppose that $\theta_i \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$, $\varepsilon_i \sim \mathcal{N}(\mu_\varepsilon, \sigma_\varepsilon^2)$ and $(\theta_i, \varepsilon_i)_i$ are independent. Then, Assumption 1 is satisfied.

The solution concept I impose is symmetric and pure Perfect Bayesian Equilibrium (PBE) with the following standard restrictions: First, search is undirected. That is, if a consumer is indifferent between searching for firms i and j then the probability that she searches for firm i is equal to the probability that she searches for firm j. Second, if a consumer encounters an off-path price $p_i \neq p^*$, she does not revise her belief about the prices of firms $j \neq i$.

2.3 Optimal Consumer Search

Suppose that the consumer expects every firm to charge price $p^* \in \mathbb{R}_+$.

In theory, a search strategy can be very complex since the decision to purchase a product or terminate search can depend on the valuations for products that have been already sampled, bought or even returned. However, as it turns out, the optimal search strategy will follow a simple stationary cutoff rule: The consumer buys the product of firm i if and only if the observable (net) valuation $\theta_i - p_i$ is larger than some cutoff $\theta^* - p^*$ and terminates search and returns all previously bought products if and only if the total valuation $v_i - p_i$ is above some cutoff $v^* - p^*$.

A stationary cutoff strategy has the property that, in equilibrium, the consumer returns product i if and only if $\theta_i \geqslant \theta^*$ and $v_i < v^*$. In other words, the decision to return product i is independent of products $j \neq i$. Thus, the consumer can also return product i as soon as she observes $v_i < v^*$. Hence, it is without loss of generality to assume that a consumer has to return a product immediately before searching for other products when focusing on stationary cutoff strategies.

Next, we will derive conditions for the optimal cutoff values (θ^*, v^*), and thereby characterize the optimal cutoff search rule. In the appendix, I show that the optimal cutoff strategy is also the optimal strategy among all strategies. Let $U(\theta^*, v^*)$ denote the consumer's expected payoff under the (θ^*, v^*) cutoff search strategy. Note that under a stationary strategy, a consumer never revisits a sampled firm. Thus, the continuation payoff of continuing to search for new products does not depend on the values of past products and is equal to the expected payoff $U^* \equiv U(\theta^*, v^*)$. When deciding to continue the search or keep the current product with value v_i and price p_i, the consumer keeps the current product and terminates search if and only if

$$\underbrace{v_i - p_i}_{\text{payoff of keeping current product}} \geqslant \underbrace{U^* - p^* - c_r}_{\text{expected payoff of returning product } i \text{ and continuing to search}}.$$

Thus, the valuation cutoff v^* satisfies $v^* = U^* - c_r$.

Now suppose that a consumer has sampled a product and observed the price p_i and the observable valuation θ_i. When the consumer buys product i, she keeps the product if and only if $v_i \geqslant U^* - c_r + \Delta_i$ where $\Delta_i := p_i - p^*$. Otherwise, she returns the product and continues to search for new products. Thus, it is optimal for the consumer to buy product i with observable valuation θ_i and price p_i if and only if

$$\underbrace{\int_{U^*-c_r+\Delta_i}^{\infty} (v_i - p_i) dF(v_i \mid \theta_i) + F(U^* - c_r + \Delta_i \mid \theta_i) \cdot (U^* - p^* - c_r)}_{\text{expected payoff of buying product } i} \geqslant \underbrace{U^* - p^*}_{\substack{\text{Expected payoff of} \\ \text{continuing to search}}}.$$

which is equivalent to

$$W(v^* + \Delta_i \mid \theta_i) := \int_{v^*+\Delta_i}^{\infty} (v_i - v^* - \Delta_i) dF(v_i \mid \theta_i) \geqslant c_r.$$

It is straightforward to show that the left-hand-side is strictly increasing in θ_i with $\lim_{\theta \to -\infty} W(v^* + \Delta_i \mid \theta) = 0$ and $\lim_{\theta \to \infty} W(v^* + \Delta_i \mid \theta) = \infty$ which implies that for any $v^* + \Delta_i \in \mathbb{R}$ there exists a unique θ^* such that $W(v^* + \Delta_i \mid \theta^*) = c_r$. To fully characterize (θ^*, v^*) note that we can express the consumer's expected (gross) payoff $U(\theta^*, v^*)$ as follows:

$$U(\theta^*, v^*) = \mathbb{E}[v \mid v \geqslant v^*, \theta \geqslant \theta^*] - n_s(\theta^*, v^*) \cdot c_s - n_r(\theta^*, v^*) \cdot c_r$$

where $n_s(\theta^*, v^*)$ denotes the expected number of sampled items and $n_r(\theta^*, v^*)$ denotes the expected number of returned items. Letting $\bar{F}(\theta^*, v^*) = \int_{v^*}^{\infty} \int_{\theta^*}^{\infty} f(v \mid \theta) g(\theta) d\theta dv$ denote the probability that a sampled product is bought and kept and letting

$$R(\theta^*, v^*) = \int_{\infty}^{v^*} \int_{\theta^*}^{\infty} f(v \mid \theta) g(\theta) d\theta dv.$$

denote the probability that a sampled product is returned, we find that the expected number of sampled products is

$$n_s(\theta^*, v^*) = \sum_{i=1}^{\infty} (1 - \bar{F}(\theta^*, v^*))^{i-1} = \frac{1}{\bar{F}(\theta^*, v^*)}$$

and the expected number of returned products is

$$n_r(\theta^*, v^*) = \sum_{i=1}^{\infty}(1 - \bar{F}(\theta^*, v^*))^{i-1}R(\theta^*, v^*) = \frac{R(\theta^*, v^*)}{\bar{F}(\theta^*, v^*)}.$$

Plugging in n_s and n_r, we obtain

$$U(\theta^*, v^*) = \frac{1}{\bar{F}(\theta^*, v^*)}\left[\int_{v^*}^{\infty}\int_{\theta^*}^{\infty} vf(v \mid \theta)g(\theta)d\theta dv - c_s - R(\theta^*, v^*)c_r\right].$$

Plugging this into $v^* = U(\theta^*, v^*) - c_r$ and simplifying we get

$$\int_{\theta^*}^{\infty} W(v^* \mid \theta)g(\theta)d\theta = c_r \cdot (1 - G(\theta^*)) + c_s.$$

Proposition 2.1. *A strategy is optimal if and only if the consumer buys the product if and only if $\theta_i - p_i \geqslant \theta^* - p^*$ and returns the product and searches for a new product if and only if $v_i - p_i \leqslant v^* - p^*$, where the threshold values (θ^*, v^*) are the unique solution to (i) purchasing indifference*

$$W(v^* \mid \theta^*) = c_r,$$

and (ii) search indifference

$$\int_{\theta^*}^{\infty} W(v^* \mid \theta)dG(\theta) = c_s + c_r \cdot (1 - G(\theta^*)).$$

In the appendix, I show that firstly purchasing indifference and search indifference define a unique solution, and secondly we provide a rigorous definition of a search strategy and show that the consumer has no incentive to deviate from the cutoff strategy (θ^*, v^*) to a non-stationary strategy.

$W(v \mid \theta)$ describes the incremental benefit of buying a product when the observable valuation is θ and the consumer already owns a product with value v. If the consumer already owns a product, the costs of buying another product are equal to the return costs since the consumer has to return one additional item regardless of which product she decides to keep. Thus, (i) implies that at the optimal thresholds (θ^*, v^*) the consumer is indifferent between purchasing a product with observable valuation θ^* and not purchasing the product when she owns a product with value v^*. The search indifference condition (ii) implies that the expected incremental benefit $\int_{\theta^*}^{\infty} W(v^* \mid \theta)dG(\theta)$ when the consumer owns a product with value v^* and buys any product with observable valuation above θ^* is equal to the expected search and return costs.[4]

[4] The purchase and search indifference conditions can also be obtained by solving for the first-order conditions of $\max_{\theta^*, v^*} U(\theta^*, v^*)$. It is straightforward to show that

$$\frac{\partial U(\theta^*, v^*)}{\partial \theta^*} = 0 = \frac{\partial U(\theta^*, v^*)}{\partial v^*}$$

2.3.1 Asymptotics of Optimal Consumer Search

It is insightful to examine the effects on consumer search behavior under two distinct scenarios: Product returns are prohibited ($c_r \to \infty$) and product returns are costless ($c_r \to 0$).

Corollary 2.1. *The optimal search strategy* (θ^*, v^*) *satisfies*

(i) $\lim_{c_r \to \infty}(\theta^*, v^*) = (w_\infty, -\infty)$ *where* w_∞ *satisfies* $c_s = \int_{w_\infty}^{\infty}(\theta - w_\infty)dG(\theta)$, *and*

(ii) $\lim_{c_r \to 0}(\theta^*, v^*) = (-\infty, w_0)$ *where* $c_s = \int_{w_0}^{\infty}(v - w_0)dF_\theta(v)$ *and* $F_\theta(v) = \int_{-\infty}^{\infty} F(v \mid \theta)dG(\theta)$.

When product returns are prohibited, the consumer will never return any product ($v^* = -\infty$) since it is too costly to do so. In this case, consumers always receive the random payoff $\varepsilon \sim H$ independent of the decision which product to buy. Therefore, only the observable valuation θ of a product influences the search and purchase decision of consumers, and a consumer buys a product if and only if the observable valuation is above the reservation value w_∞ solving $c_s = \int_{w_\infty}^{\infty}(\theta - w_\infty)dG(\theta)$. Hence, consumer search reduces to the search problem of Anderson and Renault (1999); Wolinsky (1986) where the reservation value of searching is given by w_∞.

On the other hand, when product returns are costless ($c_r = 0$), it is without loss to observe the total valuation $v_i = \theta_i + \varepsilon_i$ by purchasing any sampled products ($\theta^* = -\infty$). Therefore, the search problem is analogous to a search problem without product returns where the consumer can directly observe the valuation $v_i \sim F_\theta$ where F_θ is the ex-ante distribution of the valuation for some product i. In this case, the consumer terminates search if and only if the valuation of a product is larger than the reservation valuation w_0 solving $c_s = \int_{w_0}^{\infty}(v - w_0)dF_\theta(v)$.

Thus, the canonical model of sequential search can capture both no product returns and free product returns but not costly product returns.

2.3.2 Comparative Statics of Optimal Consumer Search

How do the costs of searching for products or returning products affect consumer search behavior? The purchase threshold θ^* optimally trades off searching for another product and purchasing the product, and the return threshold v^* optimally trades off searching for other products subsequent to returning the current product and terminating the search by keeping the current product permanently.

First, we study how search costs affect the purchase and return thresholds which jointly determine consumer search. An increase in search costs c_s reduces the continuation payoff of searching for another product. However, higher search costs also reduce the continuation payoff of buying a product since there is a positive probability that the consumer

if and only if (θ^*, v^*) jointly solve purchase and search indifference.

returns the product and continues to search for another product. Since the consumer keeps the product with positive probability, the reduction of the continuation payoff of immediately searching for another product is larger, and, as a result, the consumer is more inclined to buy any given product when search costs are large, i.e. θ^* is decreasing in search costs. Analogously, the consumer is more inclined to keep a bought product and terminate the search, i.e. the return threshold v^* is also decreasing in c_s.

Proposition 2.2. *The purchase threshold θ^* and the return threshold v^* are decreasing in search costs. In particular,*

$$\frac{\partial \theta^*}{\partial c_s} = \frac{\partial v^*}{\partial c_s} = -\frac{1}{\bar{F}(\theta^*, v^*)} < 0.$$

Next, we study the effects of changes in the return costs c_r. An increase in return costs also reduces the continuation payoff of searching for another product because searching for another product entails buying and returning this product with positive probability. Further, larger return costs also directly disincentivize consumers to return products. Unsurprisingly, this causes the return threshold v^* to decrease.

For the purchase decision, there is a trade-off: The continuation payoff of searching for another product is smaller but also the continuation payoff of buying the current product is smaller since the consumer cannot freely return the current object in case she is not satisfied with the product. To see which effect dominates, it is instructive to consider the decision of a consumer at the margin: Suppose that a consumer searches for firm i and discovers an observable valuation of $\theta_i = \theta^*$ (and price $p_i = p^*$). In other words, the consumer is exactly indifferent between buying the product of firm i and searching for another firm. Now suppose we increase the return costs c_r marginally. Will the consumer with observable valuation θ^* buy the product or search for another firm? When the consumer buys a product she returns the product with probability $1 - F(v^* \mid \theta^*)$. On the other hand, when she samples the next product she returns the next product with probability $\mathbb{E}_\theta[1 - F(v^* \mid \theta) \mid \theta \geqslant \theta^*]$ conditional on buying the next product. Therefore, a consumer at the margin is more likely to return the current product than to return the next product. As a consequence, an increase in return costs induces the consumer to be more selective when it comes to buying a product, i.e. the purchase threshold θ^* is increasing in c_r.

Proposition 2.3. *The purchase threshold θ^* is increasing in return costs and the return threshold v^* is decreasing in return costs. In particular,*

$$\frac{\partial \theta^*}{\partial c_r} = \frac{1}{1 - F(v^* \mid \theta^*)} - \frac{1}{\mathbb{E}\left[1 - F(v^* \mid \theta) \mid \theta \geqslant \theta^*\right]} > 0$$

and

$$\frac{\partial v^*}{\partial c_r} = -\frac{1}{\mathbb{E}\left[1 - F(v^* \mid \theta) \mid \theta \geqslant \theta^*\right]} < 0.$$

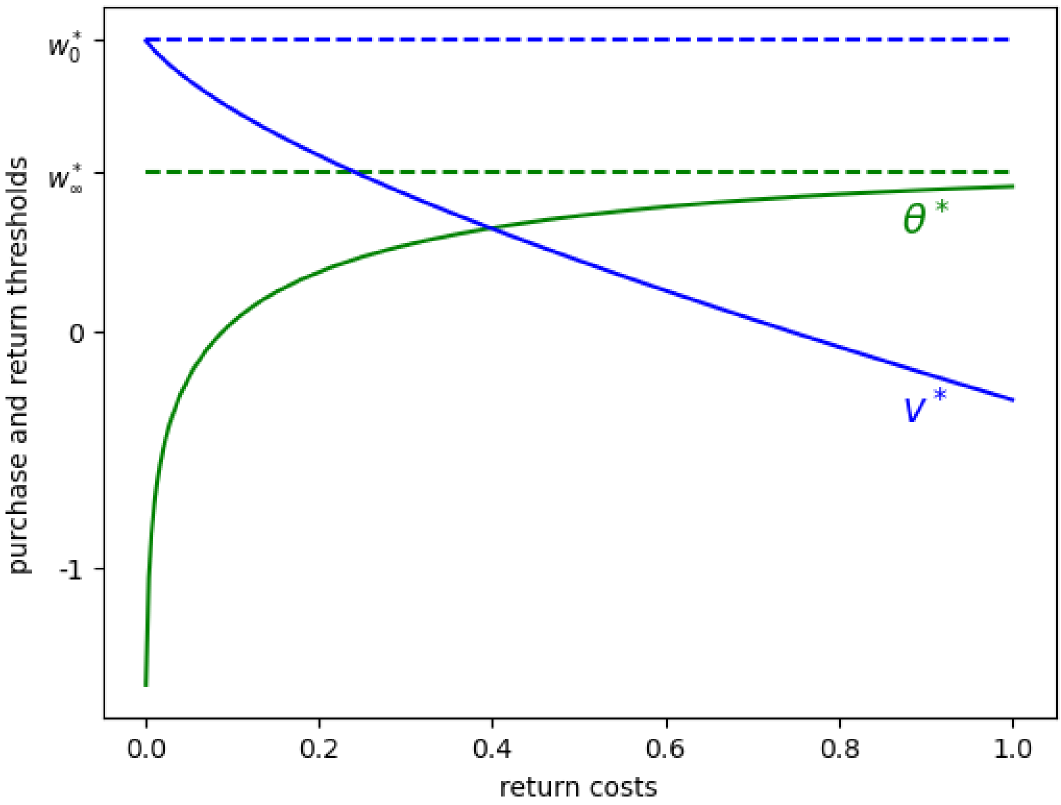

Figure 2.1: Optimal thresholds for varying return costs c_r where $(\theta_i, \varepsilon_i)_i \sim_{iid} \mathcal{N}(0,1)$ and $c_s = 0.15$.

Figure 2.1 illustrates Proposition 2.3 and Corollary 2.1 for the case when observable and residual valuations are standard normally distributed.

While it is difficult to empirically observe or estimate the threshold values (θ^*, v^*), it is much more convenient to observe the average number of products a consumer samples before buying a product and the average number of returned products. These quantities correspond to $n_s(\theta^*, v^*)$ and $n_r(\theta^*, v^*)$ in an equilibrium of the model, respectively.

We have seen in Proposition 2.2 that the consumer is more inclined to purchase a sampled product and less inclined to return a purchased product when search costs rise. Both effects imply that the expected number of sampled products decreases. Regarding the expected number of returned products $n_r(\theta^*, v^*)$ we can easily see that n_r is determined by the odds of a product return conditional on buying the product, i.e.

$$n_r(\theta^*, v^*) = \frac{R(\theta^*, v^*)/(1 - G(\theta^*))}{\bar{F}(\theta^*, v^*)/(1 - G(\theta^*))} = \frac{\mathbb{E}[F(v^* \mid \theta) \mid \theta \geqslant \theta^*]}{\mathbb{E}[1 - F(v^* \mid \theta) \mid \theta \geqslant \theta^*]}.$$

When a consumer is more inclined to purchase a product, the expected match valuation of a bought product decreases which increases the probability of a product return for a fixed return threshold v^*. On the other hand, the threshold v^* of a product return decreases in search costs as well which decreases the probability of a product return. The next Proposition shows that the latter effect dominates the first effect such that fewer products are returned when search costs are low. Intuitively, rising search costs affect the product return decision more acutely than the purchase decision because for the purchasing decision the consumer optimally trades off the decision between purchasing a product and searching for a new product where both options are affected by increasing search costs since purchasing a product entails a product return with subsequent search with positive probability.

Proposition 2.4. *The expected number of sampled products $n_s(\theta^*, v^*)$ and the expected number of returned products $n_r(\theta^*, v^*)$ are decreasing in search costs c_s.*

The consequences of rising return costs on the number of sampled and returned products are very similar to the effect of rising search costs. First, it is straightforward and unsurprising to see that larger return costs induce fewer product returns. The effect of rising return costs on the number of sampled products is more nuanced. By Proposition 2.3, the consumer is more selective regarding purchasing a product but less selective regarding returning the product. However, rising return costs affect the purchasing decision only indirectly via the probability that a product is returned after purchase, while the product return decision is directly affected. This intuition suggests that again the negative effect stemming from the product return decision outweighs the positive effect from the purchasing decision such that fewer products are sampled when product return costs are large. Proposition 2.5 confirms this intuition.

Proposition 2.5. *The expected number of sampled products $n_s(\theta^*, v^*)$ and the expected number of returned products $n_r(\theta^*, v^*)$ are decreasing in return costs c_r.*

2.4 Price Equilibrium and Welfare

I now characterize the symmetric equilibrium price. Suppose firm i assumes that every other firm charges price p^* and the consumer expects every firm to charge price p^*. As consumers do not observe deviations before searching, every firm i maximizes profits $\Pi(p_i)$ when other firms charge price p^* and the consumer searches for firm i. Since production costs are normalized to 0, profits of firm i are simply given by $\Pi(p_i) = D(p_i)p_i$ where $D(p_i)$ is firm i's demand function conditional on the consumer searches for firm i.

By Proposition 2.1, a consumer buys from firm i if and only if $\theta_i \geqslant \theta^* + \Delta_i$ and $v_i \geqslant v^* + \Delta_i$ where $\Delta_i = p_i - p^*$. Hence, firm i's demand function can be expressed by

$$D(p_i) \equiv \int_{\theta^*+\Delta_i} (1 - F(v^* + \Delta_i \mid \theta))\, g(\theta) d\theta = \bar{F}(\theta^* + \Delta_i,\, v^* + \Delta_i).$$

Any symmetric equilibrium requires $p_i = p^*$, and thus $D(p^*) = \bar{F}(\theta^*, v^*)$ denotes the equilibrium demand. The first-order condition of firm i's maximization problem $\max_{p_i} D(p_i)p_i$ is given by

$$p_i = \frac{D(p_i)}{-D'(p_i)}.$$

By inspection of $D(\cdot)$, it is immediate that

$$D'(p_i) = \frac{\partial \bar{F}(\theta^* + \Delta_i, v^* + \Delta_i)}{\partial \theta^*} + \frac{\partial \bar{F}(\theta^* + \Delta_i, v^* + \Delta_i)}{\partial v^*}.$$

Applying the symmetry condition $p_i = p^*$ yields that a candidate equilibrium price satisfies

$$p^* = \frac{\bar{F}(\theta^*, v^*)}{-\partial \bar{F}(\theta^*, v^*)/\partial \theta^* - \partial \bar{F}(\theta^*, v^*)/\partial v^*}.$$

As usual, the first-order condition only constitutes an equilibrium if the demand is "well-behaved". In particular, a sufficient condition common in the literature is that a firm's demand function is log-concave in the firm's own price. Log-concavity of the joint reliability function $\bar{F}$ suffices to ensure log-concavity of a firm's demand function.

Lemma 2.1. *Suppose that Assumption 2.1 is satisfied. Then, $D(p_i)$ is log-concave in p_i for all (p, θ^*, v^*).*

Proof. $\bar{F}(\theta, v)$ is log-concave by Assumption 2.1. Therefore, in particular

$$\log \bar{F}(\theta^* + p_i - p,\ v^* + p_i - p) \geqslant \alpha \log \bar{F}(\theta^* + p_i' - p,\ v^* + p_i' - p) + (1-\alpha) \log \bar{F}(\theta^* + p_i'' - p,\ v^* + p_i'' - p)$$

for all $\alpha \in [0, 1], p_i, p', p'' \in \mathbb{R}$ such that $p_i = \alpha p' + (1 - \alpha)p''$. This immediately implies

$$\log D(p_i) \geqslant \alpha \log D(p_i') + (1 - \alpha) \log D(p_i''),$$

i.e. the firm's demand function is log-concave in its price. □

Because of log-concavity of the demand function, the candidate equilibrium price p^* constitutes the unique symmetric equilibrium.

Proposition 2.6. *There exists a unique equilibrium price p^* which satisfies*

$$\begin{aligned} p^* &= \frac{\bar{F}(\theta^*, v^*)}{-\partial \bar{F}(\theta^*, v^*)/\partial \theta^* - \partial \bar{F}(\theta^*, v^*)/\partial v^*} \\ &= \frac{\int_{\theta^*} (1 - H(v^* - \theta)) g(\theta) d\theta}{-\int_{\theta^*} (1 - H(v^* - \theta)) g'(\theta) d\theta}. \end{aligned}$$

2.4.1 Comparative Statics of the Equilibrium Price

In search models with price discovery, it is common that larger search costs induce a larger equilibrium price. The next Proposition shows that this result continues to hold in this paper's model. As discussed, larger search costs imply that the threshold value θ^* of purchasing the product as well as the threshold value v^* of keeping the product decrease. Both effects tend to increase the equilibrium price.

Proposition 2.7. *The equilibrium price is increasing in search costs. That is,* $\frac{\partial p^*}{\partial c_s} > 0$.

In Section 2.5, I provide a detailed discussion showing that return costs can be interpreted as search costs of a two-stage consumer search model. Therefore, naive intuition suggests that the equilibrium price p^* is also increasing in return costs. However, this turns out not to be valid in general. Under mild conditions, the equilibrium price decreases when consumers become more selective when returning products ($\frac{\partial p^*}{\partial v^*} < 0$) and also decreases when consumers become more selective when purchasing products ($\frac{\partial p^*}{\partial \theta^*} < 0$). Applying the chain rule of differentiation yields

$$\frac{\partial p^*}{\partial c_r} = \underbrace{\frac{\partial p^*}{\partial \theta^*}}_{<0} \underbrace{\frac{\partial \theta^*}{\partial c_r}}_{>0} + \underbrace{\frac{\partial p^*}{\partial v^*}}_{<0} \underbrace{\frac{\partial v^*}{\partial c_r}}_{<0} .$$

This shows that two opposing forces influence equilibrium prices. First, higher return costs induce consumers to be more selective when purchasing products ($\frac{\partial \theta^*}{\partial c_r} > 0$), which lowers the equilibrium price ($\frac{\partial p^*}{\partial \theta^*}\frac{\partial \theta^*}{\partial c_r} < 0$). Second, higher return costs cause consumers to be less selective when returning products ($\frac{\partial v^*}{\partial c_r} < 0$), prompting firms to raise prices ($\frac{\partial p^*}{\partial v^*}\frac{\partial v^*}{\partial c_r} > 0$).

If the return costs are small ($c_r \approx 0$), then almost any consumer purchases the product ($\theta^* = -\infty$). Thus, marginally increasing the purchasing threshold has almost no effect ($g(\theta^*) \approx 0$) such that the increase in the equilibrium price due to a less selective product return decision dominates the negative effect on the equilibrium price due to a more selective purchasing decision.

Proposition 2.8. *There exists $\bar{c}$ such that the equilibrium price is increasing in return costs for all $c_r \leqslant \bar{c}$.*

However, in general, it is ambiguous which effect dominates such that the equilibrium price is not monotone. To see why, it is instructive to consider the asymptotic cases $c_r \to 0$ and $c_r \to \infty$. Corollary 2.1 argues that in both cases the model is equivalent to the model of Wolinsky (1986) where valuations are drawn from g if $c_r \to 0$ and drawn from $g * h$ if $c_r \to \infty$. If $\mathbb{E}[\varepsilon_i] = 0$, then $g * h$ is a mean-preserving spread of g.[5] However, a mean-preserving spread is not sufficient to induce a larger equilibrium price as Zhou (2022)

[5] $g * h$ denotes the convolution of g and h, i.e. $(g * h)(v) = \int_{-\infty}^{\infty} g(\theta)h(\theta - v)d\theta$.

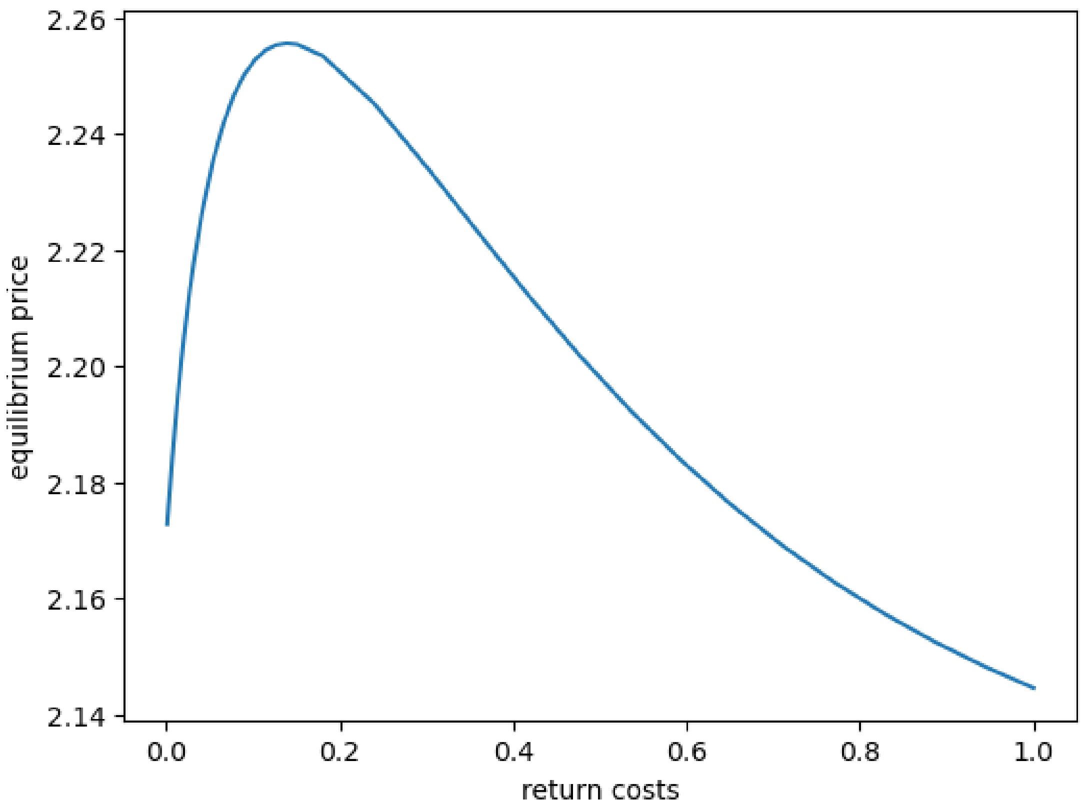

Figure 2.2: Equilibrium price p^* for varying return costs c_r where $(\theta_i, \varepsilon_i)_i \sim_{iid} \mathcal{N}(0,1)$ and $c_s = 0.75$.

shows, and furthermore there exists no simple stochastic order that induces monotone equilibrium prices. Thus, it is unsurprising that the equilibrium price is not monotone in this paper's model with costly product returns. Figure 2.2 depicts the equilibrium price for varying return costs when the observable and residual valuations are standard normally distributed.

2.4.2 Welfare Analysis

How does the ability of consumers to return products affect consumer surplus, industry profits, and total surplus?

First, since every consumer keeps exactly one product and there are no production costs, (equilibrium) industry profits are simply equal to the equilibrium price p^*. In particular, product returns do not necessarily diminish or increase the profits of sellers. For instance, Figure 2.2 illustrates that the equilibrium price, and consequently industry profits, are maximized at an intermediate level of return costs when observable and residual valuations follow a standard normal distribution and $c_s = 3/4$.

Second, consumer surplus $CS(c_s, c_r)$ is defined as the expected utility of a consumer.

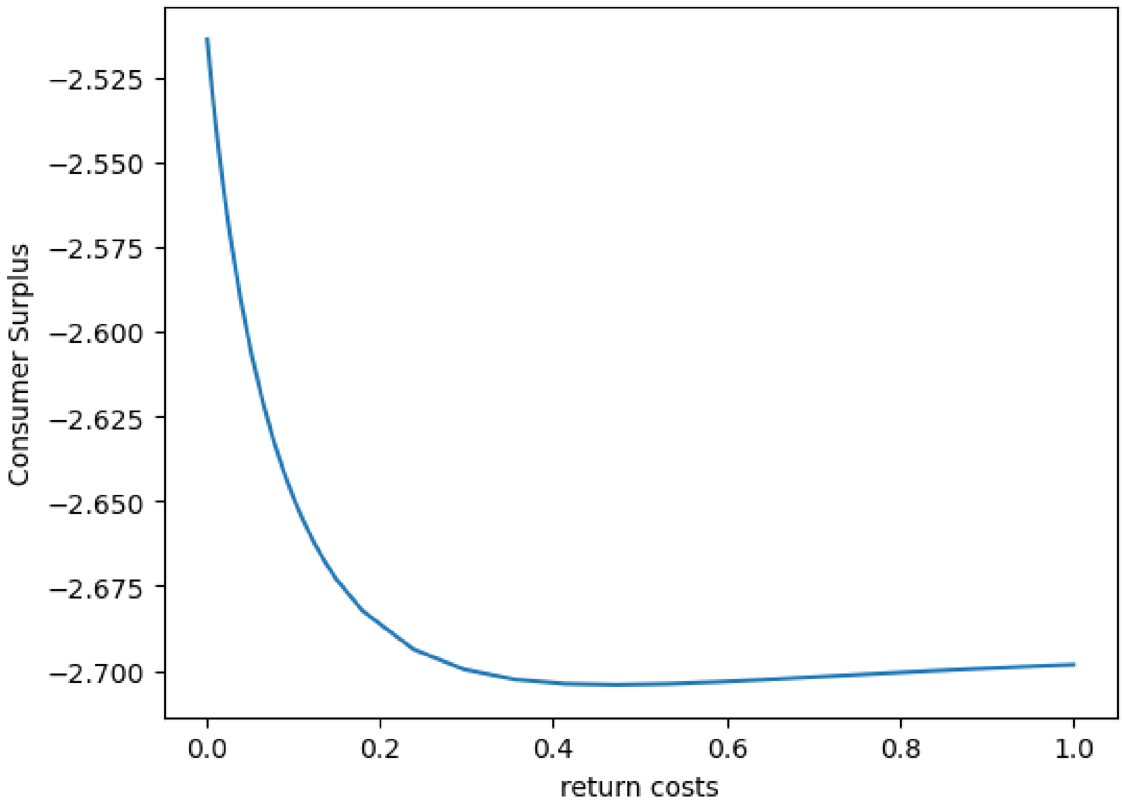

Figure 2.3: Consumer Surplus for varying return costs c_r where $(\theta_i, \varepsilon_i)_i \sim_{iid} \mathcal{N}(0,1)$ and $c_s = 0.75$.

That is,

$$\begin{aligned} CS(c_s, c_r) &:= \mathbb{E}[v \mid v \geqslant v^*,\, \theta \geqslant \theta^*] - n_s(\theta^*, v^*) \cdot c_s - n_r(\theta^*, v^*) \cdot c_r - p^* \\ &= v^* + c_r - p^* \end{aligned}$$

where the second equality stems from the search indifference condition of optimal search. Unsurprisingly, consumer surplus is decreasing in search costs. When it is more difficult to find new products, the consumer's expected match value of a kept product decreases, while the equilibrium price increases. The first effect persists when considering changes in the costs of returning a product. However, since the equilibrium price is non-monotonic, a decrease in the price of the product might offset increasing return costs. Figure 2.3 shows that consumer surplus can indeed be non-monotonic with respect to return costs.

Proposition 2.9. *If the equilibrium price is increasing in return costs, then consumer surplus is decreasing in return costs. In particular, there exists $\bar{c}$ such that consumer surplus is decreasing in return costs for all $c_r \leqslant \bar{c}$.*

Proof. The first statement immediately follows if $\partial(v^* + c_r)/\partial c_r < 0$ which is proven in Proposition 2.10. The second statement is an immediate implication of the first statement and Proposition 2.8. □

Lastly, total surplus $TS(c_s, c_r)$ is defined as the sum of consumer surplus and industry profits. That is,

$$TS(c_s, c_r) = CS(c_s, c_r) + \Pi(c_s, c_r) = v^* + c_r.$$

Although consumer surplus and industry profits are non-monotonous in the return costs, total surplus is decreasing in return costs. Since the equilibrium price does not affect total surplus, larger return costs as well as larger search costs hinder consumer from finding a well-matched product.

Proposition 2.10. *Total surplus is decreasing in return costs and search costs.*

Proof. Applying Proposition 2.3 and 2.2, respectively, immediately yields

$$\frac{\partial TS}{\partial c_r} = -\frac{1}{\mathbb{E}[1 - F(v^* \mid \theta) \mid \theta \geqslant \theta^*]} + 1 < 0$$

and

$$\frac{\partial TS}{\partial c_s} = -\frac{1}{\bar{F}(\theta^*, v^*)} < 0.$$

□

2.5 Extensions and Variations

2.5.1 Sequential Learning

In this section, I study the following variant of the model: Consumers do not have product return rights, but after they have searched for a product and observed the price p_i and the observable valuation θ_i, they can *investigate* the product by incurring costs c_r to also observe the residual valuation ε_i. Finally, only after observing ε_i can the consumer buy the product. Consequently, this variant of the model integrates a two-stage search process into the frameworks proposed by Wolinsky (1986) and Anderson and Renault (1999).

This model fits a number of sale mechanisms observed in practice: Consider a potential buyer searching for a pre-owned car. First, the buyer searches for cars and obtains superficial information about the car such as model, production year and mileage. All the information that is readily available constitutes the observable valuation θ_i. If the consumer is interested in the car, she can investigate the car more closely and do a test drive to observe the residual valuation ε_i. Only after doing the test drive can the consumer decide to buy the car or to look for another car to buy.

This model is equivalent to the model with product returns in the following way: It is optimal to buy a product in the product return model if and only if it is optimal to investigate a product in the sequential learning model, and it is optimal not to return a product in the product return model if and only if it is optimal to buy a product in the sequential learning model. To see this, let p^* be the price a consumer expects and let $U(\theta^*, v^*)$ denote the expected payoff (without the price) when the consumer investigates a product if and only if $\theta_i - p_i \geqslant \theta^* - p^*$ and buys a product if and only if $v_i - p_i \geqslant v^* - p^*$. Similar to Section 2.3, it follows that

$$U(\theta^*, v^*) = \mathbb{E}[v \mid v \geqslant v^*, \theta \geqslant \theta^*] - c_s \cdot \frac{1}{\bar{F}(\theta^*, v^*)} - c_r \cdot \frac{1 - G(\theta^*)}{\bar{F}(\theta^*, v^*)}.$$

It is optimal to buy product i if and only if

$$v_i - p_i \geqslant U(\theta^*, v^*) - p^*.$$

Therefore, it is optimal to investigate product i with observable valuation θ_i if and only if

$$\int_{U(\theta^*,v^*)+\Delta_i} (v_i - p_i) dF(v_i \mid \theta_i) + F(U(\theta^*, v^*) + \Delta_i \mid \theta_i) \cdot (U(\theta^*, v^*) - p^*) - c_r \geqslant U(\theta^*, v^*) - p^*.$$

Plugging in $U(\theta^*, v^*)$ and simplifying yields the same optimality conditions as in Proposition 2.6.

The intuitive reason for the equivalence is as follows: In the model with product returns, there is no loss for consumers to keep all bought items until deciding which product to keep. Thus, when the consumer has already bought one product, buying another product guarantees that the consumer has to eventually return one additional product which results in additional costs c_r. Thus, buying the product in the product return model reveals the residual valuation at costs c_r which is strategically equivalent to investigating the product in the sequential learning model. The only difference in the equilibrium of the models concerns the equilibrium payoff of the consumer. Since the consumer has to pay costs c_r every time she investigates a product in the model without product returns, while the consumer does not pay return costs c_r for the product she keeps in the model with product return, the equilibrium payoff without product returns is smaller by exactly c_r.

As a result, the demand functions of firms coincide in both models and thus also the equilibrium price coincides. Hence, the equilibrium price p^* is not monotone in the costs c_r of investigating a product, which is surprising because a robust result in the literature of sequential consumer search is that market prices are increasing in search costs if prices are not observable prior to search. Thus, this article contributes to the broader literature on price competition with sequential consumer search by showing that this result does not necessarily hold true if learning about a particular product unfolds sequentially and the consumer can decide to quit the process of learning and search for a new product instead.

2.5.2 Product Return Handling and Production Costs

In practice, firms incur costs both in producing products and handling product returns. In this extension, I examine how positive product return handling costs and production costs affect the equilibrium price.

First, suppose that a firm incurs costs $\beta \geqslant 0$ when producing a product. I assume that a firm can sell a returned product again such that β is only incurred when a consumer buys and keeps a product. Second, suppose that a firm incurs costs $\gamma \geqslant 0$ when handling a product return.

Product return handling and production costs do not affect consumer search, so the optimal consumer search behavior remains as defined by Proposition 2.1. Suppose that firm i and the consumer expect every other firm to charge price p^*. Firm i's profit maximization problem is given by

$$\max_{p_i} \bar{F}(\theta^* + p_i - p^*,\, v^* + p_i - p^*)(p_i - \beta) + \gamma \cdot R(\theta^* + p_i - p^*,\, v^* + p_i - p^*)$$

where

$$R(\theta^* + p_i - p^*,\, v^* + p_i - p^*) = \int_{\theta^*+\Delta_i} H(v^* + \Delta_i - \theta)g(\theta)d\theta$$

is the probability of a product return. It is straightforward to derive that if a symmetric equilibrium price $p^*(\beta, \gamma)$ exists, then

$$p^*(\beta,\gamma) = p^*(0,0) + \beta - \gamma \cdot (1 - \frac{g(\theta^*)}{\int_{\theta^*}(1 - H(v^* - \theta))(-g'(\theta))d\theta})$$

where $p^*(0,0)$ is the equilibrium price without product return handling and production costs (Proposition 2.6). Perhaps surprisingly, product return handling costs can lead to smaller market prices. Intuitively, large product return handling costs might incentivize firms to lower the number of product returns by decreasing the price.[6]

For example, consider the case where product return costs are small ($c_r \to 0$). By Corollary 2.1, this implies that a consumer always buys the product ($\theta^* \to -\infty$). Thus, $g(\theta^*) \to 0$ such that

$$\lim_{c_r \to 0} p^*(\beta,\gamma) = p^*(0,0) + \beta - \gamma.$$

For small return costs c_r, the effect of product return handling costs is exactly diametrically opposed to the effect of production costs. When consumers always buy the product and the only decision is whether to return or keep it, the firm either incurs production costs if the consumer keeps the product or product return handling costs if the consumer returns it.

[6] Note that this is not true in general. When a firm increases its price, fewer consumers buy the product. However, conditional on buying the product, more consumers might return it, meaning that the probability of a product return $R(\theta^* + p_i - p^*,\, v^* + p_i - p^*)$ does not necessarily decrease as p_i increases.

2.5.3 What if consumers can keep numerous products?

Throughout this article, we have assumed that consumers must return all but one product they have bought. This is obviously an unrealistic assumption. When a consumer finds that the price of a product is lower than the costs of returning it ($p_i < c_r$), it is more beneficial for the consumer to keep the product rather than return it because the refund the consumer would receive is not enough to cover the return costs.

A simple solution to this issue might be to impose that the return costs c_r are smaller than the equilibrium price p^*. However, the option for consumers to keep a product they do not want because returning it is not worthwhile could still disrupt the price equilibrium outlined in Proposition 2.6. This is because firms might be incentivized to set the price at $p_i = c_r$, ensuring that consumers retain the product even if they are dissatisfied with it after purchase.

To ensure that firms do not find it optimal to set prices at $p_i \leqslant c_r$, we can assume that their production costs are greater than the consumer's product return costs, i.e. $\beta > c_r$. This immediately implies that no firm would want to set prices at $p_i \leqslant c_r < \beta$ since doing so would result in a loss for each sale. Thus, we can interpret the assumption that consumers keep only one product permanently as focusing on a market where products have at least some minimal value such that the refund from returning a product outweighs the return costs.

2.6 Discussion

This paper has incorporated consumers' ability to return products into a general model of consumer search and price competition.

The findings show that optimal consumer search is defined by two threshold values. Consumers with valuations at these thresholds are indifferent between purchasing a product and continuing their search, as well as between keeping a product and searching for alternatives. Unlike research on product returns with single sellers, this chapter has shown that the option to return a product in a search market not only incentivizes consumers to make purchases but also encourages them to continue searching for alternatives. As a result, the probability that consumers purchase from a firm they have searched decreases with product returns.

The option of returning products makes consumers less selective in their initial purchases which raises market prices. However, consumers become more selective in returning products, prompting firms to lower prices. Overall, the relationship between the equilibrium price and product return costs is non-monotonic. This is surprising because return costs can be viewed as the costs of revealing a product's residual value before deciding whether to ultimately purchase the product or continue searching.

There are several directions for future research. First, future research could explore more general signal-valuation structures beyond those considered in this paper. In this paper, I analyze the specific case where the conditional CDF of the valuation based on the signal θ is given by $F(v \mid \theta) = H(v - \theta)$ for some CDF H. It is straightforward to show that (θ^*, v^*) remains the optimal stationary strategy even for general conditional distributions, as long as the following mild conditions hold: (i) (θ, v) has full support, (ii) $F(v \mid \theta') < F(v \mid \theta)$ for any $\theta' < \theta$ and (iii) $\lim_{\theta \to -\infty} F(v \mid \theta) = 1$ as well as $\lim_{\theta \to \infty} F(v \mid \theta) = 0$. Nevertheless, establishing comparative statics and showing existence and uniqueness of the equilibrium price for general signals is beyond the scope of this chapter.

Second, I have assumed that a consumer incurs costs c_r for each product she returns. In practice, the product return costs might be more nuanced. For instance, returning multiple products at once might be cheaper than returning each product separately. Another avenue for future research is to explore more general return and refund policies. Sellers are mandated in the EU to offer the option of free product returns within 14 days. However, many sellers compete by extending this period to 30 or even 100 days. Additionally, sellers vary in their exact implementation of warranty policies.

Third, I have restricted attention to infinitely many sellers. Greminger (2022) and Gibbard (2022) show that reservation value based policies remain optimal in related models with two-stage information acquisition and finitely many firms. A promising research direction would be to investigate whether the results hold when considering a finite number of sellers.

2.A Appendix

2.A.1 Proof of Remark 2.1

Proof. I show that the joint density f of (θ, v) is log-concave. This implies that the reliability function $\bar{F}(\theta, v) = \int_\theta^\infty \int_v^\infty f(\tilde{\theta}, \tilde{v}) d\tilde{\theta} d\tilde{v}$ is log-concave on its support (see for example An (1998)). Let ϕ denote the density function of the univariate standard normal distribution. The joint density f of (θ, v) is given by

$$\begin{aligned} f(\theta, v) &= \phi(\frac{\theta - \mu_\theta}{\sigma_\theta}) \cdot \phi(\frac{v - \theta - \mu_\varepsilon}{\sigma_\varepsilon}) \\ &= \frac{1}{2\sigma_\varepsilon \sigma_\theta \pi} \exp\left(-\frac{1}{2}(\frac{v - \theta - \mu_\varepsilon}{\sigma_\varepsilon})^2 - \frac{1}{2}(\frac{\theta - \mu_\varepsilon}{\sigma_\theta})^2 \right), \end{aligned}$$

Letting $c = \log(\frac{1}{2\sigma_\varepsilon \sigma_\theta \pi})$, we obtain

$$\log f(\theta, v) = c - \frac{1}{2}(\frac{v - \theta - \mu_\varepsilon}{\sigma_\varepsilon})^2 - \frac{1}{2}(\frac{\theta - \mu_\varepsilon}{\sigma_\theta})^2$$

which is concave on $\mathbb{R}^2$ if and only if

$$0 > x^2 \frac{\partial^2 \log f(\theta, v)}{\partial \theta^2} + 2xy \frac{\partial^2 \log f(\theta, v)}{\partial \theta \partial v} + y^2 \frac{\partial^2 \log f(\theta, v)}{\partial v^2}$$

for all $(x, y) \in \mathbb{R}^2$. Calculating the derivatives and plugging in yields

$$\begin{aligned} x^2 \frac{\partial^2 \log f(\theta, v)}{\partial \theta^2} + 2xy \frac{\partial^2 \log f(\theta, v)}{\partial \theta \partial v} + y^2 \frac{\partial^2 \log f(\theta, v)}{\partial v^2} & \\ &= x^2(-\frac{1}{\sigma_\varepsilon^2} - \frac{1}{\sigma_\theta^2}) + 2xy(\frac{1}{\sigma_\varepsilon^2}) + y^2(-\frac{1}{\sigma_\varepsilon^2}) \\ &= -\frac{x^2}{\sigma_\theta^2} - \frac{1}{\sigma_\varepsilon^2}(x - y)^2 \\ &< 0. \end{aligned}$$

□

2.A.2 Proof of Proposition 2.1

Outline In the following, I will provide a rigorous proof of the optimality of the cutoff strategy (θ^*, v^*) defined in Proposition 2.1. First, I show that (θ^*, v^*) is well-defined, i.e., search and purchasing indifference admit a unique solution. Next, I provide a rigorous definition of the search game and search strategies. This enables a rigorous redefinition of the cutoff strategy (θ^*, v^*) in terms of reservation values. Finally, I show that the consumer has no incentive to deviate from the optimal cutoff strategy.

To begin with, I show that (θ^*, v^*) is well-defined.

Lemma 2.2. *Let $(c_s, c_r) \in \mathbb{R}^2_+$. There exists a unique solution $(\theta^*, v^*) \in \mathbb{R}^2$ to the following system of non-linear equations:(i)*

$$W(v \mid \theta) = c_r,$$

and (ii)

$$\int_{\theta}^{\infty} W(v \mid \tilde{\theta}) dG(\tilde{\theta}) = c_s + c_r \cdot (1 - G(\theta)).$$

Proof. We impose $F(v \mid \theta) = H(v - \theta)$ such that $W(v \mid \theta) = \int_v^\infty (1 - F(z \mid \theta))dz$ is continuous in θ and, $\lim_{\theta* \to -\infty} W(v \mid \theta^*) = 0$ and $\lim_{\theta* \to \infty} W(v \mid \theta^*) \to \infty$ for any v. Also, note that

$$\frac{\partial W(v \mid \theta)}{\partial \theta} = -(1 - F(v \mid \theta)) < 0$$

for any $v \in \mathbb{R}$. Thus, the intermediate value theorem implies that for any $c_r > 0$ there exists a unique $\theta^*(v, c_r)$ such that $W(v \mid \theta^*(v, c_r)) = c_r$. Hence, the system of non-linear equations has a solution if and only if there exists $v^* \in \mathbb{R}$ satisfying

$$c_s = \int_{\theta^*(v^*, c_r)}^{\infty} (W(v^* \mid \theta) - c_r) dG(\theta).$$

We define $\mathcal{W}(v) = \int_{\theta^*(v,c_r)}^\infty (W(v \mid \theta) - c_r)dG(\theta)$ and we want to show that (i) $\mathcal{W}$ is strictly decreasing in v^*, (ii) $\mathcal{W}(-\infty) > c_s$ and (iii) $\mathcal{W}(\infty) < c_s$.

To show (i) note that

$$\begin{aligned}
\frac{\partial \mathcal{W}(v)}{\partial v} &= -\frac{\partial \theta^*(v, c_r)}{\partial v} (W(v \mid \theta^*(v, c_r)) - c_r) + \int_{\theta^*(v,c_r)}^{\infty} \frac{\partial W(v \mid \theta)}{\partial v} dG(\theta) \\
&= 0 - \int_{\theta^*(v,c_r)}^{\infty} (1 - F(v \mid \theta)) dG(\theta) \\
&< 0.
\end{aligned}$$

For (ii) and (iii) note that $\lim_{v \to -\infty} \theta^*(v, c_r) = -\infty$ and $\lim_{v \to \infty} \theta^*(v, c_r) = \infty$. Thus,

$$\begin{aligned}
\lim_{v \to -\infty} \mathcal{W}(v) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (1 - F(z \mid \theta)) dz dG(\theta) \\
&= \infty
\end{aligned}$$

and

$$\lim_{v \to \infty} \mathcal{W}(v) = 0.$$

Thus, the intermediate value theorem implies that for any $(c_r, c_s) \in \mathbb{R}^2_+$ there exists a unique solution (θ^*, v^*) such that $\mathcal{W}(v^*) = c_s$ and $\theta^* = \theta^*(v^*, c_r)$ which means that

(θ^*, v^*) is the unique solution to (i) and (ii). □

Next, I provide a definition of a search strategy. At any point of the game, we can partition the set $[0,1]$ of firms as follows: If the consumer has bought the product of firm i, then $i \in B$. If the consumer has sampled firm i but has not (yet) bought the product of firm i, then $i \in S$, and if the consumer has not yet sampled firm i, then $i \in U$. We denote the history h of search by $h = (U, S, B, (\theta_j)_{j\in\mathcal{S}}, (v_j)_{j\in\mathcal{B}}, (p_j)_{j\in\mathcal{S}\cup\mathcal{B}})$. The history of search includes all relevant information the consumer has about firms. That is, which firms she has sampled (S) and the observable valuation for sampled products $(\theta_j)_{j\in S}$, which products she has bought (B) and the (total) valuation for bought products $(v_j)_{j\in B}$ and finally the prices of firms she has encountered $(p_j)_{j\in\mathcal{S}\cup\mathcal{B}}$.

At any point of the game (that is, for any history) the consumer can decide between either sampling a firm that has not yet been sampled (denoted by action $i \in U$), buying a product of a sampled firm (denoted by action $i \in S$), and terminating search, keeping an already bought product, and returning all other bought products (denoted by action $i \in B$).[7] The search process ends once the consumer decides to take action $i \in B$. Let $\mathcal{H}$ denote the set of all histories. We define a strategy $\tilde{\sigma}$ as a function that maps histories to actions. That is $\tilde{\sigma} : \mathcal{H} \to U \cup S \cup B$.[8]

Now, I redefine the stationary cutoff strategy (θ^*, v^*). For any firm $i \in \{U, S, B\}$, the reservation value r_i of firm i is defined by

- If $i \in B$, then $r_i = v_i$
- If $i \in S$, then r_i solves $c_r = W(r_i \mid \theta_i)$
- If $i \in U$, then $r_i = v^*$.

Furthermore, define $p_i = p^*$ for any $i \in U$. Consider a strategy σ satisfying

$$\sigma(h) \in \arg\max_{i\in[0,1]} r_i - p_i.$$

First, we show that σ reduces to the cutoff strategy (θ^*, v^*) along the equilibrium path (that is, for any history h that can occur if the consumer does not deviate from σ).

Remark 2.2. Suppose that the consumer employs strategy σ. Along any equilibrium path, the consumer takes actions $i \in B$ if and only if $v_i - p_i \geqslant v^* - p^*$ and takes action $i \in S$ if and only if $\theta_i - p_i \geqslant \theta^* - p^*$.

Proof. Fix some history h. If the consumer did not deviate from σ prior to h, there can exist at most one $i \in \{S \cup B\}$ for which $r_i - p_i \geqslant v^* - p^*$.

[7] As noted, it is (weakly) dominant to return bought products when terminating search. Thus, it is without loss of generality to "bundle" product returns with the termination of search.

[8] For simplicity we consider pure strategies and assume that the consumer uniformly randomizes over actions whenever she is indifferent between actions.

We can easily see this by induction: At the start of the game this condition is vacuously satisfied since $S \cup B = \varnothing$. If this condition is satisfied for some history h, then either there exists no firm $i \in \{S \cup B\}$ for which $r_i - p_i \geqslant v - p^*$ in which case the consumer samples a new firm such that in the directly following history, there can be at most one firm (the new firm) for which $r_i - p_i \geqslant v - p^*$, or there exists exactly one firm $i \in \{S \cup B\}$ for which $r_i - p_i \geqslant v - p^*$. Then the consumer terminates search if $i \in B$ such that this condition holds since the game is over or the consumer buys the product of firm i if $i \in U$ in which case in the following history only firm i might satisfy $r_i - p_i \geqslant v - p^*$.

Since there exists at most one firm $i \in \{S \cup B\}$ for which $r_i - p_i \geqslant v^* - p^*$, the consumer either samples a new firm if no firm exists, or the consumer takes action $i \in B$ if and only if $v_i - p_i \geqslant v^* - p^*$ or she takes action $i \in S$ if $r_i \geqslant v^* + \Delta_i$ which is equivalent to $c_r \leqslant W(v^* + \Delta_i \mid \theta_i)$. Integration by parts shows that $W(v^* + \Delta_i \mid \theta_i) = \int_{v^*+\Delta_i}(1 - H(z - \theta_i))dz$ and therefore

$$W(v^* + \Delta_i \mid \theta^* + \Delta_i) = W(v^* \mid \theta^*) = c_r$$

which implies $W(v^* + \Delta_i \mid \theta_i) \geqslant c_r$ if and only if $\theta_i \geqslant \theta^* + \Delta_i$. Hence, we have shown that the consumer takes action $i \in S$ if and only if $\theta_i - p_i \geqslant \theta^* - p^*$. □

Thus, Proposition 2.1 is immediately implied by the following result.

Proposition. *Let $p_i := p^*$ for all $i \in U$. The strategy σ defined by*

$$\sigma(h) = \arg\max_{i \in [0,1]} r_i - p_i$$

is optimal.

Proof. By the one-shot deviation principle, a strategy σ is optimal if and only if the consumer has no incentive to deviate from σ at a single history h.

Hence, fix a history h and suppose the consumer has not yet deviated from σ prior to h. We will show that this implies that the consumer has no incentive to deviate from $\sigma(h)$. That is, we verify that there exists no profitable deviation for the three cases (a) $\sigma(h) \in B$, (b) $\sigma(h) \in U$ and (c) $\sigma(h) \in S$.

(a) Suppose that $\sigma(h) = i \in B$. That is, σ prescribes to terminate search, keep product i and return all other bought products. The continuation payoff from following σ is

$$V(h) = v_i - p_i - (\mid B \mid -1) \cdot c_r$$

since the net payoff from buying i is $v_i - p_i$ and the consumer return costs $(\mid B \mid -1) \cdot c_r$ from returning any product $j \in B$ except product i .

Let $\theta^*(r)$ solve $c_r = W(r \mid \theta^*(r))$. Note that $\sigma(h) = i \in B$ requires $r_i - p_i \geqslant r_j - p_j$ for all $j \in [0, 1]$ which implies (i) $v_i - p_i \geqslant v_j - p_j$ for all $j \in B$, (ii) $\theta_j \leqslant \theta^*(v_i - p_i + p_j)$ for all $j \in S$, (iii) $v_i - p_i \geqslant v^* - p^*$. (i) implies that deviating to some $j \in B$ is not profitable. Deviating to some $j \in S$ (and following σ thereafter) yields continuation payoff

$$\begin{aligned}\tilde{V}_{j\in S}(h) &= \int_{v_i - p_i + p_j} (v_j - p_j - \mid B \mid c_r) dF(v_j \mid \theta_j) + F(v_i - p_i + p_j)(v_i - p_i - \mid B \mid \cdot c_r) \\ &= -c_r + V(h) + W(v_i - p_i + p_j \mid \theta_j) \\ &\leqslant V(h)\end{aligned}$$

where the inequality is implied by $\theta_j \leqslant \theta^*(v_i - p_i + p_j)$. Thus, it is not profitable to deviate to $j \in S$. Lastly, consider a deviation to $j \in U$ and suppose that the consumer follows σ thereafter. Then, she will buy product j if and only if $r_j - p^* \geqslant v_i - p_i$ which is satisfied if and only if $\theta_j \geqslant \theta^*(v_i - \Delta_i)$. Otherwise she will keep product i and terminate search. In case she buys product j she will keep product j if and only if $v_j \geqslant v_i - \Delta_i$ and otherwise she keeps product i. Thus, a deviation to $j \in U$ yields continuation payoff

$$\begin{aligned}\tilde{V}_{j\in U}(h) &= -c_s + G(\theta^*(v_i - \Delta_i))V(h) \\ &\quad + \int_{\theta^*(v_i - \Delta_i)} \left(-c_r + \int_{v_i - \Delta_i} (v - p^*) dF(v \mid \theta) + F(v_i - \Delta_i \mid \theta) V(h) \right) dG(\theta) \\ &= -c_s + G(\theta^*(v_i - \Delta_i))V(h) + \int_{\theta^*(v_i - \Delta_i)} (-c_r + V(h) + W(v_i - \Delta_i \mid \theta)\, dG(\theta) \\ &= V(h) - c_s + \int_{\theta^*(v_i - \Delta_i)} (W(v_i - \Delta_i \mid \theta) - c_r)\, dG(\theta).\end{aligned}$$

Plugging in $c_s = \int_{\theta^*} (W(v^* \mid \theta) - c_r)\, dG(\theta)$ where $\theta^* \equiv \theta^*(v^*)$ yields

$$\begin{aligned}\tilde{V}_{j\in U}(h) &= V(h) - \int_{\theta^*} (W(v^* \mid \theta) - c_r)\, dG(\theta) + \int_{\theta^*(v_i - \Delta_i)} (W(v_i - \Delta_i \mid \theta) - c_r)\, dG(\theta) \\ &\leqslant V(h)\end{aligned}$$

where the inequality is implied by $v^* \leqslant v_i - p_i + p^*$.[9] Hence, it is not profitable to deviate when strategy σ prescribes to keep product i and terminate search ($\sigma(h) = i \in B$).

(b) Suppose that σ prescribes to sample a new firm ($\sigma(h) \in U$). This requires that $v^* - p^* \geqslant v_i - p_i$ for all $i \in B$ and $\theta_i \leqslant \theta^*(v^* - p^* + p_i)$ for all $i \in S$. Following $\sigma(h) \in U$

[9] Note that $\lambda(x) := \int_{\theta^*(x)} (W(x \mid \theta) - c_r)\, dG(\theta)$ is decreasing in x since

$$\lambda'(x) = 0 + \int_{\theta^*(x)} \frac{\partial W(x \mid \theta)}{\partial x} dG(\theta) < 0.$$

yields continuation payoff

$$\begin{aligned} V(h) &= U^* - p^* - \mid B \mid \cdot c_r \\ &= v^* - p^* - (\mid B \mid -1) \cdot c_r. \end{aligned}$$

Deviating to $i \in B$ yields continuation payoff $v_i - p_i - (\mid B \mid -1) \cdot c_r$ which is not profitable since $v^* - p^* \geqslant v_i - p_i$. Thus, consider a deviation to $i \in S$. This yields continuation payoff

$$\begin{aligned} V_{i \in S}(h) &= \int_{v^* + \Delta_i} (v_i - p_i - \mid B \mid \cdot c_r) dF(v_i \mid \theta_i) + F(v^* + \Delta_i \mid \theta_i)(-c_r + V(h)) \\ &= V(h) - c_r + W(v^* - p^* + p_i \mid \theta_i) \\ &\leqslant V(h). \end{aligned}$$

where the inequality is implied by $\theta_i \leqslant \theta^*(v^* - p^* + p_i)$. Hence, we have shown that it is not profitable to deviate when σ prescribes to sample a new firm.

(c) Suppose that $\sigma(h) = i \in S$. That is, σ prescribes to buy the product of firm i. This requires that (i) $\theta_i \geqslant \theta_j$ for all $j \in U$, (ii) $\theta_i \geqslant \theta^*(v_j - p_j + p_i)$ for all $j \in B$ and (iii) $\theta_i \geqslant \theta^*(v^* + \Delta_i)$. We will show that there exists no incentive to deviate to either $j \in U$ or $j \in B$ if $\theta_i = \theta^*(v^* + \Delta_i)$. It is straightforward to show that if the consumer has no incentive to deviate for $\theta_i = \theta^*(v^* + \Delta_i)$, she also does not deviate for $\theta_i > \theta^*(v^* + \Delta_i)$. Hence, suppose that $\theta_i = \theta^*(v^* + \Delta_i)$. Then, the continuation payoff from following $\sigma(h) = i \in S$ is

$$\begin{aligned} V(h) &= \int_{v^* + \Delta_i} (v_i - p_i - \mid B \mid \cdot c_r) dF(v_i \mid \theta_i) + F(v^* + \Delta_i)(U^* - p^* - (\mid B \mid +1) \cdot c_r) \\ &= v - p^* - \mid B \mid \cdot c_r - W(v^* + \Delta_i \mid \theta^*(v + \Delta_i)) \\ &= U^* - p^* - \mid B \mid \cdot c_r \end{aligned}$$

which is intuitive because the consumer is indifferent between buying product i and continuing to search for new firms for $\theta_i = \theta^*(v^* + \Delta_i)$. This implies in particular that the consumer has no incentive to deviate to $j \in U$ if $r_i > v^*$ and $i \in U$. Next, consider a deviation to $j \in B$. That is, the consumer terminates search and keeps an already bought product j and returns all other products instead of buying product $i \in S$. Since $j \in B$, and the consumer has not deviated previously it follows that $v_j - p_j \leqslant v^* - p^*$. Therefore, the continuation payoff of deviating to $j \in B$ is

$$\begin{aligned} v_j - p_j - (\mid B \mid -1) \cdot c_r &\leqslant v^* - p^* - (\mid B \mid -1) \cdot c_r \\ &= U^* - p^* - \mid B \mid \cdot c_r \\ &= V(h). \end{aligned}$$

Lastly, consider a deviation to $j \in U$. That is, the consumer decides to buy the product from another firm j instead of firm i. Since the consumer has not deviated prior to history h, we know that $\theta_j \leqslant \theta^*(v^* + \Delta_j)$. It can be easily shown that the continuation payoff of deviating to $j \in U$ is given by

$$\tilde{V}_{j \in U}(h) = V(h) - c_r + W(v_i + \Delta_j \mid \theta_j) \leqslant V(h).$$

Thus we have shown that the consumer has no incentive to deviate from $\sigma(h)$ for any history h which can occur under the assumption that the consumer has not deviated from σ previously. Then, σ is optimal by the one-shot deviation principle. □

2.A.3 Proof of Corollary 2.1

Proof. From Proposition 2.1, we know that the optimal thresholds (θ^*, v^*) satisfy purchase indifference

$$W(v^* \mid \theta^*) = c_r$$

and search indifference

$$\int_{\theta^*}^{\infty} (W(v^* \mid \theta) - c_r) dG(\theta) = c_s.$$

Note that

$$W(v \mid \theta) = \int_v^{\infty} (1 - F(z \mid \theta)) dz = \int_v^{\infty} (1 - H(z - \theta)) dz$$

(i) Suppose that $c_r \to \infty$. Then $W(v^* \mid \theta^*) = c_r$ implies

$$\lim_{c_r \to \infty} \int_{v^*}^{\infty} (1 - H(z - \theta^*)) dz = \infty$$

which can only be satisfied if $\lim_{c_r \to \infty} v^* = -\infty$ or $\lim_{c_r \to \infty} \theta^* = \infty$. However, if $\lim_{c_r \to \infty} v^* > -\infty$ and $\lim_{c_r \to \infty} \theta^* = \infty$, then

$$\lim_{c_r \to \infty} \int_{\theta^*}^{\infty} W(v^* \mid \theta) dG(\theta) = 0$$

which contradicts search indifference. Therefore, the optimal strategy (θ^*, v^*) must satisfy

$\lim_{c_r \to 0} v^* = -\infty$. Hence,

$$\begin{aligned}
0 &= \lim_{c_r \to \infty} W(v^* \mid \theta^*) - c_r \\
&= \lim_{c_r \to \infty} \left(\int_{v^*}^{\infty} z f(z \mid \theta^*) dz - (1 - F(v^* \mid \theta^*) v^* - c_r \right) \\
&= \int_{-\infty}^{\infty} z f(z \mid \lim_{c_r \to \infty} \theta^*) dz - \lim_{c_r \to \infty} [(1 - F(v^* \mid \theta^*) v^* - c_r] \\
&= \mathbb{E}[v \mid \theta = \lim_{c_r \to \infty} \theta^*] - \lim_{c_r \to \infty} ((1 - F(v^* \mid \theta^*) v^* - c_r)
\end{aligned}$$

provided that $\lim_{c_r \to \infty} \theta^* \in \mathbb{R}$. Thus,

$$\mathbb{E}[v \mid \theta = \lim_{c_r \to \infty} \theta^*] = \lim_{c_r \to \infty} ((1 - F(v^* \mid \theta^*) v^* - c_r).$$

Search indifference is satisfied if and only if $\int_{\theta^*}^{\infty} (W(v^* \mid \theta) - c_r) g(\theta) d\theta = c_s$. Letting $c_r \to \infty$, we get

$$\begin{aligned}
\lim_{c_r \to \infty} & \int_{\theta^*}^{\infty} (W(v^* \mid \theta) - c_r) g(\theta) d\theta \\
&= \lim_{c_r \to \infty} \int_{\theta^*}^{\infty} \left(\int_{v^*}^{\infty} z f(z \mid \theta) dz - (1 - F(v^* \mid \theta) v^* - c_r \right) g(\theta) d\theta \\
&= \int_{\lim_{c_r \to \infty} \theta^*}^{\infty} \left(\int_{-\infty}^{\infty} z f(z \mid \theta) dz - \lim_{c_r \to \infty} (1 - F(v^* \mid \theta) v^* - c_r \right) g(\theta) d\theta \\
&= \int_{\lim_{c_r \to \infty} \theta^*}^{\infty} \left(\mathbb{E}[v \mid \theta] - \mathbb{E}[v \mid \lim_{c_r \to \infty} \theta^*] \right) g(\theta) d\theta \\
&= \int_{\lim_{c_r \to \infty} \theta^*}^{\infty} \left(\theta + \mathbb{E}[\varepsilon] - (\lim_{c_r \to \infty} \theta^* + \mathbb{E}[\varepsilon]) \right) g(\theta) d\theta \\
&= \int_{\lim_{c_r \to \infty} \theta^*}^{\infty} \left(\theta - \lim_{c_r \to \infty} \theta^* \right) g(\theta) d\theta \\
&= c_s
\end{aligned}$$

where the third equality uses $\mathbb{E}[v \mid \theta = \lim_{c_r \to \infty} \theta^*] = \lim_{c_r \to \infty} ((1 - F(v^* \mid \theta^*) v^* - c_r)$ and the fourth equality stems from $\mathbb{E}[v \mid \theta] = \theta + \mathbb{E}[\varepsilon]$.

(ii) Suppose $c_r \to 0$. Then $W(v^* \mid \theta^*) = c_r$ implies

$$\lim_{c_r \to 0} \int_{v^*}^{\infty} (1 - H(z - \theta^*)) dz = 0$$

which can only be satisfied if $\lim_{c_r \to 0} v^* = \infty$ or $\lim_{c_r \to 0} \theta^* = -\infty$. However, if $\lim_{c_r \to 0} v^* = \infty$, we obtain $\lim_{c_r \to 0} \int_{\theta^*}^{\infty} W(v^* \mid \theta) dG(\theta) = 0$ for any $\theta^* \in \mathbb{R} \cup \{-\infty, \infty\}$ which is a contradiction to search indifference. Thus, the optimal strategy (θ^*, v^*) must satisfy

$\lim_{c_r \to 0} \theta^* = -\infty$. Consequently, search indifference implies

$$\lim_{c_r \to 0} \int_{-\infty}^{\infty} W(v^* \mid \theta) dG(\theta) = c_s$$

It is straightforward to show that

$$\int_{-\infty}^{\infty} W(v^* \mid \theta) dG(\theta) = \int_{v^*}^{\infty} (v - v^*) \left(\int_{-\infty}^{\infty} f(v \mid \theta) dG(\theta) \right) dv,$$

and therefore

$$\lim_{c_r \to 0} \int_{v^*}^{\infty} (v - v^*) f_\theta(v) dv = \int_{\lim_{c_r \to 0} v^*}^{\infty} (v - \lim_{c_r \to 0} v^*) f_\theta(v) dv = c_s.$$

□

2.A.4 Proof of Proposition 2.2

Proof. First note that

$$\frac{\partial W(v \mid \theta)}{\partial v} = -(1 - F(v \mid \theta))$$

and

$$\frac{\partial W(v \mid \theta)}{\partial \theta} = \int_{v}^{\infty} f(z \mid \theta) dz = 1 - F(v \mid \theta) = -\frac{\partial W(\theta, v)}{\partial v}.$$

Next, by Proposition 2.1, (θ^*, v^*) satisfy $W(v^* \mid \theta^*) = c_r$ (purchase indifference) and $\int_{\theta^*}^{\infty} (W(v^* \mid \theta) - c_r) dG(\theta) = c_s$ (search indifference). Differentiating purchase indifference w.r.t. c_s yields

$$\begin{aligned} 0 &= \frac{\partial W(v^* \mid \theta^*)}{\partial \theta^*} \frac{\partial \theta^*}{\partial c_s} + \frac{\partial W(v^* \mid \theta^*)}{\partial v^*} \frac{\partial v^*}{\partial c_s} \\ &= (1 - F(v^* \mid \theta^*)) \left(\frac{\partial \theta^*}{\partial c_s} - \frac{\partial v^*}{\partial c_s} \right) \end{aligned}$$

which immediately implies $\frac{\partial \theta^*}{\partial c_s} = \frac{\partial v^*}{\partial c_s}$. Differentiating search indifference w.r.t. c_s yields

$$1 = \int_{\theta^*}^{\infty} \frac{\partial W(v^* \mid \theta)}{\partial v^*} \frac{\partial v^*}{\partial c_s} dG(\theta) = -\frac{\partial v^*}{\partial c_s} \int_{\theta^*}^{\infty} (1 - F(v^* \mid \theta) dG(\theta).$$

Hence, we obtain

$$\frac{\partial \theta^*}{\partial c_s} = \frac{\partial v^*}{\partial c_s} = -\frac{1}{\bar{F}(\theta^*, v^*)}.$$

where $\bar{F}(\theta^*, v^*) = \int_{\theta^*}^{\infty} (1 - F(v^* \mid \theta) dG(\theta)$ denotes the probability that a consumer buys a sampled product. □

2.A.5 Proof of Proposition 2.3

Proof. Proposition 2.1 shows that (θ^*, v^*) satisfies $W(v^* \mid \theta^*) = c_r$ (purchase indifference) and $\int_{\theta^*}^{\infty}(W(v^* \mid \theta) - c_r)dG(\theta) = c_s$ (search indifference). Differentiating purchase indifference w.r.t. c_r yields

$$\begin{aligned} 1 &= \frac{\partial W(\theta^*, v^*)}{\partial \theta^*}\frac{\partial \theta^*}{\partial c_r} + \frac{\partial W(\theta^*, v^*)}{\partial v^*}\frac{\partial v^*}{\partial c_r} \\ &= (1 - F(v^* \mid \theta^*))\left[\frac{\partial \theta^*}{\partial c_r} - \frac{\partial v^*}{\partial c_r}\right] \end{aligned}$$

which is equivalent to

$$\frac{\partial \theta^*}{\partial c_r} = \frac{1}{1 - F(v^* \mid \theta^*)} + \frac{\partial v^*}{\partial c_r}.$$

Differentiating search indifference w.r.t. c_r we obtain

$$\begin{aligned} 0 &= \int_{\theta^*}^{\bar{\theta}}\left(\frac{\partial W(\theta, v^*)}{\partial v^*}\frac{\partial v^*}{\partial c_r} - 1\right)dG(\theta) \\ &= -\frac{\partial v^*}{\partial c_r}\int_{\theta^*}^{\infty}(1 - F(v^* \mid \theta))dG(\theta) - (1 - G(\theta^*)) \end{aligned}$$

This is equivalent to

$$\frac{\partial v^*}{\partial c_r} = -\frac{1}{\mathbb{E}[1 - F(v^* \mid \theta) \mid \theta \geqslant \theta^*]}$$

which in turn yields

$$\frac{\partial \theta^*}{\partial c_r} = \frac{1}{1 - F(v^* \mid \theta^*)} - \frac{1}{\mathbb{E}[1 - F(v^* \mid \theta) \mid \theta \geqslant \theta^*]} > 0.$$

□

2.A.6 Proof of Lemma 2.3

The following result is useful to prove Propositions 2.4 and 2.5. Also recall that Assumption 2.1 implies that $\bar{F}_v = 1 - G$ is log-concave.

Lemma 2.3. *Suppose that* $1 - G$ *is log-concave. Then,*

$$\int_{\theta^*}^{\infty} f(v^* \mid \theta)g(\theta)d\theta \geqslant \left(F(v^* \mid \theta^*) - \mathbb{E}[F(v^* \mid \theta) \mid \theta \geqslant \theta^*]\right)g(\theta^*)$$

for any $(\theta^*, v^*) \in \mathbb{R}^2$*. This implies in particular,*

$$\frac{\bar{F}(\theta^*, v^*)}{-\frac{\partial \bar{F}(\theta^*, v^*)}{\partial \theta^*} - \frac{\partial \bar{F}(\theta^*, v^*)}{\partial v^*}} \leqslant \frac{1 - G(\theta^*)}{g(\theta^*)}.$$

Proof. Note that

$$\begin{aligned}
F(v^* \mid \theta^*) - \mathbb{E}\left[F(v^* \mid \theta) \mid \theta \geqslant \theta^*\right] &= F(v^* \mid \theta^*) - \int_{\theta*}^{\infty} F(v^* \mid \theta)\frac{g(\theta)}{1-G(\theta^*)}d\theta \\
&= F(v^* \mid \theta^*) - \frac{[F(v^* \mid \theta)G(\theta)]_{\theta*}^{\infty} + \int_{\theta*}^{\infty} f(v^* \mid \theta)G(\theta)d\theta}{1-G(\theta^*)} \\
&= F(v^* \mid \theta^*)(1 + \frac{G(\theta^*)}{1-G(\theta^*)}) - \int_{\theta*}^{\infty} f(v^* \mid \theta)\frac{G(\theta)}{1-G(\theta^*)}d\theta \\
&= \int_{\theta*}^{\infty} f(v^* \mid \theta)\left(\frac{1-G(\theta)}{1-G(\theta^*)}\right)d\theta
\end{aligned}$$

where the second line follows from integration by parts, and the fourth line follows from

$$\begin{aligned}
F(v^* \mid \theta^*) &= \int_{-\infty}^{v*} h(v-\theta^*)dv \\
&= \int_{\theta*}^{\infty} h(v^*-\theta)d\theta = \int_{\theta*}^{\infty} f(v^* \mid \theta)d\theta.
\end{aligned}$$

Thus,

$$\begin{aligned}
\left(F(v^* \mid \theta^*) - \mathbb{E}\left[F(v^* \mid \theta) \mid \theta \geqslant \theta^*\right]\right)g(\theta^*) &= \int_{\theta*}^{\infty} f(v^* \mid \theta)\left(\frac{1-G(\theta)}{1-G(\theta^*)}g(\theta^*)\right)d\theta \\
&\leqslant \int_{\theta*}^{\infty} f(v^* \mid \theta)g(\theta)d\theta
\end{aligned}$$

where the inequality can be easily verified since log-concavity of $1-G$ implies $\frac{g(\theta^*)}{1-G(\theta^*)} \leqslant \frac{g(\theta)}{1-G(\theta)}$ for all $\theta \geqslant \theta^*$.

Next, note that applying this inequality yields

$$\begin{aligned}
\frac{\bar{F}}{-\frac{\partial\bar{F}}{\partial\theta^*} - \frac{\partial\bar{F}}{\partial v^*}} &= \frac{\int_{\theta*}^{\infty}(1-F(v^* \mid \theta))g(\theta)d\theta}{(1-F(v^* \mid \theta^*))g(\theta^*) + \int_{\theta*}^{\infty} f(v^* \mid \theta)g(\theta)d\theta} \\
&\leqslant \frac{\int_{\theta*}^{\infty}(1-F(v^* \mid \theta))g(\theta)d\theta}{(1-F(v^* \mid \theta^*))g(\theta^*) + \left(F(v^* \mid \theta^*) - \mathbb{E}\left[F(v^* \mid \theta) \mid \theta \geqslant \theta^*\right]\right)g(\theta^*)} \\
&= \frac{1}{g(\theta^*)} \cdot \frac{\int_{\theta*}^{\infty}(1-F(v^* \mid \theta))g(\theta)d\theta}{\mathbb{E}\left[1-F(v^* \mid \theta) \mid \theta \geqslant \theta^*\right]} \\
&= \frac{1-G(\theta^*)}{g(\theta^*)}.
\end{aligned}$$

□

2.A.7 Proof of Proposition 2.4

Proof. Remember that $n_s(\theta^*, v^*) = \bar{F}(\theta^*, v^*)^{-1}$ where $\bar{F}(\theta^*, v^*) = \int_{\theta^*}^{\infty}(1-F(v^* \mid \theta))g(\theta)d\theta$. Therefore, n_s is decreasing in c_s if and only if $\bar{F}$ is increasing in c_s. We have

$$\begin{aligned}\frac{\partial \bar{F}(\theta^*, v^*)}{\partial c_s} &= \frac{\partial \bar{F}(\theta^*, v^*)}{\partial \theta^*}\frac{\partial \theta^*}{\partial c_s} + \frac{\partial \bar{F}(\theta^*, v^*)}{\partial v^*}\frac{\partial v^*}{\partial c_s} \\ &= -\frac{1}{\bar{F}(\theta^*, v^*)}(\underbrace{\frac{\partial \bar{F}(\theta^*, v^*)}{\partial \theta^*}}_{<0} + \underbrace{\frac{\partial \bar{F}(\theta^*, v^*)}{\partial v^*}}_{<0}) > 0\end{aligned}$$

where the first equality is obtained by applying the chain rule of differentiation and the second equality follows from Proposition 2.2.

Next, recall that $n_r(\theta^*, v^*) = R(\theta^*, v^*)/\bar{F}(\theta^*, v^*)$ where $R(\theta^*, v^*) = \int_{\theta^*}^{\infty} F(v^* \mid \theta)g(\theta)d\theta$. Thus, n_r is decreasing in c_s if and only if

$$\begin{aligned}0 &< \frac{\partial R}{\partial c_s}\cdot \bar{F} - \frac{\partial \bar{F}}{\partial c_s}\cdot R \\ &= (\frac{\partial R}{\partial \theta^*}\frac{\partial \theta^*}{\partial c_s} + \frac{\partial R}{\partial v^*}\frac{\partial v^*}{\partial c_s})\cdot \bar{F} - R\cdot(\frac{\partial \bar{F}}{\partial \theta^*}\frac{\partial \theta^*}{\partial c_s} + \frac{\partial \bar{F}}{\partial v^*}\frac{\partial v^*}{\partial c_s}) \\ &= -\frac{\partial R}{\partial \theta^*} - \frac{\partial R}{\partial v^*} - \frac{R}{\bar{F}}\cdot(-\frac{\partial \bar{F}}{\partial \theta^*} - \frac{\partial \bar{F}}{\partial v^*})\end{aligned}$$

where the last equality follows from $\frac{\partial \theta^*}{\partial c_s} = \frac{\partial v^*}{\partial c_s} = -\frac{1}{\bar{F}}$. Note that $R(\theta^*, v^*) = (1-G(\theta^*)) - \bar{F}(\theta^*, v^*)$. Therefore,

$$\begin{aligned}-\frac{\partial R}{\partial \theta^*} - \frac{\partial R}{\partial v^*} - \frac{R}{\bar{F}}(-\frac{\partial \bar{F}}{\partial \theta^*} - \frac{\partial \bar{F}}{\partial v^*}) &= g(\theta^*) + \frac{\partial \bar{F}}{\partial \theta^*} + \frac{\partial \bar{F}}{\partial v^*} - \left(\frac{1-G(\theta^*)}{\bar{F}} - 1\right)(-\frac{\partial \bar{F}}{\partial \theta^*} - \frac{\partial \bar{F}}{\partial v^*}) \\ &= g(\theta^*) - \frac{1-G(\theta^*)}{\bar{F}}(-\frac{\partial \bar{F}}{\partial \theta^*} - \frac{\partial \bar{F}}{\partial v^*})\end{aligned}$$

which is strictly positive by Lemma 2.3. □

2.A.8 Proof of Proposition 2.5

Proof. Remember that $n_s(\theta^*, v^*) = \bar{F}(\theta^*, v^*)^{-1}$ where $\bar{F}(\theta^*, v^*) = \int_{\theta^*}^{\infty}(1-F(v^* \mid \theta))g(\theta)d\theta$. Therefore, n_s is decreasing in c_r if and only if $\bar{F}(\theta^*, v^*)$ is increasing in c_r. We have

$$\frac{\partial \bar{F}(\theta^*, v^*)}{\partial c_r} = \frac{\partial \bar{F}(\theta^*, v^*)}{\partial \theta^*}\frac{\partial \theta^*}{\partial c_r} + \frac{\partial \bar{F}(\theta^*, v^*)}{\partial v^*}\frac{\partial v^*}{\partial c_r}.$$

By Proposition 2.3, we have

$$\begin{aligned}\frac{\partial v^*}{\partial c_r} &= -\frac{1}{\mathbb{E}[1 - F(v^* \mid \theta) \mid \theta \geqslant \theta^*]} \\ &= \frac{1 - G(\theta^*)}{\bar{F}(\theta^*, v^*)}\end{aligned}$$

and

$$\begin{aligned}\frac{\partial \theta^*}{\partial c_r} &= \frac{1}{1 - F(v^* \mid \theta^*)} - \frac{1}{\mathbb{E}[1 - F(v^* \mid \theta) \mid \theta \geqslant \theta^*]} \\ &= \frac{g(\theta^*)}{-\frac{\partial \bar{F}(\theta^*, v^*)}{\partial \theta^*}} - \frac{1 - G(\theta^*)}{\bar{F}(\theta^*, v^*)}.\end{aligned}$$

Therefore,

$$\begin{aligned}\frac{\partial \bar{F}(\theta^*, v^*)}{\partial c_r} &= g(\theta^*) - (1 - G(\theta^*)) \left(\frac{-\frac{\partial \bar{F}(\theta^*, v^*)}{\partial \theta^*} - \frac{\partial \bar{F}(\theta^*, v^*)}{\partial v^*}}{\bar{F}(\theta^*, v^*)} \right) \\ &\geqslant g(\theta^*) - (1 - G(\theta^*)) \cdot \frac{g(\theta^*)}{1 - G(\theta^*)} \\ &= 0.\end{aligned}$$

where the inequality follows from Lemma 2.3.

Next, recall that $n_r(\theta^*, v^*) = R(\theta^*, v^*)/\bar{F}(\theta^*, v^*)$ where $R(\theta^*, v^*) = \int_{\theta^*}^{\infty} F(v^* \mid \theta) g(\theta) d\theta$. First, note that

$$\frac{\partial R}{\partial c_r} = \underbrace{\frac{\partial R}{\partial \theta^*}}_{<0} \underbrace{\frac{\partial \theta^*}{\partial c_r}}_{>0} + \underbrace{\frac{\partial R}{\partial v^*}}_{>0} \underbrace{\frac{\partial v^*}{\partial c_r}}_{<0} < 0$$

This implies that n_r is decreasing in c_r since

$$\frac{\partial n_r}{\partial c_r} = \frac{1}{\bar{F}^2} \cdot (\bar{F} \cdot \underbrace{\frac{\partial R}{\partial c_r}}_{<0} - R \cdot \underbrace{\frac{\partial \bar{F}}{\partial c_r}}_{>0}) < 0.$$

□

2.A.9 Proof of Proposition 2.7

Proof. Applying the chain rule of differentiation, we obtain

$$\frac{\partial p^*}{\partial c_s} = \frac{\partial \theta^*}{\partial c_s} \frac{\partial p^*}{\partial \theta^*} + \frac{\partial v^*}{\partial c_s} \frac{\partial p^*}{\partial v^*}.$$

By Proposition 2.2, we have

$$\frac{\partial p^*}{\partial c_s} = -\frac{1}{\bar{F}(\theta^*, v^*)}\left(\frac{\partial p^*}{\partial \theta^*} + \frac{\partial p^*}{\partial v^*}\right).$$

From Proposition 2.6, it is immediate that $\frac{\partial p^*}{\partial \theta^*} + \frac{\partial p^*}{\partial v^*}$ is proportional to

$$-\left(\frac{\partial \bar{F}}{\partial \theta^*}\right)^2 - \left(\frac{\partial \bar{F}}{\partial v^*}\right)^2 - 2\frac{\partial \bar{F}}{\partial \theta^*} \cdot \frac{\partial \bar{F}}{\partial v^*} + \bar{F} \cdot \left[\frac{\partial^2 \bar{F}}{(\partial \theta^*)^2} + \frac{\partial^2 \bar{F}}{(\partial v^*)^2} + 2\frac{\partial^2 \bar{F}}{\partial \theta^* \partial v^*}\right]$$

which is negative by log-concavity of $\bar{F}$.[10] The equilibrium price p^* is therefore increasing in search costs, i.e. $\frac{\partial p^*}{\partial c_s} > 0$. □

2.A.10 Proof of Proposition 2.8

Proof. As aforementioned, we have

$$\frac{\partial p^*}{\partial c_r} = \frac{\partial p^*}{\partial v^*}\frac{\partial v^*}{\partial c_r} + \frac{\partial p^*}{\partial \theta^*}\frac{\partial \theta^*}{\partial c_r}.$$

We want to show (i)

$$\lim_{c_r \to 0} \frac{\partial p^*}{\partial v^*}\frac{\partial v^*}{\partial c_r} > 0$$

and (ii)

$$\lim_{c_r \to 0} \frac{\partial p^*}{\partial \theta^*}\frac{\partial \theta^*}{\partial c_r} = 0$$

which jointly imply $\lim_{c_r \to 0} \frac{\partial p^*}{\partial c_r} > 0$.

By Corollary 2.1, we know that $\lim_{c_r \to 0} \theta^* = -\infty$ and $\lim_{c_r \to 0} v^* = w_0$. Note that this implies $\lim_{c_r \to 0} \mathbb{E}[1 - F(v^* \mid \theta) \mid \theta \geqslant \theta^*] = 1 - F_\theta(w_0)$. Thus,

$$\begin{aligned}
\frac{\partial p^*}{\partial \theta^*}\frac{\partial \theta^*}{\partial c_r} &= \left(\frac{-\frac{\partial \bar{F}}{\partial \theta^*}\left(\frac{\partial \bar{F}}{\partial v^*} + \frac{\partial \bar{F}}{\partial \theta^*}\right) + \left(\frac{\partial^2 \bar{F}}{\partial \theta^{*2}} + \frac{\partial^2 \bar{F}}{\partial v^* \partial \theta^*}\right) \cdot \bar{F}}{\left(\frac{\partial \bar{F}}{\partial v^*} + \frac{\partial \bar{F}}{\partial \theta^*}\right)^2}\right)\left(\frac{g(\theta^*)}{-\frac{\partial \bar{F}}{\partial \theta^*}} - \frac{1 - G(\theta^*)}{\bar{F}}\right) \\
&= \frac{1}{\left(\frac{\partial \bar{F}}{\partial v^*} + \frac{\partial \bar{F}}{\partial \theta^*}\right)^2}\left(A(\theta^*, v^*) + B(\theta^*, v^*)\right)
\end{aligned}$$

---

[10] $\bar{F}$ is log-concave if and only if

$$\alpha^2 \cdot \left(\frac{\partial \bar{F}}{\partial \theta^*}\right)^2 + \beta^2 \cdot \left(\frac{\partial \bar{F}}{\partial v^*}\right)^2 + 2\alpha\beta \cdot \frac{\partial \bar{F}}{\partial \theta^*} \cdot \frac{\partial \bar{F}}{\partial v^*} \geqslant \bar{F} \cdot \left[\alpha^2 \frac{\partial^2 \bar{F}}{(\partial \theta^*)^2} + \beta^2 \frac{\partial^2 \bar{F}}{(\partial v^*)^2} + 2\alpha\beta \frac{\partial^2 \bar{F}}{\partial \theta^* \partial v^*}\right]$$

for all $\alpha, \beta \in \mathbb{R}$. Setting $\alpha = \beta = 1$ yields the desired inequality.

where

$$
\begin{aligned}
A(\theta^*, v^*) &= g(\theta^*)(\frac{\partial \bar{F}}{\partial v^*} + \frac{\partial \bar{F}}{\partial \theta^*}) + \frac{(\frac{\partial^2 \bar{F}}{\partial \theta^{*2}} + \frac{\partial^2 \bar{F}}{\partial v^* \partial \theta^*})}{-\frac{\partial \bar{F}}{\partial \theta^*}} g(\theta^*) \\
&= g(\theta^*)(\frac{\partial \bar{F}}{\partial v^*} + \frac{\partial \bar{F}}{\partial \theta^*}) + g'(\theta^*) \\
&\to 0
\end{aligned}
$$

as $\theta^* \to -\infty$ and $v^* \to w_0$ and

$$
B(\theta^*, v^*) = \underbrace{\frac{\partial \bar{F}}{\partial \theta^*}}_{\to 0} (\frac{\partial \bar{F}}{\partial v^*} + \frac{\partial \bar{F}}{\partial \theta^*}) \frac{1 - G(\theta^*)}{\bar{F}} - \underbrace{(\frac{\partial^2 \bar{F}}{\partial \theta^{*2}} + \frac{\partial^2 \bar{F}}{\partial v^* \partial \theta^*})}_{\to 0} \cdot (1 - G(\theta^*)) \to 0
$$

such that $\lim_{c_r \to 0} \frac{\partial p^*}{\partial \theta^*} \frac{\partial \theta^*}{\partial c_r} = 0$. Next, note that

$$
\begin{aligned}
(\frac{\partial \bar{F}}{\partial v^*} + \frac{\partial \bar{F}}{\partial \theta^*})^2 \cdot \frac{\partial p^*}{\partial v^*} \frac{\partial v^*}{\partial c_r} &= \left(-\frac{\partial \bar{F}}{\partial v^*} \cdot (\frac{\partial \bar{F}}{\partial v^*} + \frac{\partial \bar{F}}{\partial \theta^*}) + (\frac{\partial^2 \bar{F}}{\partial v^{*2}} + \frac{\partial^2 \bar{F}}{\partial v^* \partial \theta^*}) \cdot \bar{F} \right) \frac{-(1 - G(\theta^*))}{\bar{F}} \\
&= -\frac{\partial \bar{F}}{\partial v^*} (-\frac{\partial \bar{F}}{\partial v^*} - \frac{\partial \bar{F}}{\partial \theta^*}) \frac{1 - G(\theta^*)}{\bar{F}} - (\frac{\partial^2 \bar{F}}{\partial v^{*2}} + \frac{\partial^2 \bar{F}}{\partial v^* \partial \theta^*})(1 - G(\theta^*))
\end{aligned}
$$

Note that $\lim_{c_r \to 0} \frac{\partial \bar{F}}{\partial v^*} = \int_{-\infty}^{\infty} f(w_0 \mid \theta) g(\theta) d\theta = f_\theta(w_0)$, $\lim_{c_r \to 0} -\frac{\partial \bar{F}}{\partial \theta^*} = 0$, $\lim_{c_r \to 0}(1 - G(\theta^*)) = 1$, $\lim_{c_r \to 0} \bar{F}(\theta^*, v^*) = 1 - F_\theta(w_0)$. Thus,

$$
\begin{aligned}
\lim_{c_r \to 0} \left(\frac{\partial^2 \bar{F}}{\partial v^{*2}} + \frac{\partial^2 \bar{F}}{\partial v^* \partial \theta^*} \right) &= \lim_{c_r \to 0} \left(-\int_{\theta^*} f'(v^* \mid \theta) g(\theta) d\theta + f(v^* \mid \theta^*) g(\theta^*) \right) \\
&= -\int_{-\infty}^{\infty} f'(w_0 \mid \theta) g(\theta) d\theta \\
&= f'_\theta(w_0).
\end{aligned}
$$

Therefore,

$$
\lim_{c_r \to 0} \frac{\partial p^*}{\partial v^*} \frac{\partial v^*}{\partial c_r} = \frac{1}{1 - F_\theta(w_0)} + \frac{f'_\theta(w_0)}{f_\theta(w_0)^2},
$$

which is strictly positive since log-concavity of $\bar{F}$ implies that the marginal survival function $\bar{F}_\theta = 1 - F_\theta$ is also log-concave.

Thus, we have shown that

$$
\lim_{c_r \to 0} \frac{\partial p^*}{\partial c_r} = \underbrace{\lim_{c_r \to 0} \frac{\partial p^*}{\partial v^*} \frac{\partial v^*}{\partial c_r}}_{>0} + \underbrace{\lim_{c_r \to 0} \frac{\partial p^*}{\partial \theta^*} \frac{\partial \theta^*}{\partial c_r}}_{=0} > 0.
$$

□

References

Mark Yuying An. Logconcavity versus logconvexity: a complete characterization. *Journal of economic theory*, 80(2):350–369, 1998. 2.A.1

Simon P Anderson and Regis Renault. Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics*, pages 719–735, 1999. 2.1, 2.1.1, 2.3.1, 2.5.1

Michael Choi, Anovia Yifan Dai, and Kyungmin Kim. Consumer search and price competition. *Econometrica*, 86(4):1257–1281, 2018. 2.1.1

Pascal Courty and Li Hao. Sequential screening. *The Review of Economic Studies*, 67(4): 697–717, 2000. 2.1.1

Peter Gibbard. A model of search with two stages of information acquisition and additive learning. *Management Science*, 68(2):1212–1217, 2022. 2.1.1, 2.6

John Gittins, Kevin Glazebrook, and Richard Weber. *Multi-armed bandit allocation indices*. John Wiley & Sons, 2011. 2.1.1

Rafael P Greminger. Optimal search and discovery. *Management Science*, 68(5):3904–3924, 2022. 2.1.1, 2.6

Marco A Haan, José L Moraga-González, and Vaiva Petrikaitė. A model of directed consumer search. *International Journal of Industrial Organization*, 61:223–255, 2018. 2.1.1

Toomas Hinnosaar and Keiichi Kawai. Robust pricing with refunds. *The RAND Journal of Economics*, 51(4):1014–1036, 2020. 2.1.1

Roman Inderst and Gilad Tirosh. Refunds and returns in a vertically differentiated industry. *International Journal of Industrial Organization*, 38:44–51, 2015. 2.1.1

Maarten Janssen and Cole Williams. Consumer search and product returns in e-commerce. *American Economic Journal: Microeconomics*, 16(2):387–419, 2024. 2.1.1

Kinshuk Jerath and Qitian Ren. Consumer search and product returns. *Working Paper*, 2023.

Daniel Krähmer and Roland Strausz. Optimal sales contracts with withdrawal rights. *The Review of Economic Studies*, 82(2):762–790, 2015. 2.1.1

Steven A Matthews and Nicola Persico. Information acquisition and refunds for returns. 2007. 2.1.1

Vaiva Petrikaitė. A search model of costly product returns. *International Journal of Industrial Organization*, 58:236–251, 2018. 2.1.1

Qitian Ren and Kinshuk Jerath. Consumer search and product returns. *revise and resubmit at Marketing Science*, 2022. 2.1.1

Jessica Schiffer. The unsustainable cost of free returns. *Vogue Business*, 2019. 2.1

Jian Shen. Ex-ante preference in a consumer search market. *Available at SSRN 2826011*, 2015. 2.1.1

Jeffrey D Shulman, Anne T Coughlan, and R Canan Savaskan. Optimal restocking fees and information provision in an integrated demand-supply model of product returns. *Manufacturing & Service Operations Management*, 11(4):577–594, 2009. 2.1.1

Jonas von Wangenheim. Optimal information design of online marketplaces with return rights. Technical report, University of Bonn and University of Mannheim, Germany, 2024. 2.1.1

Asher Wolinsky. True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics*, 101(3):493–511, 1986. 2.1, 2.1.1, 2.3.1, 2.4.1, 2.5.1

Jidong Zhou. Improved information in search markets. Technical report, Cowles Foundation for Research in Economics, Yale University, 2022. 2.1.1, 2.4.1

Chapter 3

Information Design in Selection Problems

3.1 Introduction

Economic decision problems often consist of selection problems under uncertainty: Sellers select products to advertise with limited information about consumer preferences. Politicians implement one out of many alternatives without perfect information about the identity of the most beneficial alternative. Employers choose which applicant to hire based on noisy information about applicants' productivity. In all of these situations, decision-makers must weigh the available information and make a choice based on what they believe to be the best option.

This paper employs a Bayesian Persuasion model to study how a sender can manipulate a receiver's beliefs about the best alternative. More specifically, I assume that the receiver takes a decision consisting of a selection and an action. The objective of information disclosure by the sender is twofold: The sender wants to persuade the receiver to select a favorable alternative and take a favorable action. For example, a politician may choose one out of many policy proposals to implement (selection) and decide how much money to allocate towards it (action). In this scenario, a think tank may disclose information about the proposals' viability in order to convince the politician to select a particular proposal and to maximize the funding allocated toward it.

The main result decomposes the multi-dimensional Bayesian Persuasion problem into two sub-problems: The *Selection Persuasion* problem, which involves persuading the receiver to select the favorable alternative and the *Action Persuasion* problem, which entails persuading the receiver to take a favorable action conditional on the selection of the favorable alternative.

Selection Persuasion and Action Persuasion may form a trade-off: Suppose that the politician implements the policy proposal he expects to be the most viable and suppose that the amount of money allocated towards the chosen proposal is increasing in the expected viability of the proposal. If a think tank wants to increase the probability that the politician chooses a particular proposal, it can pool together states of the world in which that proposal is the most viable with states in which it is not the most viable. However, the downside is that this pooling decreases the expected viability of the particular proposal conditional on selection of the proposal, which in turn reduces the amount of money that the politician is willing to invest. Therefore, the objective of persuading the politician to select the think tank's preferred proposal may interfere with the think tank's objective of increasing the funding allocated toward the preferred proposal.

Formally, this paper studies a Bayesian Persuasion model with an N-dimensional state of the world $(\theta_1, ..., \theta_N)$ where θ_k is interpreted as the viability of alternative k which affects the sender's and receiver's payoffs only if the receiver selects alternative k. It is well known that solving multi-dimensional Bayesian Persuasion models is challenging. The main problems are first that concavification methods are not helpful due to a large state space, and second, restricting attention to posterior means does not per se significantly simplify the problem since there is no tractable characterization of the convex order when the dimensionality is larger than one.

When it comes to selection problems, the distribution of a receiver's posterior beliefs can only affect the sender's expected payoffs in two ways. First, it can impact the probability that the receiver chooses the sender's preferred option. Second, it can influence the distribution of the receiver's belief regarding the viability of the sender's preferred alternative, given that it has been selected by the receiver. We refer to this distribution as the conditional selection distribution. This observation facilitates a transformation of the multi-dimensional Bayesian persuasion problem into a problem that involves selecting pairs (p, σ) of a selection probability $p \in [0, 1]$ and a conditional selection distribution $\sigma \in \Delta([0, 1])$. It is natural to ask which pairs of selection probability and conditional selection distribution can be induced through some signal structure.

I provide a tractable solution to this question: First, I argue that the optimal selection probability p must be at least as large as the selection probability induced by full information, p_F. The main theorem demonstrates that for any fixed selection probability p in the range of $[p_F, 1]$, there is a unique most-informative conditional selection distribution, denoted as F_p^{CMSS}. In other words, a pair of selection probability and conditional distribution (p, σ) can be induced by some signal structure if and only if σ is a mean-preserving contraction of F_p^{CMSS}. The most informative conditional selection distribution F_p^{CMSS} is induced by a *cutoff-mean signal structure* (CMSS), which reveals the average viability of alternatives $1, ..., N$ if the average viability is below a certain threshold and fully discloses the viability of alternatives $1, ...N$ if the average viability is above the threshold.

The characterization of feasible decomposition pairs disentangles selection persuasion from action persuasion. Selection persuasion refers to the sender's selection of a probability p within the range $[p_F, 1]$, which determines the probability that the receiver selects the sender's preferred alternative. On the other hand, action persuasion concerns the sender's choice of a mean-preserving contraction of F_p^{CMSS}, which determines the receiver's action conditional on the selection of the preferred alternative of the sender. Thus, the action persuasion problem for a fixed selection probability resembles a classical uni-dimensional Bayesian Persuasion problem. Therefore, methods provided by the literature on Bayesian Persuasion can be used to address the action persuasion problem for fixed selection probabilities. Some examples of such methods include those developed by Dworczak and Martini (2019); Gentzkow and Kamenica (2016); Kleiner, Moldovanu, and Strack (2021); Kolotilin (2018); Kolotilin, Mylovanov, and Zapechelnyuk (2022). Once a solution to the action persuasion problem is found, solving the selection persuasion problem becomes a straightforward optimization task.

I analyze two economic applications of the model. Example 3.1 studies the politician-think tank application. The think tank derives a fixed payoff $c > 0$ if the politician selects policy proposal $i^* = 1$ out of N policy proposals, and an additional payoff which is linear in the funding allocated towards policy proposal 1. Using the decomposition approach, I show that a cutoff-mean signal structure solves the think tank's persuasion problem, implying that the politician is either fully informed about the viability of all proposals or fully informed about their average viability. Cutoff-mean signal structures range from only revealing the average viability of the policy proposal to complete information depending on the specific cutoff. I establish that the optimal cutoff is determined by the think tank's relative incentive to maximize the funding allocated to policy proposal 1 versus maximizing the probability of the politician selecting proposal 1.

Example 3.2 considers the case of a salesperson who aims to persuade consumers to buy a particular product, say product 1, out of N products by disclosing information about the products' quality. For the consumers to purchase product 1, they must prefer product 1 over all other products as well as prefer buying product 1 to not buying any product. In this context, selection persuasion determines the probability that consumers prefer product 1 over all other products, and action persuasion determines the probability that the consumers prefer buying product 1 to not buying the product. I use the decomposition approach to derive the optimal information disclosure by the salesperson. Under optimal information disclosure, the consumers strictly prefer product 1 to all other products whenever they purchase product 1 but are indifferent between buying product 1 and not buying any product.

This chapter is organized as follows: In the remainder of this section, we discuss the related literature. In Section 3.2, the model is presented. Section 3.3 characterizes which pairs of selection probability and conditional selection distribution can be induced by some signal structure and thereby disentangles the persuasion problem into selection and action persuasion. Section 3.4 studies optimal persuasion. In Section 3.5, we discuss comparative statics of the model. In Section 3.6, I analyze economic applications, and in Section 3.7, I provide some extensions and limitations of the approach. Section 3.8 discusses the related literature, and Section 3.9 concludes. All omitted proofs can be found in the appendix.

3.2 The Model

The Setup There are N alternatives. Each alternative is either viable ($\tilde{\theta}_i = \bar{\theta}$) or not viable ($\tilde{\theta}_i = \underline{\theta}$). The joint probability function f of $(\tilde{\theta}_1, ..., \tilde{\theta}_N)$ is assumed to be exchangeable.[1] [2] We denote the probability of $\tilde{\theta}_i = \bar{\theta}$ by $q_F \in (0, 1)$.[3] Unless stated otherwise, I normalize $\bar{\theta} = 1$ and $\underline{\theta} = 0$.

We study a Bayesian Persuasion game with two players: A sender (S) transmits information about $(\tilde{\theta}_1, ..., \tilde{\theta}_N)$ to a receiver (R) who observes a signal realization and optimally takes a decision. A decision consists of a selected alternative $k \in \mathcal{N} := \{1, ..., N\}$, and an action $a \in \mathcal{A}$ where $\mathcal{A}$ is an (arbitrary) action set.

Given the decision (k, a) and state of the world $(\theta_1, ..., \theta_N)$, the receiver's payoff satisfies $u^R((k, a), (\theta_1, .., \theta_N)) \equiv u^R(a, \theta_k)$ for some payoff function $u^R : \mathcal{A} \times \{\underline{\theta}, \bar{\theta}\} \to \mathbb{R}$. That is, the selection of an alternative affects the payoffs of the receiver only through the underlying viability of the selected alternative, and the viabilities of other alternatives do not affect payoffs.

The sender's payoff function is assumed to satisfy

$$u^S((k, a), (\theta_1, ..., \theta_N)) = \begin{cases} u^S(a, \theta_k) & \text{if } k = 1 \\ 0 & \text{if } k \neq 1 \end{cases},$$

where $u^S(a, \theta_k) \geqslant 0$ for all $(a, \theta_k) \in \mathcal{A} \times \{\underline{\theta}, \bar{\theta}\}$. Thus, we assume that the sender obtains a positive payoff which depends on the receiver's action and the underlying viability only if the receiver selects a particular alternative which is (without loss of generality) alternative 1. Hence, the sender's primary objective is to persuade the receiver to select alternative 1. The sender's secondary objective is to persuade the receiver to select an action maximizing $u^S(a, \theta_k)$.

The game is as follows: First, the sender chooses a signal structure τ consisting of a (finite) signal space S and a disclosure rule $\tau(\cdot \mid \cdot) : \{\underline{\theta}, \bar{\theta}\}^N \to \Delta(S)$ where $\tau(\mathbf{s} \mid \theta)$ denotes the probability that signal $\mathbf{s} \in S$ is sent if the state of the world is $\theta \in \{\underline{\theta}, \bar{\theta}\}^N$.[4]

[1] To avoid special cases in the proofs we assume that f has full support on $\{\underline{\theta}, \bar{\theta}\}^N$. All results hold without the full support assumption.

[2] f is exchangeable if $f(\theta_1, ..., \theta_N) = f(\theta_{s(1)}, ..., \theta_{s(N)})$ for any $(\theta_1, ..., \theta_N) \in \{0, 1\}^N$ and permutation s on $\{1, ..., N\}$.

[3] Formally, $q_F = \sum_{\theta_{-i} \in \{\underline{\theta}, \bar{\theta}\}^{N-1}} f(\bar{\theta}, \theta_{-i})$ which does not depend on i since f is exchangeable

[4] The restriction to finite signal spaces is without loss of generality since the state space is finite

Then, the state of the world $(\theta_1, .., \theta_N) \in \{\underline{\theta}, \bar{\theta}\}^N$ is realized, and the receiver observes a signal realization $s \in S$ drawn according to $\tau(\cdot \mid \theta)$. Finally, the receiver takes a decision $(k, a) \in \mathcal{N} \times \mathcal{A}$ and payoffs are realized. Both, receiver and sender, are risk-neutral expected utility maximizers. The solution concept is sender-preferred Bayesian Nash Equilibrium.

Receiver's Decision Problem Suppose that the receiver has some posterior belief $(q_1, ..., q_N)$ where q_i denotes the probability the receiver assigns to $\tilde{\theta}_i = 1$.[5] The optimal decision (k^*, a^*) maximizes the receiver's expected payoff given posterior belief $(q_1, ..., q_N)$. That is, the receiver solves

$$\max_{(k,a) \in \mathcal{N} \times \mathcal{A}} q_k u^R(a, \bar{\theta}) + (1 - q_k) u^R(a, \underline{\theta}).$$

Let $U^R(q_k)= \max_{a \in \mathcal{A}} q_k u^R(\bar{\theta}, a) + (1 - q_k) u^R(\underline{\theta}, a)$ denote the continuation payoff of the receiver when he selects an alternative with posterior belief q_k. I assume that U^R is strictly increasing in q_k. This assumption is trivially satisfied when $u^R(a, \bar{\theta}) > u^R(a, \underline{\theta})$, i.e. the receiver strictly prefers viable alternatives over not-viable alternatives for any action.[6] Under this assumption, the receiver selects the alternative that is the most likely to be viable. That is, an optimally selected alternative k^* satisfies

$$k^* \in \arg\max_{i \in \mathcal{N}} q_i. \tag{3.1}$$

The set of optimal actions can then be expressed as

$$a^*(q_1, ..., q_N) \equiv a^*(\max_i q_i) = \arg\max_{a \in \mathcal{A}} \Big((\max_{i \in \mathcal{N}} q_i) u^R(\bar{\theta}, a) + (1 - \max_{i \in \mathcal{N}} q_i) u^R(\underline{\theta}, a) \Big). \tag{3.2}$$

Sender's Persuasion Problem The sender commits to a signal structure τ consisting of a finite signal space S and a disclosure rule $\tau(\cdot \mid \cdot) : \{\underline{\theta}, \bar{\theta}\}^N \to \Delta(S)$. Each signal $s \in S$ induces a posterior belief $\mathbf{q}^s \in [0, 1]^N$ obtained by Bayes' rule, and each signal structure (S, π) induces a distribution of posterior beliefs. It is well known that a distribution $G \in \Delta[0, 1]^N$ can be induced by some signal structure if and only if G is dominated by the prior F in the convex stochastic order, denoted by $G \leqslant_{cx} F$ (Elton and Hill, 1992; Strassen, 1965).[7] Given the restriction to sender-preferred equilibria, and that the

(Kamenica and Gentzkow, 2011).

[5]Formally, a posterior belief is a joint distribution $\eta \in \Delta(\{0, 1\}^N)$. It is straightforward to observe that only the N uni-dimensional marginal distributions $\eta_1, ..., \eta_N$ affect payoffs such that we can identify each posterior belief by the tuple $(q_1, ..., q_N) \in [0, 1]^N$ where $q_i = \eta_i(\bar{\theta})$.

[6]To see that this property is not necessary, consider the payoff function $u^R(a, \theta) = \theta - (a - \theta)^2$ and let $\mathcal{A} = \mathbb{R}$, $\bar{\theta} = 1$, $\underline{\theta} = 0$. It is straightforward to verify that $U^R(q) = q^2$ and that $u^R(a, \bar{\theta}) > u^R(a, \underline{\theta})$ is not satisfied for any $a \leqslant 0$.

[7]It is said that G is dominated by F in the convex order, $G \leqslant_{cx} F$, if and only if $\int c(\mathbf{q}) dG(\mathbf{q}) \leqslant \int c(\mathbf{q}) dF(\mathbf{q})$ for any convex $c : \mathbb{R}^N \to \mathbb{R}$.

receiver's optimal selection and action satisfy (3.1) and (3.2), we can express the sender's expected continuation payoff from inducing posterior belief $\mathbf{q} = (q_1, ..., q_N)$ as

$$U^S(\mathbf{q}) = \begin{cases} U^S(q_1) := \max_{a \in a^*(\mathbf{q})} q_1 u^S(a, \bar{\theta}) + (1 - q_1) u^S(a, \underline{\theta}) & \text{if } q_1 = \max_{i \in \mathcal{N}} q_i \\ 0 & \text{otherwise} \end{cases}. \quad (3.3)$$

The sender's persuasion problem can now be stated as maximizing her expected payoff over all distributions of posterior beliefs that can be induced by some signal structure. Formally,

$$\max_{G: G \leqslant_{cx} F} \int U^S(\mathbf{q}) dG(\mathbf{q}). \quad (3.4)$$

We study two economics applications of the general framework.

Example 3.1 (Lobbying)**.** There is a politician (receiver) who decides (i) which policy proposal $k = 1, ..., N$ to implement and (ii) how much money $a \in \mathbb{R}_+$ to spend on the selected proposal. Each proposal can either be viable ($\theta_i = \bar{\theta}$) or not viable ($\theta_i = \underline{\theta}$). The politician's payoff from (k, a) is assumed to be $a^\gamma/\gamma - a$ if $\tilde{\theta}_k = \bar{\theta}$ and $-a$ if $\tilde{\theta}_k = \underline{\theta}$ where $\gamma \in (0, 1)$. There is a think tank (sender) which has payoff $c + d \cdot a$ for some $c, d > 0$ if the receiver selects alternative 1, and 0 otherwise. The think tank can persuade the politician by designing a signal structure (S, π).

Example 3.2 (Advertising)**.** There is a salesperson (sender) who can disclose information about the quality $\theta_1, ..., \theta_N$ of products $1, ..., N$. There is a consumer (receiver) who decides whether to purchase a product or exercise an outside option which yields a payoff of $t \in (0, 1)$. The consumer's payoff from purchasing product k is $\theta_k \in \{\underline{\theta}, \bar{\theta}\}$. If the consumer purchases product 1, the salesperson receives a positive commission such that the salesperson's objective is to maximize the probability that the consumer buys product 1.

3.3 Decomposing the Problem

3.3.1 Decomposition Approach

The multi-dimensional Bayesian persuasion problem (3.4) for arbitrary continuation payoff $U^S(\cdot)$ is known to be a difficult problem because there is no convenient characterization of the convex order for dimensions two or greater.[8] However, in the specific model of this paper the sender's expected continuation payoff $U^S(\mathbf{q})$ has a simple structure: The continuation payoff is either 0 or only depends on the posterior belief q_1 about the first alternative if the receiver selects alternative 1. Hence, the expected payoff of the sender

[8] If $N = 1$, then $G \leqslant_{cx} F$ if and only if G is a mean-preserving contraction of F defined by $\int_{-\infty}^{x} G(z)dz \leqslant \int_{-\infty}^{x} F(z)dz$ for all x with equality at $x = \max \operatorname{supp} F$.

only depends on the probability that alternative 1 is selected and the distribution of the belief about alternative 1 conditional on selection of alternative 1.

Therefore, we decompose any multi-dimensional distribution G of posterior beliefs into the induced probability p_G that the receiver selects alternative 1 and into the induced distribution of q_1 conditional on selection of alternative 1. Formally, for any $G \in \Delta[0,1]^N$ we define

1. The selection probability $p_G = Pr_G(\tilde{q}_1 = \max_i \tilde{q}_i)$, and
2. The conditional selection distribution $\sigma_G(x) = Pr_G(\tilde{q}_1 \leqslant x \mid \tilde{q}_1 = \max_i \tilde{q}_i)$ for any $x \in \mathbb{R}$.

Proposition 3.1. *The sender's expected continuation payoff from G only depends on (p_G, σ_G). In particular,*

$$\int U^S(\mathbf{q})dG(\mathbf{q}) = p_G \int U^S(x)d\sigma_G(x).$$

Proof. Since it is without loss of generality to restrict attention to finite signal spaces, it is without loss of generality to restrict attention to discrete distributions of posterior beliefs. Thus, G admits a probability function g, and

$$\begin{aligned}
\int U^S(\mathbf{q})dG(\mathbf{q}) &= \sum_{\mathbf{q} \in \text{supp } G} g(\mathbf{q})U^S(\mathbf{q}) \\
&= \sum_{\mathbf{q} \in \text{supp } G,\, q_1 = \max_i q_i} g(\mathbf{q})U^S(q_1) + 0 \\
&= p_G \sum_{\mathbf{q} \in \text{supp } G,\, q_1 = \max_i q_i} \frac{g(\mathbf{q})}{p_G} U^S(q_1) \\
&= p_G \sum_{q_1 \in \text{supp } \sigma_G} Pr_G(\tilde{q}_1 = q_1 \mid \tilde{q}_1 = \max_i \tilde{q}_i) U^S(q_1) \\
&= p_G \int U^S(x) d\sigma_G(x).
\end{aligned}$$

where the second equality follows from (3.3). □

The approach of this paper is to maximize $p_G \int U^S(x)d\sigma_G(x)$ over all feasible decomposition pairs (p_G, σ_G) as opposed to the standard approach of maximizing $\int U^S(\mathbf{q})dG(\mathbf{q})$ over all feasible distributions of posterior beliefs G. We say that a pair (p, σ) of selection probability $p \in [0,1]$ and conditional selection distribution $\sigma \in \Delta[0,1]$ is feasible if $(p, \sigma) = (p_G, \sigma_G)$ for some $G \in \Delta[0,1]^N$ satisfying $G \leqslant_{cx} F$. This definition gives rise to the *decomposed persuasion problem*

$$\max_{(p,\sigma)\, \textit{feasible}} p \int U^S(x)d\sigma(x). \tag{3.5}$$

Proposition 3.1 implies that the decomposed persuasion problem is equivalent to the persuasion problem.[9] In the remaining section, we will show that the decomposed persuasion problem is tractable and related to the standard uni-dimensional Bayesian Persuasion problem over posterior means.

3.3.2 Undominated (p, σ)

As a first step, I argue that it can never be optimal for the sender to choose a signal structure for which the receiver selects an alternative $k \neq 1$ in a state of the world $(\theta_1, ..., \theta_N) \in \{\underline{\theta}, \bar{\theta}\}^N$ satisfying either $\theta_1 = \bar{\theta}$ or $\theta_1 = ... = \theta_N = \underline{\theta}$. Note that in those states of the world, the receiver would select alternative 1 if he knew the state.

We say that a signal structure (S, π) is undominated if for any $s \in S$ such that $q_1^s \neq \max_i q_i^s$, then $\pi(s \mid \theta) = 0$ for all $\theta \in \{\underline{\theta}, \bar{\theta}\}^N$ satisfying either $\theta_1 = \bar{\theta}$ or $\theta_1 = ... = \theta_N = \underline{\theta}$. Thus, a signal structure is undominated if the receiver selects alternative 1 with probability 1 in states of the world for which he would choose to select alternative 1 if he knew the state. Consequently, we say that (p, σ) is undominated if (p, σ) is induced by some undominated signal structure. The next Lemma argues that we can restrict attention to undominated (p, σ) when searching for a solution to the decomposed persuasion problem (3.5).

Lemma 3.1. *There exists an undominated (p, σ) that solves the decomposed persuasion problem* (3.5). *If $U^S(x) > 0$ for any $x \in [0, 1]$, any solution (p, σ) to the decomposed persuasion problem is undominated.*

The proof works as follows: Suppose that a signal structure (S, π) is dominated. We can improve upon this signal structure by fully disclosing a state of the world θ satisfying $\theta_1 = \bar{\theta}$ or $\theta_1 = ... = \theta_N = \underline{\theta}$ whenever the signal structure (S, π) sends a signal s for which the receiver does not select alternative 1. This construction increases the selection probability without changing the continuation payoff from other signals. Thus, the sender weakly prefers the newly constructed signal structure and strictly prefers this signal structure if the continuation payoff when the receiver selects alternative 1 is positive.

Undominated (p, σ) have two important properties: First, the selection probability p is bounded below by the selection probability of full information $p_F = f(\underline{\theta}, ..., \underline{\theta}) + \sum_{\theta_{-i}} f(\bar{\theta}, \theta_{-i})$ since the receiver always selects alternative 1 whenever he would select alternative 1 if he had full information. Second, the expected value of any conditional selection distribution σ only depends on the selection probability p.

Lemma 3.2. *Any undominated (p, σ) satisfies $p \geqslant p_F$ and $\int x d\sigma(x) = \frac{q_F}{p}$.*

[9] Equivalence refers to (i) G^* solves (3.4) if and only if (p_{G^*}, σ_{G^*}) solves (3.5), and (ii) the value of maximization problem (3.4) is equal to the value of maximization problem (3.5).

3.3.3 Characterization of Feasibility

In order to solve the decomposed persuasion problem (3.5), we have to characterize which decomposition pairs (p, σ) are feasible.

First, consider any fixed selection probability p. It is always possible to decrease the informativeness of the conditional selection distribution in the sense of mean-preserving contraction by pooling signal realizations for which the receiver selects alternative 1, and leave signal realizations for which the receiver selects alternative $k \neq 1$ unchanged. Lemma 3.3 formalizes this intuition.

Lemma 3.3. *Suppose that σ' is a mean-preserving contraction of σ. If (p, σ) is feasible, then (p, σ') is feasible.*

While it is always possible to decrease informativeness of the conditional selection distribution without changing the selection probability p, increasing the informativeness of the conditional selection distribution is bounded. Intuitively, for any fixed selection probability p, the sender cannot reveal information that allows the receiver to select a viable alternative with higher precision.

In particular, suppose that the receiver always selects alternative 1, i.e. $p = 1$. Since the prior distribution F is symmetric, the receiver has to be indifferent between selecting alternative $1, ..., N$ with probability 1. That is, a minimal requirement on any distribution of posterior beliefs G which induces a selection probability equal to 1 is that $q_1 = ... = q_N$ for all $\mathbf{q} \in \operatorname{supp} G$. The most informative distribution of posterior belief satisfying this condition is the distribution of posterior beliefs induced by disclosing the average state of the world $\sum \theta_i/N$ since any information that is not less informative than revealing the average state can be used to distinguish alternative 1 from other alternatives, and thereby decreases the selection probability to less than 1. As a consequence, Lemma 3.3 implies that if $p = 1$, then (p, σ) is feasible if and only if σ is a mean-preserving contraction of the CDF of the average viability F^{Mean}.[10]

The main theorem extends this logic to any fixed selection probability $p \in [q_F, 1]$. We define a cutoff-mean signal structure (CMSS) by a cutoff $\theta^* \in [0, 1]$ on the average viability of a state such that the signal structure discloses the average state $\sum \theta_i/N$ if $\sum \theta_i/N$ is below some cutoff $\bar{\theta}$, and discloses the complete state of the world $(\theta_1, ..., \theta_N)$ if $\sum \theta_i/N$ is above the cutoff θ^*. If the average state $\sum \theta_i/N$ is exactly equal to the cutoff θ^*, a cutoff-mean signal structure randomizes between disclosing the average state of the world and the complete state of the world.

Symmetry of the receiver's prior belief implies that the receiver is indifferent between selecting alternatives $1, ..., N$, whenever the average state of the world is revealed. Since

[10] The probability function f^{Mean} is defined by $f^{Mean}(k/N) = \sum_{\theta : \sum \theta_i = k} f(\theta_1, ..., \theta_N)$ for $k = 0, 1, ..., N$. Since f is exchangeable, we get $f^{Mean}(k/N) = \binom{N}{k} f(\theta_1, ..., \theta_N)$ for any $(\theta_1, ...\theta_N)$ such that $\sum \theta_i = k$.

we focus on sender-preferred equilibria, the receiver selects alternative 1 whenever he is indifferent. Thus, the larger the cutoff θ^* of a cutoff-mean signal structure is, the larger the probability is that the receiver selects alternative 1. For every $p \in [0, 1]$, there exists a unique cutoff-mean signal structure such that the receiver selects alternative 1 with probability p. I refer to the cutoff mean signal structure that induces a selection probability of p as (the) p-cutoff-mean signal structure (p-CMSS).

The p-CMSS induces a conditional selection distribution F_p^{CMSS} defined by

$$F_p^{CMSS}(x) = \begin{cases} \frac{F^{Mean}(x)}{p} & \text{if } x < (F^{Mean})^{-1}(y_p) \\ \frac{y_p}{p} & \text{if } (F^{Mean})^{-1}(y_p) \leqslant x < 1 \\ 1 & \text{if } x \geqslant 1 \end{cases}$$

where $y_p \in [F^{Mean}(0), p]$ solves $\int x dF_p^{CMSS}(x) = \frac{q_F}{p}$.[11] Lemma 3.4 in the appendix confirms that F_p^{CMSS} is well-defined, i.e. y_p exists and is unique, and F_p^{CMSS} is a CDF for any $p \in [p_F, 1]$.

Theorem 3.1. *An undominated (p, σ) is feasible if and only if σ is a mean-preserving contraction of F_p^{CMSS}.*

Two key observations indicate why the conditional selection distribution F_p^{CMSS} induced by the p-CMSS is more informative than any other feasible conditional selection distribution. First, whenever the receiver selects alternative 1, he is either indifferent between selecting alternative 1 and selecting any other alternative or is fully convinced that alternative 1 is viable. Second, compared to any other signal structure satisfying the first property, the receiver is the least convinced of the viability of alternative 1 whenever the receiver is indifferent between selecting alternative 1 and selecting any other alternative. Thus, among all signal structures for which the receiver selects alternative 1 with probability p, a cutoff-mean signal structure spreads out the posterior belief about alternative 1 as much as possible, which results in the most informative feasible conditional selection distribution.

3.3.4 Selection and Action Persuasion

Theorem 3.1 and Lemma 3.1 disentangle the sender's persuasion problem into a *selection persuasion* and *action persuasion* problem.

Corollary 3.1. *The persuasion problem* (3.4) *is equal to*

$$\max_{p \in [p_F, 1]} p \cdot \left[\max_{\sigma \in MPC(F_p^{CMSS})} \int U^S(x) d\sigma(x) \right]. \tag{3.6}$$

[11] $(F^{Mean})^{-1}$ denotes the generalized inverse of F^{Mean}. That is, $(F^{Mean})^{-1}(y) = \sup\{x : F^{Mean}(x) \leqslant y\}$.

We can interpret the sender's persuasion problem as follows: First, the sender chooses the probability that the receiver will select alternative 1. Then, in the second stage, the sender decides how much information to disclose about alternative 1 while keeping the selection probability p fixed. The second stage resembles a standard uni-dimensional Bayesian Persuasion problem in which the receiver's action only depends on the posterior mean and the prior distribution is given by F_p^{CMSS}. We refer to the inner-maximization problem *(p-)action persuasion*, and given the value function

$$V(p) := \max_{\sigma \in \mathrm{MPC}(F_p^{CMSS})} \int U^S(x) d\sigma(x),$$

we refer to the outer-maximization problem $\max_{p \in [p_F, 1]} pV(p)$ as *selection persuasion.* The focus of this article is on the selection persuasion problem since there are a number of articles characterizing solutions to the action persuasion problem, e.g. Dworczak and Martini (2019); Gentzkow and Kamenica (2016); Kleiner, Moldovanu, and Strack (2021); Kolotilin (2018). In the next subsection, I derive some properties of optimal persuasion which will be useful in applications.

3.4 Optimal Persuasion and Trade-off

Providing a solution to the selection persuasion problem is difficult since the selection persuasion problem depends on the value of the action persuasion problem which is not known in the general case. In this section, we provide solutions for special classes of the sender's expected continuation payoffs.

Non-Increasing Sender Payoffs First, suppose that the continuation payoff U^S is non-increasing, which means that the sender wants to persuade the receiver to select alternative 1 and conditionally on selection of alternative 1, she prefers that the receiver believes that alternative 1 is not viable. In this case, there is no trade-off: Relative to full information, the sender can increase the selection probability, and reduce the receiver's belief about the viability of alternative 1 at the same time by pooling states in which alternative 1 is viable with states for which alternative 1 is not viable, for instance by disclosing the average viability of alternatives.

Proposition 3.2. *Suppose that U^S is non-increasing. Then, there exists a solution (p, σ) to* (3.6) *for which $p = 1$ (any solution (p, σ) satisfies $p = 1$ if U^S is strictly decreasing).*

Proof. We show that the value of the p-action persuasion problem, $V(p)$ is non-decreasing in p which implies that $pV(p)$ is non-decreasing in p, and hence, $p = 1$ solves $\max_{p \in [p_F, 1]} pV(p)$.

Let $p', p \in [p_F, 1]$, $p' \geqslant p$. From the definition of F_p^{CMSS}, it is straightforward to see that $F_p^{CMSS} \geqslant_{fosd} F_{p'}^{CMSS}$. Consider an arbitrary $\sigma \in \text{MPC}(F_p^{CMSS})$ and note that $\sigma \in \text{MPC}(F_p^{CMSS})$ and $F_p^{CMSS} \geqslant_{fosd} F_{p'}^{CMSS}$ implies $F_{p'}^{CMSS} \leqslant_{icv} \sigma$. where $\leqslant_{icv}$ denotes the increasing concave order. Then, by Shaked and Shanthikumar (2007) Theorem 4.A.6.(c), there exists σ' such that $\sigma' \leqslant_{fosd} \sigma$ and $\sigma' \in \text{MPC}(F_p^{CMSS})$. Therefore, if U^S is non-increasing, then $\int U^S(x)d\sigma(x) \leqslant \int U^S(x)d\sigma'(x)$ such that $V(p) \leqslant V(p')$ which implies that $pV(p)$ is non-decreasing in p. □

If U^S is non-decreasing, the sender faces a trade-off: She wants to persuade the receiver to select alternative 1 but also prefers that the receiver is convinced of the viability of alternative 1 conditional on selection of alternative 1. Any increase of the selection probability above p_F requires pooling states in which alternative 1 is viable with states for which alternative 1 is not viable which in turn decreases the expected viability of alternative 1 conditional on selection.

Sufficiently Concave Sender Payoffs We identify a large class of sender preferences such that the trade-off is always resolved in favor of increasing the selection probability, and in particular in favor of not revealing any information. Let $\bar{co}(U^S)(\cdot)$ denote the concavification of $U^S(\cdot)$.[12]

Proposition 3.3. *If $U^S(q_F) = \bar{co}(U^S)(q_F)$, then $(1, \delta_{q_F})$ (which corresponds to an uninformative signal structure) is optimal.*

Proof. By 3.1, the optimal (p, σ) solves $\max_{p \in [p_F,1]} p \cdot V(p)$ where $V(p) = \max_{\sigma \in \text{MPC}(F_p^{CMSS})} \int U^S(x)d\sigma(x)$. Any $\sigma \in \text{MPC}(F_p^{CMSS})$ satisfies $\int x d\sigma(x) = q_F/p$, and thus $V(p) \leqslant \bar{co}(U^S)(\frac{q_F}{p})$ which implies that $\max_{p \in [p_F,1]} p\bar{co}(U^S)(\frac{q_F}{p})$ is an upper bound on the sender's persuasion value. Note that

$$\begin{aligned}
\max_{p \in [p_F,1]} p\bar{co}(U^S)(\frac{q_F}{p}) &\leqslant \max_{p \in [p_F,1]} p\bar{co}(U^S)(\frac{q_F}{p}) + (1-p)\bar{co}(U^S)(0) \\
&\leqslant \max_{p \in [p_F,1]} \bar{co}(U^S)(p\frac{q_F}{p} + (1-p) \cdot 0) \\
&= \bar{co}(U^S)(q_F) = U^S(q_F)
\end{aligned}$$

where the first inequality is satisfied since $\bar{co}(U^S)(0) \geqslant U^S(0) \geqslant 0$ and the second inequality is satisfied since $\bar{co}(U^S)$ is concave by definition. Thus, the sender's continuation payoff from inducing any (p, σ) is bounded above by $U^S(q_F)$ which is the no-information value. □

Convex Sender Payoffs When U^S is convex, the p-action persuasion problem has the trivial solution F_p^{CMSS}. Thus, some cutoff signal structure must be optimal for convex

[12]That is, $\bar{co}(U^S)$ is the smallest concave function that is everywhere larger than U^S.

continuation payoffs. The next proposition characterizes the exact cutoff x^*, and the corresponding selection probability p^*.

Consider the set $\bar{\Theta}(k)$ of states of the world with average state $x = k/N$ for some $k = 1, ..., N$. This set contains all states of the world for which exactly k alternatives are viable. Since the prior is exchangeable, all states of the world in $\bar{\Theta}(k)$ are equally likely. If the sender pools all of these states together, she receives a continuation payoff $U^S(x)$ since the receiver has posterior belief $(x, ..., x)$ and is indifferent among all alternatives, so selects alternative 1 with probability 1. On the other hand, if the sender separates all of these states, the receiver selects alternative 1 if and only if a state realizes in which alternative 1 is viable which occurs with probability x. Therefore, she receives an expected continuation payoff of $xU^S(1)$.

Let $\bar{x} = \min\{x : U^S(x) \leqslant xU^S(1)\}$ denote the threshold average viability for which the sender is indifferent between revealing the average viability and disclosing the state of the world. Since U^S is convex and $U^S(0) \geqslant 0$, the sender prefers fully disclosing the viability of alternatives to disclosing the average viability if and only if the average viability exceeds $\bar{x}$. Finally, since an alternative is either viable ($\theta = \bar{\theta} \equiv 1$) or not viable ($\theta = \underline{\theta} \equiv 0$), the average viability has positive support only on $\{0, 1/N, ..., 1\}$ such that the exact cutoff x^* cannot be equal to $\bar{x}$. Therefore, we define $x^* = \frac{\lfloor N\bar{x} \rfloor}{N}$ as the largest average viability for which the sender prefers disclosing the average viability to disclosing the state of the world fully. The next Proposition, illustrated in Figure 3.1, shows that it is optimal to pool all states with the same average viability if the average viability is below x^* and disclose states with an average viability above x^*.

Proposition 3.4. *Suppose that U^S is convex. Then, $(p^*, F_{p^*}^{CMSS})$ is optimal where $p^* = F^{Mean}(x^*) + \sum_{x>x^*} f^{Mean}(x)x$.*

As an immediate consequence, we obtain a sufficient condition for the optimality of full information.

Corollary 3.2. *Suppose that U^S is convex and $U^S(0) = 0$. Then, $(p_F, F_{p_F}^{CMSS})$ (which corresponds to full information) is optimal.*

We close the section by arguing that under mild conditions, full information can never be optimal provided that the number of alternatives is sufficiently large. If the sender obtains a fixed payoff guarantee $c > 0$, when the receiver selects alternative 1 irrespective of the selected action, then full information is never optimal with sufficiently many alternatives.

Proposition 3.5. *Suppose there exists some $c > 0$ such that $U^S(x) > c$ for all x, and suppose that $N \geqslant \frac{U^S(1)}{c}$. Then, full information is not optimal.*

Proof. I show that the sender can improve upon full information. Consider the signal structure that pools all states $(\theta_1, ..., \theta_N)$ such that $\sum \theta_i/N = 1/N$, and fully discloses

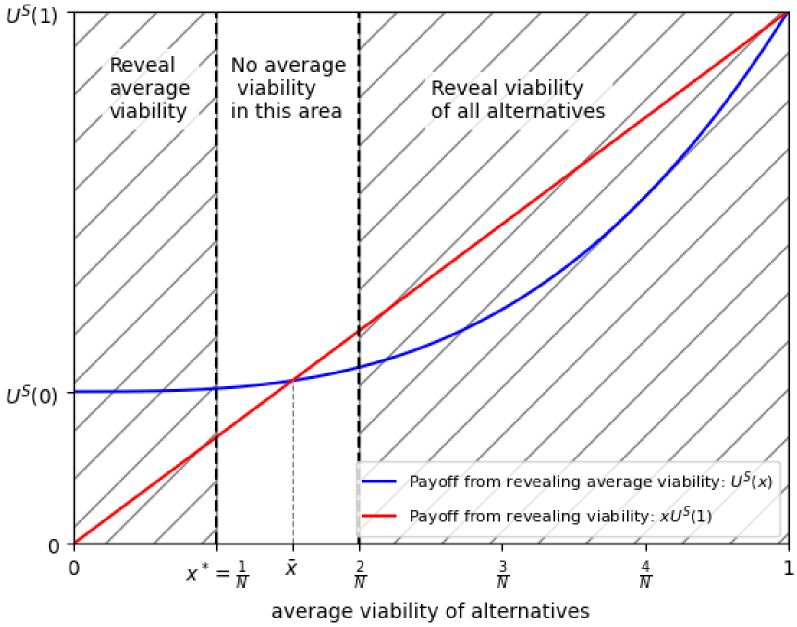

Figure 3.1: Proposition 3.4 illustrated for $N = 5$ alternatives

any other state of the world. This signal structure differs from full information if and only if nature draws a state $(\theta_1, ..., \theta_N)$ with average state equal to $1/N$. Among all of these N states, the receiver selects alternative 1 under full information only in the state $(1, 0, ...0)$ while under this signal structure the receiver always selects alternative 1, and beliefs alternative 1 is viable with probability $1/N$. Therefore, this signal structure yields a larger sender's expected payoff if and only if $U^S(\frac{1}{N}) > \frac{1}{N}U^S(1)$, which is satisfied under the assumptions of the Proposition. □

3.5 Comparative Statics

How does the value of persuasion V_N^* change in the number N of alternatives? The main comparative statics result establishes that the sender's value of persuasion is non-increasing in the number of alternatives. A naive intuition is that it is more difficult for the sender to persuade the receiver to select alternative 1 when there are more alternatives. However, note that this intuition is misleading because the sender can always induce a selection probability of 1 by not revealing any information, or by revealing the average state of the world.

The correct intuition is that given some selection probability p, it becomes more difficult to persuade the receiver to select favorable actions when the number of alternatives increases. To convey the idea behind this intuition, suppose the following: The optimal selection probability is $p = 1$ for any number of alternatives (for example because $U^S(\cdot)$ is non-increasing by Proposition 3.2) and the viabilities $(\tilde{\theta}_1, ..., \tilde{\theta}_N)$ are identically and independently distributed. Theorem 3.1 implies that for $p = 1$ any conditional selection distribution $\sigma \in \text{MPC}(F_N^{Mean})$ is feasible where F_N^{Mean} denotes the CDF of $\sum \tilde{\theta}_i/N$. The weak law of large numbers implies that F_N^{Mean} converges to the degenerate distribution δ_{q_F}. Furthermore, it is known (Müller and Stoyan (2002), Corollary 1.5.24) that F_{N+1}^{Mean} is a MPC of F_{N+1}^{Mean}, i.e. the convergence of F_N^{Mean} to δ_{q_F} is monotone in the convex stochastic order which implies that the set of feasible conditional distributions collapses as N increases, and the value of the persuasion problem decreases. The proof of Theorem 3.6 extends this argument to all exchangeable random variables and any selection probabilities.

Suppose that the N states of the world are distributed according to some exchangeable probability function f_N for all $N \in \mathbb{N}$. We assume that increasing the number of alternatives from N to $N+1$ does not alter the joint distribution of the first N alternatives.[13] That is,

$$f_{N+1}(\theta_1, ..., \theta_N, \underline{\theta}) + f_{N+1}(\theta_1, ..., \theta_N, \bar{\theta}) = f_N(\theta_1, ..., \theta_N) \text{ for all } N \in \mathbb{N}.$$

Let V_N^* denote the value of persuasion when there are $N \in \mathbb{N}$ alternatives. That is, $V_N^* = \max_{p \in [p_{F;N}, 1]} p \cdot \max_{\sigma \in \text{MPC}(F_{p;N}^{CMSS})} \int U^S(x) d\sigma(x)$ where $p_{F;N}$ denotes the probability that alternative 1 is selected from $N \in \mathbb{N}$ alternatives under full information and $F_{p;N}^{CMSS}$ denotes the p-cutoff-mean conditional distribution when there are N alternatives.

Proposition 3.6. *The value of persuasion V_N^* is non-increasing in the number of alternatives N.*

3.6 Applications

3.6.1 Lobbying

This section applies the decomposition approach to analyze Example 3.1. Given some posterior belief $(q_1, ..., q_N) \in [0,1]^N$, the politician solves the decision problem $\max_{(k,a)} q_k a^\gamma/\gamma - a$. It is straightforward to verify that the optimal decision (k^*, a^*) satisfies $k^* \in \arg\max_i q_i$ and $a^*(q_1, ..., q_N) \equiv a^*(\max q_i) = (\max q_i)^{1/(1-\gamma)}$. Therefore, the think tank's continuation payoff from inducing $(q_1, ..., q_N)$ is

$$U^S(q_1, ..., q_N) = \begin{cases} c + d \cdot (q_1)^{1/(1-\gamma)} & \text{if } q_1 = \max_i q_i \\ 0 & \text{if } q_1 \neq \max_i q_i \end{cases}.$$

[13] This is a minimal assumption to compare the persuasion problem with $N+1$ alternatives to the persuasion problem with N alternatives. Intuitively, this implies that the N dimensional problem is equivalent to the $N+1$ dimensional problem if the sender and receiver were forced to ignore the $(N+1)$th alternative.

Corollary 3.1 implies that the think tank's decomposed persuasion problem is

$$\max_{p\in[p_F,1]} p \cdot \left[\max_{\sigma\in\mathrm{MPC}(F_p^{CMSS})} \int \left(c + dx^{1/(1-\gamma)}\right) d\sigma(x)\right].$$

Since $x \mapsto c + dx^{1/(1-\gamma)}$ is convex for any $\gamma \in (0,1)$, we can apply Proposition 3.4 to solve the think tank's problem. Let $\bar{x} = \min\{x \in [0,1] : x(c+d) \geqslant c + dx^{1/(1-\gamma)}\}$ and $x^* = \frac{\lfloor N\bar{x}\rfloor}{N}$.

Proposition 3.7. *$(p^*, F_{p^*}^{CMSS})$ is optimal where $p^* = F^{Mean}(x^*) + \sum_{x>x^*} f^{Mean}(x)x$, and x^* is non-decreasing in $\frac{c}{d}$. In particular,*

(i) Suppose that $\frac{c}{d} > \gamma/(1-\gamma)$. Then, $(1, F^{Mean})$ (revealing average state) is optimal, and

(ii) Suppose that $\frac{c}{d} < \frac{1-(\frac{1}{N})^{\gamma/(1-\gamma)}}{N-1}$. Then, $(p_F, F_{p_F}^{CMSS})$ (full information) is optimal.

There are a couple of noteworthy points: First, in any state of the world, it is optimal to either fully disclose the state, or pool the state together with all other states of the world with the same average state of the world. Further, the optimal signal structure discloses the state of the world if and only if the average state of the world is above a cutoff. Second, the optimal cutoff x^* is determined by the number of alternatives N, the relative incentive $\frac{c}{d}$ of maximizing money spent versus selection of alternative 1, and the politician's risk aversion parameter γ.

If the think tank's priority is to persuade the politician to select alternative 1 (that is, $\frac{c}{d}$ is sufficiently large), then it is optimal to reveal the average viability of an alternative. In this case, the politician is indifferent among all alternatives, and selects alternative 1 with probability 1. If the think tank's priority is to maximize the expected money spent by the politician on alternative 1 (that is, $\frac{c}{d}$ is sufficiently small), then full information is optimal. Full information maximizes the expected money spent on alternative 1 among all signal structures, since the politician is risk averse and as a result, the optimal amount $a^*(\cdot)$ of money allocated towards the selected proposal is convex in the politician's belief that the selected alternative is viable.

Third, when the number of alternatives is sufficiently large, full information is never optimal.[14] As the number of alternatives gets large, the loss of pooling all states of the world in which exactly one alternative is viable instead of disclosing these states, is given by $\frac{N-1}{N}c > 0$ while the gain is given by $\frac{1}{N}da^*(1)$ which converges to 0.

3.6.2 Advertising

In this section, we will analyze Example 3.2. First, note that a consumer with posterior belief $(q_1, ..., q_N)$ purchases product 1 if and only if $q_1 = \max_i q_i$ and $q_1 \geqslant t$. Therefore,

[14] This is also implied by Proposition 3.5 since $\lim_{N\to\infty}(\frac{1-(\frac{1}{N})^{\gamma/(1-\gamma)}}{N-1}) = 0$.

the salesperson's payoff from inducing a posterior belief $(q_1, ..., q_N)$ is

$$U^S(q_1, ..., q_N) = \begin{cases} \mathbf{1}_{q_1 \geqslant t} & \text{if } q_1 = \max_i q_i \\ 0 & \text{if } q_1 \neq \max_i q_i \end{cases}.$$

Corollary 3.1 implies that the salesperson's persuasion problem is

$$\max_{p \in [p_F, 1]} p \cdot \left[\max_{\sigma \in \text{MPC}(F_p^{CMSS})} \int_t^1 d\sigma(x) \right].$$

If $t \leqslant q_F$, the consumer will buy product 1 with a probability of 1, even if he receives no additional information beyond the prior information. Therefore, to ensure that the problem is non-trivial, we assume that $q_F < t$. It is straightforward to solve the salesperson's decomposed persuasion problem:

Proposition 3.8. *The salesperson's persuasion value (the maximal probability with which consumer buys product* 1*) is* $\frac{q_F}{t}$*. In particular,*

(i) Suppose that $t \leqslant \frac{q_F}{p_F}$*. Then,* $(\frac{q_F}{t}, \delta_t)$ *is optimal.*

(ii) Suppose that $t \geqslant \frac{q_F}{p_F}$*. Then,* $(p_F, (1-\alpha)\delta_0 + \alpha\delta_t)$ *is optimal where* $\alpha = \frac{1}{t}\frac{q_F}{p_F}$*.*

Under the non-triviality assumption, the threshold t is larger than the prior belief that a product is viable such that the salesperson has to persuade the consumer. In case (i), the optimal decomposition pair $(\frac{q_F}{t}, \delta_t)$ can be interpreted as follows: First, the salesperson reveals all states of the world perfectly if the average state of the world is above some cutoff and discloses the mean state of the world if the average state of the world is below some cutoff where the cutoff is set such that the consumer selects alternative 1 with probability q_F/t. In the next step, the salesperson pools all signal realizations together for which the receiver prefers product 1 over all other products.

As an example, consider $N = 3$ i.i.d. products where product $i = 1, 2, 3$ is viable with probability 1/2, and suppose that the consumer's threshold is $t = 4/7$. Then, the optimal signal structure can be obtained by first considering the cutoff-mean signal structure for which the consumer prefers product 1 over products 2 and 3 with probability $q_F/t = 7/8$. This cutoff-mean signal structure is the signal structure that reveals the state of the world if $\sum \theta_i/3 \geqslant 2/3$ (either 2 or all 3 products are viable) and reveals the average state of the world if $\sum \theta_i/3 \leqslant 1/3$ (one or no products are viable). For this signal structure, the consumer only prefers product 2 or 3 over product 1 if the state of the world is $(\underline{\theta}, \bar{\theta}, \bar{\theta})$ which occurs with probability 1/8. In the next step, the salesperson pools all signal realizations together on a signal realization except for the signal realization revealing $(\underline{\theta}, \bar{\theta}, \bar{\theta})$. It is straightforward to calculate that the receiver's posterior belief conditionally on receiving the pooled signal realization is $t = 4/7$ such that the consumer buys product 1 with probability $\frac{q_F}{t} * 1 = 7/8$.

If the threshold t is sufficiently large (Case (ii)), the procedure of case (i) is no longer optimal: If t is large, pooling all states for which the consumer would purchase product 1 on a single signal might not be sufficient to ensure that the consumer prefers to purchase product 1 over selecting the outside option. The optimal signal structure first separates all states in which the consumer purchases product 1 from states in which the consumer purchases products 2 or 3, and then either separates $(\underline{\theta}, ..., \underline{\theta})$ from states of the world in which product 1 is viable or pools states for which product 1 is viable with $(\underline{\theta}, ..., \underline{\theta})$ in such a way that the receiver is exactly indifferent between purchasing product 1 and selecting his outside option.

3.7 Extensions and Limitations

3.7.1 Asymmetric Priors

Up until now, it has been assumed that the prior belief about the states of the world $(\theta_1, ..., \theta_N)$ is exchangeable. If there are $N = 2$ alternatives, this assumption is satisfied if and only if $(\bar{\theta}, \underline{\theta})$ is equally likely as $(\underline{\theta}, \bar{\theta})$.

I illustrate that the approach of this paper can be adjusted to account for asymmetric priors, at least for two alternatives. Suppose that a priori it is less likely that the sender's preferred alternative (alternative 1) is viable, i.e. $f(\bar{\theta}, \underline{\theta}) < f(\underline{\theta}, \bar{\theta})$.[15]

Further, suppose first that the sender wants to maximize the probability that the receiver selects alternative 1. Disclosing the mean state of the world $\frac{1}{2}(\theta_1 + \theta_2)$ or disclosing no information induces the receiver to be indifferent between choosing alternative 1 and alternative 2 with an exchangeable prior. However, the receiver selects alternative 2 when alternative 2 is ex-ante more likely to be viable, and the receiver only knows that the average viability. As a consequence, achieving a selection probability of $p = 1$ is not feasible.

Instead, the sender can maximize the selection probability by fully disclosing $(\bar{\theta}, \underline{\theta})$ with probability $1 - f(\underline{\theta}, \bar{\theta})/f(\bar{\theta}, \underline{\theta})$, and otherwise disclosing the mean state of the world. This results in a selection probability of $p_{max} = f(\bar{\theta}, \bar{\theta}) + 2f(\underline{\theta}, \bar{\theta}) + f(\underline{\theta}, \underline{\theta})$, while revealing as much information about θ_1 as possible. Conversely, with full information the receiver selects alternative 1 whenever $(\bar{\theta}, \bar{\theta})$, $(\underline{\theta}, \underline{\theta})$ or $(\bar{\theta}, \underline{\theta})$ is realized. Therefore, full information reveals as much information as possible about θ_1 for a selection probability of $p_F = f(\bar{\theta}, \bar{\theta}) + f(\bar{\theta}, \underline{\theta}) + f(\underline{\theta}, \underline{\theta})$.

For any $p \in [p_F, p_{max}]$, define

$$\alpha(p) := \frac{p - f(\bar{\theta}, \bar{\theta}) - f(\underline{\theta}, \underline{\theta})}{2f(\bar{\theta}, \underline{\theta})}.$$

[15] The diametric case in which the sender's preferred alternative is more likely to be viable can be analyzed similarly.

The adjusted $p \in [p_F, p_{max}]$ cutoff mean signal structure is defined as follows: If $(\theta_1, \theta_2) = (\bar{\theta}, \bar{\theta})$ or $(\theta_1, \theta_2) = (\underline{\theta}, \underline{\theta})$ reveal the states of the world perfectly. If $(\theta_1, \theta_2) = (\underline{\theta}, \bar{\theta})$ disclose (θ_1, θ_2) with probability $1 - \alpha(p) \cdot f(\bar{\theta}, \underline{\theta})/f(\underline{\theta}, \bar{\theta})$ and with probability $\alpha(p) \cdot f(\bar{\theta}, \underline{\theta})/f(\underline{\theta}, \bar{\theta})$ send a signal that reveals the average state of the world $s = \frac{1}{2}(\bar{\theta} + \underline{\theta})$. Finally, if $(\theta_1, \theta_2) = (\bar{\theta}, \underline{\theta})$ disclose the states of the world perfectly with probability $1 - \alpha(p)$ and reveal the average state of the world by sending signal $s = \frac{1}{2}(\bar{\theta} + \underline{\theta})$ with probability $\alpha(p)$.

It is straightforward to verify that (i) the receiver remains indifferent upon receiving a signal that reveals the average state of the world, (ii) an adjusted p-CMSS indeed induces a selection probability of p. Additionally, it can be shown that analogously to Theorem 3.1, the adjusted p-CMSS reveals as much information as possible about alternative 1 in the sense that the conditional selection distribution of any other signal structure inducing a selection probability p must be a mean-preserving contraction of the conditional distribution induced by the adjusted p-CMSS.

This result clarifies that two properties are important for the role of the CMSS as most informative signal structures. First, under any CMSS, whenever the receiver selects alternative 1, he is indifferent between selecting alternative 1 and any other alternative. Second, any CMSS achieves the first property without pooling any two states of the world $(\theta_i)_{i=1,..,N}$ and $(\theta'_i)_{i=1,...,N}$ for which the receiver would select alternative 1 if he knew the state since pooling these states degrades the information about alternative 1 without changing the selection probability. If the viability of any alternative is binary ($\theta_i \in \{\underline{\theta}, \bar{\theta}\}$), the receiver selects alternative 1 whenever he expects alternative 1 to be the most viable ($q_1 = \max_i q_i$) and there are two alternatives, a unique signal structure satisfies these two properties and induces a selection probability of $p \in [p_F, p_{max}]$ - the (adjusted) p-CMSS.

3.7.2 Continuous State Space

One critical assumption is that each alternative is either completely viable ($\theta_i = \bar{\theta} = 1$) or not viable at all ($\theta_i = \underline{\theta} = 0$). In this subsection, I illustrate that the characterization of feasibility fails to have a tractable structure for any selection probability $p \in (p_F, 1)$ when the viabilities are continuously distributed. For this subsection, I assume that there are two alternatives and $(\tilde{\theta}_1, \tilde{\theta}_2)$ are distributed on $[0,1]^2$ according to some differentiable and exchangeable CDF F. First, if the receiver's optimal action only depends on the expected viability $(q_1, ..., q_n)$ of alternatives, Proposition 3.1 holds and we can still apply the decomposition approach. However, this decomposition approach is only useful because the set of feasible conditional selection distributions has a familiar and well-studied structure for binary states of the world, i.e. (p, σ) is feasible if σ is a mean-preserving contraction of F_p^{CMSS}.

Following the discussion of the last subsection, the most informative conditional selection distribution is induced by a signal structure that leaves the receiver indifferent between selecting alternative 1 and 2 whenever he selects alternative 1 and does not pool two or more states (θ_1, θ_2) for which $\theta_1 \geqslant \theta_2$.

For a selection probability of $p = 1$, there is an (essentially) unique signal structure that satisfies these properties: The signal structure that pools state (θ_1, θ_2) with state (θ_2, θ_1) for all $(\theta_1, \theta_2) \in [0, 1]^2$. For this signal structure, the receiver is indifferent between selecting alternative 1 and 2 since $f(\theta_1, \theta_2) = f(\theta_2, \theta_1)$ by exchangeability and it obviously does not pool any two states for which $\theta_1 \geqslant \theta_2$. Also, for a selection probability of $p = p_F = 1/2$, the fully informative signal is the unique signal structure that satisfies these properties.

However, for an intermediate selection probability $p \in (p_F, 1)$ there does not exist a unique signal structure. To provide an intuition as to why this is the case, suppose that the sender wants to induce a selection probability just below 1, i.e. $p = 1 - \varepsilon$ for a vanishing ε. A natural starting point is the most-informative signal structure for $p = 1$, which pools (θ_1, θ_2) with (θ_2, θ_1). In order to reduce the selection probability and still satisfy the necessary conditions for a most-informative signal structure, some permutation pairs $(\theta_1, \theta_2), (\theta_2, \theta_1)$ must be disclosed instead of pooled.

We compare two alternative ways of breaking up the pooled permutation pairs. First, pooling all permutation pairs except for a small neighborhood around the pairs $(0, 1), (1, 0)$, and second, pooling all permutation pairs except for a small neighborhood around the pairs $(1, 0), (0, 1/2)$. Separating pairs around $(0, 1)$ and $(1, 0)$ reduces the probability mass of the conditional selection probability near $1/2$ and transfers this probability mass to 1 since when $(1, 0)$ is pooled with $(0, 1)$ the receiver expects the viability of 1 to be $1/2$ and when it is separated the receiver selects alternative 1 for the state of the world $(1, 0)$ and expects the viability of alternative 1 to be equal to 1. On the other hand separating pairs near $(0, 1/2)$ and $(1/2, 0)$ transfers probability mass from $1/4$ to $1/2$. The resulting conditional selection distributions cannot be ordered in terms of mean-preserving contractions.

This limitation clarifies why it is crucial that the state space is binary for each alternative. Note that if there are more than two alternatives and the state space is binary for each alternative, there also exists more than one signal structure for a fixed selection probability p, which leaves the receiver indifferent between selecting different alternatives whenever he selects alternative 1 and does not pool two or more states for which the receiver would select alternative 1.

However, the crucial difference between a binary state space and a continuous state space is that for a binary state space all conditional selection distributions obtained by signal structures satisfying the two above mentioned properties can be ordered in terms of mean-preserving spreads: To illustrate, let $N = 3$ and let $\theta_i \in \{0, 1\}$ for all $i = 1, 2, 3$. Instead of pooling all states of the world with the same mean with and below the cutoff mean $1/3$, the sender could pool all states of the world with a mean above $1/3$. The first signal structure (the CMSS), transfers probability mass from 2/3 to 1 relative to pooling all states of the world with the same mean, while the second signal structure transfers probability mass from 1/3 to 1, relative to pooling all states of the world with the same mean. It can be easily verified that the first conditional selection distribution is a mean-preserving spread of the second conditional selection distribution since we can obtain the second distribution from the first by pooling probability mass from 1/3 with the new probability mass on 1.

3.7.3 Asymmetric Selection

Lastly, I have assumed that the receiver selects the alternative that is the most likely to be viable. That is, if the receiver has a posterior belief (q_1, q_2), the receiver selects alternative 1 if and only if $q_1 \geqslant q_2$. Now, suppose that the receiver is biased against alternative 1 such that the receiver selects alternative 1 if and only if $q_1 \geqslant q_2 + \delta$ for some $\delta > 0$. Importantly, and contrary to the case with a symmetric selection, the receiver would not select alternative 1 if he knew that the state of the world is $(\underline{\theta}, \underline{\theta})$ or $(\bar{\theta}, \bar{\theta})$. Thus, the selection probability under full information p_F is equal to $f(\bar{\theta}, \underline{\theta})$. If the sender wants to induce a larger selection probability, then it is necessary to pool $(\bar{\theta}, \underline{\theta})$ with one of the remaining states. However, unlike the case with symmetric selection, there is no most informative way of doing this.

It seems natural that the sender might first pool $(\bar{\theta}, \underline{\theta})$ with $(\bar{\theta}, \bar{\theta})$ so that the receiver is exactly indifferent between alternative 1 and alternative 2, and next (if possible) pool $(\bar{\theta}, \underline{\theta})$ with $(\underline{\theta}, \underline{\theta})$. However, the optimality of this scheme heavily depends on the specific properties of the sender's continuation payoff $U^S(\cdot)$. A characterization of feasible selection pairs, similar to Theorem 3.1, is beyond the scope of this chapter.

3.8 Related Literature

This paper is related to the literature on Bayesian Persuasion initiated by Kamenica and Gentzkow (2011) and Rayo and Segal (2010). Closely related are articles that take a belief-based approach to information disclosure. Kamenica and Gentzkow (2011) and Aumann, Maschler, and Stearns (1995) show that a distribution of posterior beliefs can originate from Bayesian updating given some signal structure if and only if the distribution of posterior belief averages back to the prior belief.

Depending on the application, only some statistic of the distribution of posterior beliefs might be of interest. A common assumption is that only the mean of a posterior belief is relevant. Strassen's theorem (Strassen, 1965) implies that a distribution of posterior means can be induced by some signal structure if and only if the distribution of posterior means is dominated in the convex order by the distribution F of the posterior means induced by full information. If the state of the world is drawn from a subset of the real line, the convex order is equivalent to the mean-preserving spread order. Kleiner et al. (2021) characterize the extreme points of the set of CDFs dominated in the mean-preserving spread order by F. In a recent contribution, Yang and Zentefis (2022) consider problems in which preferences only depend on posterior quantiles. They show that a distribution of posterior quantiles is consistent with some signal structure if and only if the distribution is in a first-order stochastic interval by an upper and a lower truncation of the prior. In contrast, this article studies selection problems in which for any posterior belief the sender's expected payoff only depends on (i) whether the receiver selects the sender's favorite alternative and (ii) the marginal distribution of sender's favorite alternative conditional on (i), and we show that a decomposition pair (p, σ) is consistent with some signal structure if and only if σ is a MPC of a most-informative distribution F_p^{CMSS}.

The literature on multi-dimensional Bayesian persuasion is relatively sparse. Tamura (2018) studies Bayesian persuasion with a multi-dimensional state space and assumes a linear-quadratic specification of sender and receiver payoffs which reduces the sender's information disclosure problem to a choice of a covariance matrix of the receiver's posterior expectation. In contrast, I assume that the receiver solves a selection problem, i.e. the receiver selects one out of many alternatives and chooses an action for the selected alternative. In this setting, the sender's information disclosure problem reduces to a choice of a selection probability and a conditional selection distribution.

Applications of multi-dimensional Bayesian persuasion include Armstrong and Zhou (2022) who study how information disclosure of consumers' valuations over products affects competition between firms, and Lee (2021) who analyzes how an intermediary optimally persuades a consumer to buy a product from a recommended seller by strategically releasing information about match values as well as sellers' types. None of these papers offer a general approach to multi-dimensional Bayesian persuasion problems: Lee (2021) employs a guess-and-verify approach to solve the problem. Armstrong and Zhou (2022) reduce their problem to a uni-dimensional problem by restricting attention to two firms and cases in which the consumer always wants to buy the product such that only the (one-dimensional) difference in valuations is relevant.

Closest to this paper is Ichihashi (2020), which investigates the disclosure problem of a consumer who provides a multi-product seller with information about their valuations for various products. This information is then used by the seller to recommend a product to the consumer and price discriminate. Thus, Ichihashi (2020) studies a specific selection problem: the seller's selection is a recommended product and the seller's action is the price charged for the recommended product. However, there are several important differences between my work and Ichihashi's. First, while Ichihashi's model assumes that the sender is ex-ante indifferent between products, in my model the sender seeks to persuade the receiver to select a specific option. Second, whereas Ichihashi's focus is solely on the recommendation and pricing problem of multi-product sellers, I consider general selection problems. As a caveat, I restrict attention to binary state spaces. Finally, Ichihashi (2020) analyzes properties of the optimal information structure without characterizing it, whereas the approach of this paper can be used to explicitly characterize optimal information structures.

3.9 Conclusion

This chapter has studied the information design problem of a sender who wants to persuade a receiver to select favorable alternatives and take actions favorable to the sender. Under the assumption that the receiver finds it optimal to select the alternative he believes to be the most viable, I have used a decomposition approach to reformulate the multi-dimensional information design problem into a selection persuasion and an action persuasion problem. This decomposition separates the problem of finding the optimal signal structure that maximizes the sender's payoff from the receiver's selection from finding the optimal signal structure that maximizes the sender's payoff from the receiver's actions, given some selection probability. For particular sender preferences, I have characterized the optimal solutions, and established that the sender never benefits from increasing the number of alternatives. The results were applied to analyze economic applications.

The approach in this paper requires three restrictions: First, the prior distribution f is exchangeable. Second, the payoff function of the receiver is symmetric across alternatives. Third, every alternative is either viable ($\theta_k = \bar{\theta}$) or not viable ($\theta_k = \underline{\theta}$). The approach can be generalized to asymmetric prior distributions. However, the second and third restrictions are more substantial. Generalizations to non-binary state spaces and asymmetries in the payoff function are left as directions for future research.

3.A Appendix

3.A.1 Proof of Lemma 3.1

Proof. I show that we can improve upon any dominated signal structure (S, π) by constructing a related undominated signal structure $(\tilde{S}, \tilde{\pi})$. Let $\mathbf{q}^{s;\pi}$ denote the posterior beliefs induced by signal structure (S, π) and signal realization s. We define $\mathcal{Q} = \{\mathbf{q} \in [0,1]^N : q_1 = \max_i q_i\}$ as the set of all posterior beliefs for which the receiver selects alternative 1.

The persuasion value of (S, π) is

$$\begin{aligned} V((S,\pi)) &= \sum_{s \in S} \Big(\sum_{\theta \in \{\underline{\theta}, \bar{\theta}\}^N} f(\theta)\pi(s \mid \theta)\Big) U^S(\mathbf{q}^{s;\pi}) \\ &= \sum_{s \in S: \mathbf{q}^s \in \mathcal{Q}} \sum_{\theta \in \{\underline{\theta}, \bar{\theta}\}^N} f(\theta)\pi(s \mid \theta) U^S(q_1^{s;\pi}). \end{aligned}$$

Let $T_{(S,\pi)}$ denote all signal realization and state of the world pairs (s, θ) such that the receiver would select alternative 1 if he knew the state θ but does not select alternative upon learning that the signal realization is s. That is,

$$\mathcal{T}_{(S,\pi)} = \Big\{(s,\theta) \in S \times \{\underline{\theta}, \bar{\theta}\}^N : \pi(s \mid \theta) > 0, \big(\theta_1 = \bar{\theta} \text{ or } \theta_1 = ... = \theta_N = \underline{\theta}\big), \mathbf{q}^{s;\pi} \notin \mathcal{Q}\Big\}.$$

A signal structure (S, π) is said to be undominated if $T_{(S,\pi)} = \varnothing$ and a signal structure (S, π) is dominated if and only if $\mathcal{T}_{(S,\pi)} \neq \varnothing$.

Let $(s^*, \theta^*) \in \mathcal{T}_{(S,\pi)}$ and consider $(\tilde{S}, \tilde{\pi})$ defined by $\tilde{S} = S \cup \{s_{\theta^*}\}$ and

$$\tilde{\pi}(s \mid \theta) = \begin{cases} 0 & \text{if } s = s^* \text{ and } \theta = \theta^* \\ \pi(s^* \mid \theta^*) & \text{if } s = s_{\theta^*} \text{ and } \theta = \theta^*. \\ \pi(s \mid \theta) & \text{otherwise} \end{cases}$$

In words, $\tilde{\pi}$ mirrors π except that in state θ^*, the new signal structure $\tilde{\pi}$ fully discloses θ^* whenever π sends signal s^*. Note that $\mathbf{q}^{s^*;\pi} \notin Q$ implies $\mathbf{q}^{s^*;\tilde{\pi}} \notin \mathcal{Q}$ and $\mathbf{q}^{s;\pi} = \mathbf{q}^{s;\tilde{\pi}}$ for any $s \in S \backslash \{s^*\}$. Further,

$$\mathcal{Q} \ni \mathbf{q}^{s_{\theta^*};\tilde{\pi}} = \begin{cases} (1, q_2^{s_{\theta^*};\tilde{\pi}}, ..., q_N^{s_{\theta^*};\tilde{\pi}},) & \text{if } \theta_1^* = \bar{\theta} \\ (0, 0, ...0) & \text{if } \theta_1^* = ... = \theta_N^* = \underline{\theta} \end{cases}.$$

The persuasion value of $(\tilde{S}, \tilde{\pi})$ is

$$
\begin{aligned}
V((\tilde{S}, \tilde{\pi})) &= \sum_{s \in \tilde{S}} \sum_{\theta \in \{\underline{\theta}, \bar{\theta}\}^N} f(\theta)\pi(s \mid \theta) U^S(\mathbf{q}^{s;\tilde{\pi}}) \\
&= \sum_{s \in S: \mathbf{q}^s \in \mathcal{Q}} \sum_{\theta \in \{\underline{\theta}, \bar{\theta}\}^N} f(\theta)\pi(s \mid \theta) U^S(q_1^{s;\pi}) + f(\theta^*)\pi(s \mid \theta^*) U^S(q_1^{s_{\theta^*};\tilde{\pi}}) \\
&= V((S, \pi)) + f(\theta^*)\pi(s \mid \theta^*) U^S(q_1^{s_{\theta^*};\tilde{\pi}}). \\
&\geqslant V((S, \pi))
\end{aligned}
$$

where the inequality is strict if $U^S(x) > 0$ for any x. We can use this procedure iteratively to eliminate any $(s, \theta) \in \mathcal{T}_{(S,\pi)}$, and end up with an undominated signal structure that improves upon (S, π). □

3.A.2 Proof of Lemma 3.2

Proof. Suppose that (p, σ) can be induced by an undominated signal structure (S, π). Recall that the receiver only selects alternative 1 upon learning a signal realization $s \in S$ if the induced posterior belief $\mathbf{q}^{s,\pi}$ satisfies $\mathbf{q}^{s,\pi} \in Q = \{\mathbf{q} \in [0,1]^N : q_1 = \max_i q_i\}$. The selection probability p is given by

$$
\begin{aligned}
p &= \sum_{\theta_1, ..., \theta_N} f(\theta_1, ..., \theta_N) \sum_{s \in S: \mathbf{q}^{s,\pi} \in \mathcal{Q}} Pr(s \mid \theta_1, ..., \theta_N) \\
&\geqslant \sum_{\theta: \theta_1 = \bar{\theta} \text{ or } \theta_1 = ... = \theta_N = \underline{\theta}} f(\theta_1, ..., \theta_N) \sum_{s \in S: \mathbf{q}^s \in Q} Pr(s \mid \theta_1, ..., \theta_N) \\
&= \sum_{\theta: \theta_1 = \bar{\theta} \text{ or } \theta_1 = ... = \theta_N = \underline{\theta}} f(\theta_1, ..., \theta_N) = p_F.
\end{aligned}
$$

Further note that by definition $(p, \sigma) = (p_G, \sigma_G)$ for some undominated G and any undominated G has to satisfy $q_1 = 0$ for any $\mathbf{q} \in (\text{supp } G) \backslash \mathcal{Q}$. Thus,

$$
\begin{aligned}
p \int x d\sigma(x) &= \sum_{\mathbf{q} \in \text{supp G} \cap \mathcal{Q}} g(\mathbf{q}) q_1 \\
&= \sum_{\mathbf{q} \in \text{supp G}} g(\mathbf{q}) q_1 \\
&= q_F.
\end{aligned}
$$

where the last equality is implied by $G \leqslant_{cx} F$.[16] □

[16] $G \leqslant_{cx} F$ immediately implies that $\sum_{\mathbf{q} \in \text{supp G}} g(\mathbf{q}) l(\mathbf{q}) = \sum_{\mathbf{q} \in \text{supp G}} f(\mathbf{q}) l(\mathbf{q})$ for any linear $l(\cdot)$.

3.A.3 Proof of Lemma 3.3

Proof. Suppose that $\sigma' \in \mathrm{MPC}(\sigma)$. Let supp $\sigma' = \{x'_1, ..., x'_M\}$. Then, for any $x \in$ supp σ there exists a probability vector $\alpha^x = (\alpha^x_1, ..., \alpha^x_M)$ such that

$$\sum_{x \in \text{supp } \sigma} \alpha_i^x x = x'_i$$

and

$$\sum_{x \in \text{supp } \sigma} \alpha_i^x \sigma(x) = \sigma(x'_i).$$

Now suppose that (p, σ) is feasible. I.e., there exists some G for which $G \leqslant_{cx} F$ and $(p, \sigma) = (p_G, \sigma_G)$. Recall that the receiver only selects alternative 1 if the posterior belief satisfies $\mathbf{q} \in Q = \{\mathbf{q}' \in [0,1]^N : q'_1 = \max_i q'_i\}$. For any $x \in$ supp σ, define the set of posterior beliefs corresponding to x as $Q(x) = \{\mathbf{q} \in \mathcal{Q} : q_1 = x\}$. For any $\mathbf{q} \in Q(x) = \{\mathbf{q} \in [0,1]^N : q_1 = \max_i q_i\}$, we define the relative proportion of $\mathbf{q}$ as $\beta_{\mathbf{q}} = g(\mathbf{q}) / \sum_{\mathbf{q}' \in Q(x)} g(\mathbf{q}')$. Now, define the distribution of posterior beliefs G' by

$$g'(\mathbf{q}') = \begin{cases} g(\mathbf{q}') & \text{if } \mathbf{q}' \notin \mathcal{Q} \\ \sum_{x \in \text{supp } \sigma} \alpha_i^x \sum_{\mathbf{q} \in Q(x)} g(\mathbf{q}) & \text{if } \mathbf{q}' = \sum_{x \in \text{supp } \sigma} \alpha_i^x \sum_{\mathbf{q} \in \mathcal{Q}(x)} \beta_{\mathbf{q}} \mathbf{q} \text{ for some } i = 1, ..., N \\ 0 & \text{otherwise} \end{cases}.$$

Note that we construct G' by pooling probability mass from points on the support of G onto its barycenter. Thus, G' is a fusion of G (Elton and Hill, 1992) which is equivalent to $G' \leqslant_{cx} G$.

Observe that $p_{G'} = p_G = p$, and

$$\begin{aligned} \sigma_{G'}(x'_i) &= \frac{1}{p_G} \Big(\sum_{x \in \text{supp } \sigma} \alpha_i^x \sum_{\mathbf{q} \in Q(x)} g(\mathbf{q})\Big) \\ &= \sum_{x \in \text{supp } \sigma} \alpha_i^x \left(\frac{1}{p_G} \sum_{\mathbf{q} \in Q(x)} g(\mathbf{q}) \right) \\ &= \sum_{x \in \text{supp } \sigma} \alpha_i^x \sigma(x) = \sigma(x'_i). \end{aligned}$$

Hence, G' has decomposition (p, σ'), which shows that (p, σ') is feasible. □

3.A.4 Proof of Theorem 3.1

First, we show that F_p^{CMSS} is well defined. Let

$$F_p^{CMSS}(x;y) = \begin{cases} \frac{F^{Mean}(x)}{p} & \text{if } x < (F^{Mean})^{-1}(y) \\ \frac{y}{p} & \text{if } (F^{Mean})^{-1}(y) \leqslant x < 1 \\ 1 & \text{if } x \geqslant 1 \end{cases} \tag{3.7}$$

Note that $F_p^{CMSS}(x) := F_p^{CMSS}(x; y_p)$ where y_p solves $\int x dF_p^{CMSS}(x; y_p) = \frac{q_F}{p}$.

Lemma 3.4. *F_p^{CMSS} is well-defined. That is, $\int x dF_p^{CMSS}(x; y) = \frac{q_F}{p}$ has a unique solution y_p and $F_p^{CMSS}(\cdot) = F_p^{CMSS}(\cdot; y_p)$ is a CDF.*

Proof. First, we argue that $F_p^{CMSS}(x; y)$ is a CDF for any $y \in [0, p]$. By definition, F_p^{CMSS} is right-continuous, and $F_p^{CMSS}(x) = 0$ for any $x < 0$ as well as $F_p^{CMSS}(x) = 1$ for any $x \geqslant 1$. Note that $\frac{y}{p} \geqslant \frac{F^{Mean}(x)}{p}$ for any $x \leqslant (F^{Mean})^{-1}(y)$ and $\frac{y}{p} \leqslant 1$ since $y \leqslant p$. Thus, $F_p^{CMSS}(x; y)$ is a CDF for any $y \in [0, p]$.

Next, we show that $\bar{\mu}_p(y) := \int x dF_p^{CMSS}(x; y) = \frac{q_F}{p}$ has a unique solution $y_p \in [F^{Mean}(0), p]$. Note that $\bar{\mu}_p(y)$ is continuous and strictly decreasing in y, and

$$\begin{aligned} \bar{\mu}_p(F^{Mean}(0)) &= 1 - \frac{F^{Mean}(0)}{p} \\ &= \frac{1}{p}(p - f(0, ..., 0)) \geqslant \frac{q_F}{p} \end{aligned}$$

where the second equality follows from $F^{Mean}(0) = f(0, ..., 0)$, and the inequality follows from $p \geqslant p_F = q_F + f(0, ..., 0)$. Also, observe that

$$\begin{aligned} \bar{\mu}_p(p) &= \frac{1}{p} \int_0^{(F^{Mean})^{-1}(p)} x dF^{Mean}(x) \\ &\leqslant \frac{1}{p} \int_0^1 x dF^{Mean}(x) \\ &= \frac{q_F}{p}. \end{aligned}$$

where the first equality follows from supp $F_p^{CMSS}(\cdot, y = p) \subseteq [0, (F^{Mean})^{-1}(p)]$ and the final equality follows since F^{Mean} has expected value q_F.

The intermediate value theorem implies that for any $p \in [p_F, 1]$ there exists some $y_p \in [F^{Mean}(0), p]$ such that $\bar{\mu}_p(y_p) = \frac{q_F}{p}$, and it is unique since $\bar{\mu}_p(\cdot)$ is strictly decreasing in y. □

We need the following Lemma to prove Theorem 3.1.

Lemma 3.5. *Let (p, σ) be undominated. The following statements are equivalent:*

(i) $\int_0^x F^{Mean}(z)dz \geqslant \int_0^x p \cdot \sigma(z)dz$ *for all* $x \in [0,1]$.
(ii) $\int_0^x F_p^{CMSS}(z)dz \geqslant \int_0^x \sigma(z)dz$ *for all* $x \in [0,1]$ *with equality at* $x = 1$.

Proof. Note that by Lemma 3.2, we know that $p \in [p_F, 1]$ and $\int x d\sigma(x) = q_F/p$.

"$(ii) \implies (i)$" Suppose that $\int_0^x F_p^{CMSS}(z)dz \geqslant \int_0^x p \cdot \sigma(z)dz$ with equality at $x = 1$. From the definition of F_p^{CMSS}, it is immediate that $\frac{1}{p}F^{Mean}(z) \geqslant F_p^{CMSS}(z)$ for all $z \in [0,1]$ which implies (i).

"$(i) \implies (ii)$" Suppose that $\int_0^x F^{Mean}(z)dz \geqslant \int_0^x p \cdot \sigma(z)dz$ for all $x \in [0,1]$. We want to show that $\int_0^x F_p^{CMSS}(z)dz \geqslant \int_0^x \sigma(z)dz$ for all $x \in [0,1]$ with equality at $x = 1$.

Case 1: Let $x = 1$. By definition of F_p^{CMSS}, we have $\int x dF_p^{CMSS}(x) = q_F/p$. Integration by parts yields that $\int_0^1 \sigma(z)dz = 1 - \int x d\sigma(x) = 1 - q_F/p$, and $\int_0^1 F^{Mean}(z)dz = 1 - \int x dF_p^{CMSS}(x) = 1 - q_F/p$ Thus,

$$\int_0^1 F_p^{CMSS}(x)dx = \int_0^1 \sigma(z)dz.$$

Case 2: Let $x \in [(F^{Mean})^{-1}(y_p), 1)$. Note that $\int_0^x F_p^{CMSS}(z)dz \geqslant \int_0^x \sigma(z)dz$ if and only if

$$\int_x^1 \sigma(z)dz \geqslant \int_x^1 F_p^{CMSS}(z)dz = (1-x)\frac{y_p}{p}.$$

which is satisfied for all $x \in [(F^{Mean})^{-1}(y_p), 1)$ if and only if it is satisfied for $x = (F^{Mean})^{-1}(y_p)$ (left-hand side concave and decreasing, right-hand side linear and decreasing in x). Now,

$$\begin{aligned}
\int_{(F^{Mean})^{-1}(y_p)}^1 \sigma(z)dz &= \int_0^1 \sigma(z)dz - \int_0^{(F^{Mean})^{-1}(y_p)} \sigma(z)dz \\
&\geqslant 1 - \frac{q_F}{p} - \int_0^{(F^{Mean})^{-1}(y_p)} \frac{F^{Mean}(z)}{p}dz \\
&= \int_0^1 F_p^{CMSS}(z)dz - \int_0^{(F^{Mean})^{-1}(y_p)} F_p^{CMSS}(z)dz \\
&= \int_{(F^{Mean})^{-1}(y_p)}^1 F_p^{CMSS}(z)dz
\end{aligned}$$

where the inequality is implied by (i) and the second equality follows from $F_p^{CMSS}(z) = F^{Mean}(z)/p$ for all $z < (F^{Mean})^{-1}(y_p)$.

Case 3: Assume $x < (F^{Mean})^{-1}(y_p)$. In this case, $F_p^{CMSS} = F^{Mean}(z)$ for all $z \leqslant x$ such that $\int_0^x F_p^{CMSS}(z)dz \geqslant \int_0^x \sigma(z)dz$ immediately follows from (i). □

Now we are ready to present the proof of Theorem 3.1.

Proof. " undominated (p, σ) feasible $\implies$ σ is a MPC of F_p^{CMSS}" :

We argue that $\int_0^x F^{Mean}(z)dz \geqslant \int_0^x p \cdot \sigma(z)dz$ for any $x \in [0,1]$, which is equivalent to $\sigma \in \text{MPC}(F_p^{CMSS})$ by Lemma 3.5. If (p, σ) is feasible, then there exists $G \in \Delta[0,1]^N$

satisfying $G \leqslant_{cx} F$ such that $(p, \sigma) = (p_G, \sigma_G)$. If $G \leqslant_{cx} F$, then $\bar{G}$ is a MPC of F^{Mean}. In order to see this, let $c : \mathbb{R} \to \mathbb{R}$ be convex and note that

$$\int c(x) dG^{Mean}(x) = \int c(\sum_i \mathbf{x}_i/N) dG(\mathbf{x}) \leqslant \int c(\sum_i \mathbf{x}_i/N) dF(\mathbf{x}) = \int c(x) dF^{Mean}(x).$$

Next, we argue that $p_G \cdot \sigma_G(x) \leqslant G^{Mean}(x)$ for any $x \in [0,1]$. Observe that

$$\begin{aligned}
\frac{1}{p_G} G^{Mean}(x) &= \frac{1}{p_G} Pr_G(\sum q_i/N \leqslant x) \\
&= \frac{1}{p_G} \Big(p_G Pr_G(\sum q_i/N \leqslant x \mid \mathbf{q} \in \mathcal{Q}) + (1 - p_G) Pr_G(\sum q_i/N \leqslant x \mid \mathbf{q} \notin \mathcal{Q}) \Big) \\
&\geqslant Pr_G(\sum q_i/N \leqslant x \mid \mathbf{q} \in \mathcal{Q}). \\
&\geqslant Pr_G(q_1 \leqslant x \mid \mathbf{q} \in \mathcal{Q}) \\
&= \sigma_G(x)
\end{aligned}$$

where the second equality follows from the law of total probability, and the second inequality holds since $q_1 \geqslant \sum \mathbf{q}_i/N$ for any $\mathbf{q} \in \mathcal{Q}$. Since G^{Mean} is a MPC of F^{Mean} , we know that $\int G^{Mean}(z)dz \leqslant \int_0^x F^{Mean}(z)dz$ for any $x \in [0,1]$. Thus, we obtain

$$\int_0^x p_G \sigma_G(z) dz \leqslant \int_0^x G^{Mean}(z) dz \leqslant \int_0^x F^{Mean}(z) dz.$$

Hence, $\sigma \in \mathrm{MPC}(F_p^{CMSS})$ by Lemma 3.5.

"σ is a MPC of $F_p^{CMSS} \implies$ undominated (p, σ) feasible ":

Suppose that σ is a MPC of F_p^{CMSS} where $p \in [p_F, 1]$. Lemma 3.3 implies that (p, σ) is feasible if (p, F_p^{CMSS}) is feasible. Thus, it remains to show that (p, F_p^{CMSS}) is feasible. We construct a signal structure that induces (p, F_p^{CMSS}). We define the *cutoff-mean signal structure* (S, τ) with parameters $(\bar{x}, \alpha) \in [0,1]^2$ by signal space $S = \{\frac{1}{N}, ..., \frac{N-1}{N}\} \bigcup \{\underline{\theta}, \bar{\theta}\}^N$ and disclosure rule

$$\tau(s \mid \theta_1, ..., \theta_N) = \begin{cases} 1 & \text{if } \sum \theta_i/N < \bar{x} \text{ and } s = \sum \theta_i/N \\ \alpha & \text{if } \sum \theta_i/N = \bar{x} \text{ and } s = \sum \theta_i/N \\ 1 - \alpha & \text{if } \sum \theta_i/N = \bar{x} \text{ and } s = (\theta_1, .., \theta_N) \\ 1 & \text{if } \sum \theta_i/N > \bar{x} \text{ and } s = (\theta_1, .., \theta_N) \\ 0 & \text{otherwise} \end{cases}.$$

That is, the *cutoff-mean signal structure* with parameters $(\bar{x}, \alpha)$ fully discloses a state $(\theta_1, ..., \theta_N)$ if the average state is $\sum \theta_i/N$ is above a cutoff $\bar{x}$, discloses the average state $\sum \theta_i/N$ if the average state is below $\bar{x}$, and randomizes between disclosing state and average state if $\sum \theta_i/N = \bar{x}$.

Remember that y_p solves $\int x dF_p^{CMSS}(x;y) = q_F/p$ and $F_p^{CMSS}(\cdot) := F_p^{CMSS}(\cdot\,; y_p)$. Let $x_p = (F_p^{CMSS})^{-1}(y_p)$ and $\alpha_p = (y_p - \lim_{x \nearrow x_p} F_p^{CMSS}(x))/f_p^{CMSS}(x_p)$. Therefore,

$$\int x dF_p^{CMSS}(x) = \sum_{x<x_p} \frac{f^{Mean}(x)}{p} x + \frac{f^{Mean}(x_p)\alpha_p}{p} \cdot x_p + (1 - \frac{y_p}{p}) \cdot 1 = \frac{q_F}{p}$$

which is equivalent to $p = y_p + q_F - \sum_{x<x_p} f^{Mean}(x)x - f^{Mean}(x_p)\alpha_p \cdot x_p$.

We show that the (x_p, α_p)-cutoff-mean signal structure induces (p, F_p^{CMSS}).

First, note that the receiver selects alternative 1 whenever he receives some signal $s = \sum \theta_i/N$ or a signal $s = (\theta_1, ..., \theta_N)$ for which $\theta_1 = 1$. Thus, the induced selection probability is

$$\begin{aligned}
p_{\alpha_p,x_p} &= \sum_{x<x_p} f^{Mean}(x) \cdot 1 + f^{Mean}(x_p)\alpha_p + f^{Mean}(x_p)(1-\alpha_p)x_p + \sum_{x>x_p} f^{Mean}(x)x \\
&= y_p + f^{Mean}(x_p)(1-\alpha_p)x_p + \sum_{x>x_p} f^{Mean}(x) \cdot x \\
&= y_p + \sum_{x} f^{Mean}(x)x - \sum_{x<x_p} f^{Mean}(x)x - f^{Mean}(x_p)\alpha_p x_p \\
&= q_F + y_p - \sum_{x<x_p} f^{Mean}(x)x - f^{Mean}(x_p)\alpha_p x_p \\
&= p.
\end{aligned}$$

Then, the induced selection distribution σ_{α_p,x_p} is given by

$$\begin{aligned}
\sigma(x) &= \begin{cases} \frac{1}{p} F^{Mean}(x) & \text{if } x < x_p \\ \frac{1}{p}\left(F^{Mean}(x_p - \frac{1}{N}) + F^{Mean}(x_p)\alpha_p\right) = \frac{y_p}{p} & \text{if } x \in [x_p, 1) \\ 1 & \text{if } x \geqslant 1 \end{cases} \\
&= F_p^{CMSS}(x).
\end{aligned}$$

which completes the proof. □

3.A.5 Proof of Proposition 3.4

Proof. Since U^S is convex, it immediately implies that F_p^{CMSS} is optimal in the p-action persuasion problem for any $p \in [p_F, 1]$. Each p corresponds to some (x_p, α_p) where $x_p \in \{0, \frac{1}{N}, ..., 1\}$ and $\alpha_p \in [0, 1]$ such that $p = p_{\alpha_p,x_p}$ (see proof of Theorem 3.1). The sender's expected payoff from F_p^{CMSS} is

$$\sum_{x<x_p} f^{Mean}(x)U^S(x) + \alpha_p f^{Mean}(x_p)U^S(x_p) + (1-\alpha_p)f^{Mean}(x_p)x_p U^S(1) + U^S(1) \sum_{x>x_p} x_p f^{Mean}(x).$$

It is straightforward to see that the maximum must be achieved at some $\alpha_p \in \{0,1\}$. Further, note that $(x_p, 0)$ achieves the same payoff as $(x_p - \frac{1}{N}, 1)$, so setting $\alpha = 1$ is without loss of generality. Assuming that $p = p_{1,x_p}$, the payoff from F_p^{CMSS} is

$$\sum_{x \leqslant x_p} f^{Mean}(x) U^S(x) + U^S(1) \sum_{x > x_p} x_p f^{Mean}(x).$$

Therefore, the optimal selection probability is the selection probability $p \in [p_F, 1]$ corresponding to the cutoff x_p where x_p solves

$$\max_{x_p \in \{0, \frac{1}{N}, \ldots, 1\}} \Lambda(x_p) := \sum_x f^{Mean}(x) \left[\mathbf{1}_{x \leqslant x_p} U^S(x) + \mathbf{1}_{x > x_p} x U^S(1)\right].$$

Thus, the optimal cutoff x satisfies $\Lambda(x) - \Lambda(x') \geqslant 0$ for all $x' \in \{0, \frac{1}{N}, \ldots, 1\}$.

Let $\bar{x} = \min\{x \in [0,1] : U^S(x) = xU^S(1)\}$ and $x^* = \lfloor N\bar{x} \rfloor / N$. We show $\Lambda(x) \geqslant \Lambda(x - \frac{1}{N})$ for any $x \leqslant x^*$ such that $x - \frac{1}{N} \geqslant 0$, and $\Lambda(x) \geqslant \Lambda(x + \frac{1}{N})$ for any $x \geqslant x^*$ such that $x + \frac{1}{N} \leqslant 1$, which implies by induction that $x^* \in \arg\max_{x_p \in \{0, \frac{1}{N}, \ldots, 1\}} \Lambda(x_p)$.

Suppose that $1 - \frac{1}{N} \geqslant x \geqslant x^*$. Note that $(x + \frac{1}{N}) \in [x^*, 1]$ implies that $U^S(x + \frac{1}{N}) \leqslant (x + \frac{1}{N}) U^S(1)$ since U^S is convex. Thus,

$$\begin{aligned}\Lambda(x) - \Lambda(x + \frac{1}{N}) &= f^{Mean}(x + \frac{1}{N}) \left[U^S(1)(x + \frac{1}{N}) - U^S(x + \frac{1}{N})\right] \\ &\geqslant 0\end{aligned}$$

for any $x \in \{x^*, \ldots, 1 - \frac{1}{N}\}$.

Suppose that $x \leqslant x^*$. Note that $x^* \leqslant \bar{x}$, and thus $x \leqslant \bar{x}$ which implies $U^S(x) \geqslant xU^S(1)$ by convexity of U^S and by $U^S(0) \geqslant 0$. Hence,

$$\begin{aligned}\Lambda(x) - \Lambda(x - \frac{1}{N}) &= f^{Mean}(x) \left[U^S(x) - xU^S(1)\right] \\ &\geqslant 0\end{aligned}$$

for any $x \in \{\frac{1}{N}, \ldots, x^*\}$. □

3.A.6 Proof of Proposition 3.6

Outline Proposition 3.6 is proven by Lemma 3.6, 3.7, and 3.8. Lemma 3.6 shows that the CDF of the average viability F_N^{Mean} is monotonically decreasing in N in the sense of the stochastic convex order. This Lemma generalizes a result of Müller and Stoyan (2002) (Corollary 1.5.24) to exchangeable (and binary) random variables. Lemma 3.7 shows that this property is inherited by any cutoff-mean conditional selection distributions. That is, $F_{p;N+1}^{CMSS} \in \mathrm{MPC}(F_{p;N}^{CMSS})$ for any $p \in [p_{F;N}, 1]$.

Lemma 3.7 immediately implies that

$$V_N^* \geqslant \max_{p\in[p_{F;N},1]} p \cdot \max_{\sigma\in \mathrm{MPS}(F_{p;N+1}^{CMSS})} \int U^S(x)d\sigma(x)$$

which implies $V_N^* \geqslant V_{N+1}^*$ if a selection probability $p \in [p_{F;N}, 1]$ is optimal in the $(N+1)$-dimensional problem. Lemma 3.8 implies that $V_N^* \geqslant V_{N+1}^*$ if $p \in [p_{F;(N+1)}, p_F)$ is optimal in the $(N+1)$-dimensional problem, and thus completes the proof of Proposition 3.6.

Lemma 3.6. *F_N^{Mean} is a MPS of F_{N+1}^{Mean} for all $N \in \mathbb{N}$*

Proof. Note that f_N^{Mean} has positive probability on $\{0, \frac{1}{N}, ..., 1\}$ while f_{N+1}^{Mean} has positive probability on $\{0, \frac{1}{N+1}, ..., 1\}$. Furthermore,

$$\begin{aligned} f_N^{Mean}(\frac{i}{N}) &:= \sum_{\theta_1,...,\theta_N:\sum\theta_j=i} f_N(\theta_1,...,\theta_N) \\ &= \binom{N}{i} f_N(\underbrace{1,...,1}_{i},\underbrace{0,...,0}_{N-i}) \end{aligned}$$

where the second equality holds since f is exchangeable.

Let $i \in \{1, ..., N\}$, and note that $\sum_{j=1}^{N+1}\theta_j = i$ can only be satisfied if either $\sum_{j=1}^{N}\theta_j = i$ and $\theta_{N+1} = 0$ or $\sum_{j=1}^{N}\theta_j = i-1$ and $\theta_{N+1} = 1$. We denote the distribution of θ_{N+1} conditional on $\theta_1, ..., \theta_N$ by $f_{N+1}(\theta_{N+1} \mid \theta_1, ..., \theta_N)$. Since $(\theta_i)_i$ are exchangeable, $f_{N+1}(\cdot \mid \theta_1, ..., \theta_N)$ can only depend on $\sum_{j=1}^{N}\theta_j$. To save on notation, we denote $\alpha_i = f_{N+1}(\theta_{N+1} = 1 \mid \sum_{j=1}^{N}\theta_j = i)$ which yields

$$f_N(\underbrace{1,...,1}_{i-1},\underbrace{0,...,0}_{N-i+1})\alpha_{i-1} = f_{N+1}(\underbrace{1,...,1}_{i},\underbrace{0,...,0}_{N+1-i}) = f_N(\underbrace{1,...,1}_{i},\underbrace{0,...,0}_{N-i})(1-\alpha_i).$$

With these insights, we can rewrite f_{N+1}^{Mean} as follows

$$\begin{aligned} f_{N+1}^{Mean}(\frac{i}{N+1}) &= \binom{N+1}{i} f_{N+1}(\underbrace{1,...,1}_{i},\underbrace{0,...,0}_{N+1-i}) \\ &= \left[\binom{N}{i-1} + \binom{N}{i}\right] f_{N+1}(\underbrace{1,...,1}_{i},\underbrace{0,...,0}_{N+1-i}) \\ &= \binom{N}{i-1} f_N(\underbrace{1,...,1}_{i-1},\underbrace{0,...,0}_{N-i+1})\alpha_{i-1} + \binom{N}{i} f_N(\underbrace{1,...,1}_{i},\underbrace{0,...,0}_{N-i})(1-\alpha_i) \\ &= f_N^{Mean}(\frac{i-1}{N})\alpha_{i-1} + f_N^{Mean}(\frac{i}{N})(1-\alpha_i) \end{aligned}$$

Thus, the probability mass of f_{N+1}^{Mean} on $\frac{i}{N+1}$ can be obtained by taking fraction α_{i-1}

of the probability mass of f_N^{Mean} on $\frac{i-1}{N}$ and fraction $(1-\alpha_i)$ of the probability mass of f_N^{Mean} on $\frac{i}{N}$.

Further, note that

$$\frac{f_N^{Mean}(\frac{i}{N})(1-\alpha_i)}{\binom{N}{i}} = f_{N+1}(\underbrace{1,...,1}_{i},\underbrace{0,...,0}_{N+1-i}) = \frac{f_N^{Mean}(\frac{i-1}{N})\alpha_{i-1}}{\binom{N}{i-1}}$$

which implies

$$f_N^{Mean}(\frac{i-1}{N})\alpha_{i-1} = \frac{i}{N+1-i} f_N^{Mean}(\frac{i-1}{N})(1-\alpha_i)$$

Pooling fraction α_{i-1} of the probability mass of $\frac{i-1}{N}$ and fraction $(1-\alpha_i)$ of the probability mass of $\frac{i}{N}$ yields a barycenter of

$$\frac{f_N^{Mean}(\frac{i-1}{N})\alpha_{i-1}}{f_N^{Mean}(\frac{i-1}{N})\alpha_{i-1} + f_N^{Mean}(\frac{i}{N})(1-\alpha_i)}\frac{i-1}{N} + \frac{f_N^{Mean}(\frac{i}{N})(1-\alpha_i)}{f_N^{Mean}(\frac{i-1}{N})\alpha_{i-1} + f_N^{Mean}(\frac{i}{N})(1-\alpha_i)}\frac{i}{N} = \frac{i}{N+1}.$$

This shows that the probability function f_{N+1}^{Mean} can be obtained by pooling probability mass of neighboring points on the support of f_N^{Mean} together on its barycenter. Thus, f_{N+1}^{Mean} is a fusion of f_N^{Mean}, which is equivalent to $F_{N+1}^{Mean} \in \text{MPC}(F_N^{Mean})$ (Elton and Hill, 1992). □

Lemma 3.7. *Let $p \in [p_{F;N}, 1]$. Then, $F_{p;N+1}^{CMSS}$ is a MPC of $F_{p;N}^{CMSS}$.*

Proof. From Lemma 3.6 we know that F_N^{Mean} is a MPS of F_{N+1}^{Mean} for all $N \in \mathbb{N}$. Therefore $\int_0^x F_N^{Mean}(z)dz \geqslant \int_0^x F_{N+1}^{Mean}(z)dz$ for all $x \in [0,1]$ with equality at $x = 1$.

Remember that $F_{p;N}^{CMSS}(x) := F_{p;N}^{CMSS}(x; y_{p;N})$ where $F_{p;N}^{CMSS}(x;y)$ is defined in (3.A.4) and $y_{p;N}$ solves $\int F_{p;N}^{CMSS}(x; y_{p;N})dx = 1 - \frac{q_F}{p}$. Hence, by definition, $F_{p;N}^{CMSS}$ and $F_{p;N+1}^{CMSS}$ both yield an expected value of $\frac{q_F}{p}$. Therefore, $\int_0^1 F_{p;N}^{CMSS}(z)dz = \int_0^1 F_{p;N+1}^{CMSS}(z)dz$. Suppose for the moment that $y_{p;N+1} \geqslant y_{p;N}$. Denote $x_{p;N} = (F_{p;N}^{CMSS})^{-1}(y_{p;N})$ and define $x_{p;N+1}$ analogously. Then,

$$\int_0^x F_{p;N}^{CMSS}(z)dz = \begin{cases} \int_0^x F_N^{Mean}(z)dz & \text{if } x \leqslant x_{p;N} \\ \int_0^{x_{p;N}} F_N^{Mean}(z)dz + (x - x_{p;N})y_{p;N} & \text{if } x \geqslant x_{p;N} \end{cases}$$

Let

$$\begin{aligned} \psi_p(x) &:= \int_0^x F_{p;N}^{CMSS}(z)dz - \int_0^x F_{p;N+1}^{CMSS}(z)dz \\ &= \int_x^1 F_{p;N+1}^{CMSS}(z)dz - \int_x^1 F_{p;N}^{CMSS}(z)dz \end{aligned}$$

We want to show that $\psi_p(x) \geqslant 0$ for all $x \in [0,1]$ and $p \in [p_{F;N}, 1]$ which implies

$F^{CMSS}_{p;N+1} \in \mathrm{MPC}(F^{CMSS}_{p;N})$. Suppose first that $x_{p;N} \leqslant x_{p;N+1}$. Then,

$$\psi_p(x) = \begin{cases} \int_0^x (F^{Mean}_N(z) - F^{Mean}_{N+1}(z)dz & \text{if } x \leqslant x_{p;N} \\ \psi_p(x_{p;N}) + \int_{x_{p;N}}^x (y_{p;N} - F^{Mean}_{N+1}(z))dz & \text{if } x_{p;N} < x < x_{p;N+1} \\ -(1-x)(y_{p;N} - y_{p;N+1}) & \text{if } x \geqslant x_{p;N+1} \end{cases}.$$

In the first case, $\psi_p(x) \geqslant 0$ since $F^{Mean}_{N+1} \in \mathrm{MPC}(F^{Mean}_N)$. In the last case $\psi_p(x) \geqslant 0$ since $y_{p;N+1} \geqslant y_{p;N}$, and for the second case, $\psi_p(x) \geqslant \min\{\psi_p(x_{p;N}), \psi_p(x_{p;N+1})\}$,[17] and thus $\psi_p(x) \geqslant 0$ for all x. Now, suppose that $x_{p;N} \geqslant x_{p;N+1}$. Then,

$$\psi_p(x) = \begin{cases} \int_0^x (F^{Mean}_N(z) - F^{Mean}_{N+1}(z)dz & \text{if } x \leqslant x_{p;N+1} \\ \psi_p(x_{p;N+1}) + \int_{x_{p;N+1}}^x (F^{Mean}_N(z) - y_{p;N+1})dz & \text{if } x_{p;N+1} < x < x_{p;N} \\ -(1-x)(y_{p;N} - y_{p;N+1}) & \text{if } x \geqslant x_{p;N+1} \end{cases}.$$

The first and last cases are analogous, and in the second case we have that $\psi_p(x) \geqslant \psi_p(x_{p;N}) \geqslant 0$.[18] Hence, $\int_0^x F^{CMSS}_{p;N}(z)dz - \int_0^x F^{CMSS}_{p;N+1}(z)dz \geqslant 0$ with equality at $x = 1$ such that $F^{CMSS}_{p;N}$ is a mean-preserving contraction of $F^{CMSS}_{p;N+1}$.

The remainder of the proof argues that $y_{p,N+1} \geqslant y_{p,N}$ is indeed the case. First, note that $\int_0^1 F^{CMSS}_{p;N}(x;y)dx$ is strictly increasing in y, and second, it holds that for any $y \in [0,1]$,

$$\int_0^1 F^{CMSS}_{p;N}(x;y)dz \geqslant \int_0^1 F^{CMSS}_{p;N+1}(x;y)dz.$$

First, suppose that y satisfies $(F^{CMSS}_{p;N+1})^{-1}(y) \geqslant (F^{CMSS}_{p;N})^{-1}(y)$. In this case,

$$\begin{aligned} \int_0^1 F^{CMSS}_{p;N}(z;y)dz &= \int_0^{(F^{CMSS}_{p;N})^{-1}(y)} F^{Mean}_N(z)dz + \int_{(F^{CMSS}_{p;N})^{-1}(y)}^1 ydz \\ &\geqslant \int_0^{(F^{CMSS}_{p;N})^{-1}(y)} F^{Mean}_{N+1}(z)dz + \int_{(F^{CMSS}_{p;N})^{-1}(y)}^1 ydz \\ &\geqslant \int_0^{(F^{CMSS}_{p;N+1})^{-1}(y)} F^{Mean}_{N+1}(z)dz + \int_{(F^{CMSS}_{p;N+1})^{-1}(y)}^1 ydz \\ &= \int_0^1 F^{CMSS}_{p;N+1}(z;y)dz. \end{aligned}$$

where the first inequality is implied by $F^{Mean}_{N+1} \in \mathrm{MPC}(F^{Mean}_N)$ and the second inequality is obtained by replacing $(F^{CMSS}_{p;N})^{-1}(y)$ with $(F^{CMSS}_{p;N+1})^{-1}(y)$.

---

[17] Let $x \in (x_{p;N}, x_{p;N+1})$. Then, $\psi'_p(x) = y_{p;N} - F^{Mean}_{N+1}(x)$ which implies that ψ_p is concave in x. By continuity, $\lim_{x \to x_{p;N}} \psi_p(x) = \psi_p(x_{p;N})$ and $\lim_{x \to x_{p;N+1}} \psi_p(x) = \psi(x_{p;N+1})$. Concavity implies quasi-concavity, and thus $\psi_p(x) \geqslant \min\{\psi(x_{p;N}), \psi(x_{p;N+1})\}$.

[18] Let $x \in (x_{p;N+1}, x_{p;N})$. Then, $\psi'_p(x) = F^{Mean}_N(x) - y_{p;N+1} < 0$ since $y_{p;N+1} \geqslant y_{p;N} = F^{Mean}_N(x_{p;N})$. By continuity, $\lim_{x \to x_{p;N+1}} \psi_p(x) = \psi(x_{p;N+1})$ such that $\psi_p(x) \geqslant \psi_p(x_{p;N}) \geqslant 0$.

If $(F_N^{Mean})^{-1}(y) \geqslant (F_N^{Mean})^{-1}(y)$, then

$$\begin{aligned}\int_0^1 F_{p;N}^{CMSS}(z;y)dz - \int_0^1 F_{p;N+1}^{CMSS}(z;y)dz &= \int_0^{(F_{p;N+1}^{CMSS})^{-1}(y)} ((F_{N+1}^{Mean})(z) - (F_N^{Mean})(z))dz \\ &\quad + \int_{(F_{p;N+1}^{CMSS})^{-1}(y)}^{(F_{p;N}^{CMSS})^{-1}(y)} ((F_{N+1}^{Mean})(z) - y))dz + 0 \\ &\leqslant \int_0^{(F_{p;N}^{CMSS})^{-1}(y)} ((F_{N+1}^{Mean})(z) - (F_N^{Mean})(z))dz \\ &\leqslant 0\end{aligned}$$

where the inequality is obtained by replacing y with $(F_N^{Mean})^{-1}(z)$ and $(F_N^{Mean})^{-1}(z) \leqslant y$ for all $z \leqslant (F_N^{Mean})^{-1}(y)$.

In any case, we obtain $\int_0^1 F_{p;N}^{CMSS}(z;y)dz \geqslant \int_0^1 F_{p;N+1}^{CMSS}(z;y)dz$ for any y. Since $\int_0^1 F_{p;N}^{CMSS}(z;y)dz$ and $\int_0^1 F_{p;N+1}^{CMSS}(z;y)dz$ are increasing in y as well as

$$\int_0^1 F_{p;N+1}^{CMSS}(z;y_{p;N+1})dz = \int_0^1 F_{p;N}^{CMSS}(z;y_{p;N})dz,$$

this implies $y_{p;N+1} \geqslant y_{p;N}$. □

Lemma 3.8. *Suppose that the selection probability $p \in [p_{F;N+1}, p_{F;N})$ solves the sender's $(N+1)$-dimensional selection persuasion problem. Then, the value of persuasion in the $(N+1)$- dimensional problem is bounded above by the value of selecting $p_{F;N}$ in the N-dimensional problem. That is,*

$$p \cdot \max_{\sigma \in MPC(F_{p;N+1}^{CMSS})} \int U^S(q)d\sigma(q) \leqslant p_{F;N} \cdot \max_{\sigma \in MPC(F_{p_{F;N};N}^{CMSS})} \int U^S(x)d\sigma(x).$$

Proof. Suppose $V_{N+1}^* = p \max_{\sigma \in \mathrm{MPC}(F_{p;N+1}^{CMSS})} \int U^S(q)d\sigma(q)$ for some $p \in [p_{F;(N+1)}, p_{F;N})$. Note that $\sigma \in \mathrm{MPC}(F_{p;N+1}^{CMSS})$ implies $\int x d\sigma(x) = q_F/p$ such that

$$\begin{aligned}V_{N+1}^* &\leqslant p \cdot \max_{\sigma: \int x d\sigma(x) = q_F/p} \int U^S(q)d\sigma(q) \\ &= p\bar{co}(U^S)(\frac{q_F}{p}) \\ &= p\left[\frac{q_F/p - q_2}{q_1 - q_2}U^S(q_1) + \frac{q_1 - q_F/p}{q_1 - q_2}U^S(q_2)\right] \\ &= \frac{1}{q_1 - q_2}\left[q_F(U^S(q_1) - U^S(q_2)) + p(q_1 U^S(q_2) - q_2 U^S(q_1))\right]\end{aligned}$$

where $q_1 \geqslant q_F/p \geqslant q_2$ and the existence of q_1, q_2 is implied by Carathéodory's theorem.

Optimality of p implies that $q_1 U^S(q_2) \geqslant q_2 U^S(q_1)$. Too see this, suppose for a contradiction that $q_2 U^S(q_1) > q_1 U^S(q_2)$. Under this assumption, we obtain

$$\begin{aligned} V^*_{N+1} &< \frac{1}{q_1 - q_2} q_F(U^S(q_1) - U^S(q_2)) \\ &< \frac{q_F}{q_1 - q_2}(U^S(q_1) - \frac{q_2}{q_1}U^S(q_1)) \\ &= \frac{q_F}{q_1}U^S(q_1). \end{aligned}$$

Thus, the sender's optimal payoff is smaller than $\frac{q_F}{q_1}U^S(q_1)$. However, this is a contradiction since the sender can induce some payoff $U^* \geqslant \frac{q_F}{q_1}U^S(q_1)$ by choosing (p, σ) where $p = \max\{p_{F;N+1}, \frac{q_F}{q_1}\}$ and

$$\sigma = \begin{cases} \delta_{q_1} & \text{if } p = q_F/q_1 \\ \frac{q_F/p_{F;N+1}}{q_1}\delta_{q_1} + (1 - \frac{q_F/p_{F;N+1}}{q_1})\delta_0 & \text{if } p = p_{F;N+1} \end{cases}.$$

Now, note that the full information conditional distribution $F^{CMSS}_{p_{F;N};N}$ has support on $\{0, 1\}$. Therefore, $\sigma \in \mathrm{MPC}(F^{CMSS}_{p_{F;N};N})$ if and only if σ has the same expected value as $F^{CMSS}_{p_{F;N};N}$ (Elton and Hill, 1992). That is, $\int x d\sigma(x) = \frac{q_F}{p_{F;B}}$. Thus, the concavification result of Kamenica and Gentzkow (2011) implies that

$$V^*_N(p_{F;N}) := \max_{\sigma \in \mathrm{MPC}(F^{CMSS}_{p_{F;N};N})} \int U^S(x)d\sigma(x) = \bar{co}(U^S)(\frac{q_F}{p_{F;N}}).$$

Therefore, we can bound the value of selecting $p_{F;N}$ in the N-dimensional problem, $p_{F;N}V^*_N(p_{F;N})$, as follows:

$$\begin{aligned} p_{F;N}V^*_N(p_{F;N}) &= p_{F;N}\bar{co}(U^S)(\frac{q_F}{p_{F;N}}) \\ &\geqslant p_{F;N}\left[\frac{q_F/p_{F;N} - q_2}{q_1 - q_2}U^S(q_1) + \frac{q_1 - q_F/p_{F;N}}{q_1 - q_2}U^S(q_2)\right] \\ &= \frac{1}{q_1 - q_2}\left[q_F(U^S(q_1) - U^S(q_2)) + p_{F;N}(q_1U^S(q_2) - q_2U^S(q_1))\right] \\ &\geqslant \frac{1}{q_1 - q_2}\left[q_F(U^S(q_1) - U^S(q_2)) + p(q_1U^S(q_2) - q_2U^S(q_1))\right] \\ &\geqslant V^*_{N+1} \end{aligned}$$

where the first inequality stems from concavity of $\bar{co}(\cdot)$, the second inequality immediately follows from $p_{F;N} \geqslant p$ and $q_1U^S(q_2) \geqslant q_2U^S(q_1)$ and the final inequality follows from the bound on V^*_{N+1} derived above. □

3.A.7 Proof of Proposition 3.7

Proof. Recall that by Proposition 3.4, $(p^*, F_{p^*}^{CMSS})$ is optimal where $p^* = F^{Mean}(x^*) + \sum_{x>x^*} f^{Mean}(x)x$, and $x^* = \lfloor N\bar{x} \rfloor / N$ where $\bar{x} = \min\{x \in [0,1] : U^S(x) \leqslant xU^S(1)\}$.

Let $\psi(x) := \frac{1}{d}(U^S(x) - xU^S(1))$ such that $\bar{x} = \min\{x \in [0,1] : \psi(x) \leqslant 0\}$. First, note that

$$\psi(x) = \frac{c}{d} + x^{\frac{1}{1-\gamma}} - x(\frac{c}{d} + 1) = (1-x)\frac{c}{d} + x^{\frac{1}{1-\gamma}} - x,$$

which is non-decreasing in c/d for any $x \in [0,1]$ which immediately implies that $\bar{x}$ is non-decreasing in c/d.

(i) Suppose that $\gamma/(1-\gamma) < c/d$. Then,

$$\begin{aligned}
\psi'(x) =& \frac{1}{1-\gamma} x^{\frac{\gamma}{1-\gamma}} - (\frac{c}{d} + 1) \\
\leqslant& \frac{1}{1-\gamma} - (\frac{c}{d} + 1) \\
=& \frac{\gamma}{1-\gamma} - \frac{c}{d} \\
<& 0
\end{aligned}$$

which implies $\psi(x) > \psi(1) = 0$ for all $x < 1$. Thus, $x^* = \bar{x} = 1$, which implies $p^* = F^{Mean}(1) = 1$.

(ii) Next, suppose that $c/d < \frac{1-(\frac{1}{N})^{\gamma/(1-\gamma)}}{N-1}$. We want to show that $\psi(x) < 0$ for all $x \geqslant \frac{1}{N}$ which implies $\bar{x} < 1/N$ and consequently $x^* = 0$. Note that $\psi(1) = 0$ and since ψ is convex, $\psi(x) < 0$ for all $x \in [\frac{1}{N}, 1]$ if $\psi(\frac{1}{N}) < 0$. It is straightforward to verify that $\psi(\frac{1}{N}) < 0$ if and only if $\frac{c}{d} < \frac{1-(\frac{1}{N})^{\gamma/(1-\gamma)}}{N-1}$. Thus, $x^* = 0$ which implies $p^* = F^{Mean}(0) + \sum_{x>x^*} f^{Mean}(x)x = f^{Mean}(0) + q_F = p_F$. □

3.A.8 Proof of Proposition 3.8

Proof. Suppose the non-triviality condition $t > q_F$ is satisfied. Let the value of the p-action persuasion problem be denoted by

$$V(p) = \max_{\sigma \in \mathrm{MPC}(F_p^{CMSS})} \int U^S(x) d\sigma(x)$$

where $U^S(x) = \mathbf{1}_{x \geqslant t}$. Note that

$$V(p) \leqslant \bar{co}(U^S)(\frac{q_F}{p}) = \begin{cases} \frac{1}{t}\frac{q_F}{p} & \text{if } q_F/p \leqslant t \\ 1 & \text{if } q_F/p \geqslant t \end{cases}.$$

Therefore, we can bound the sender's persuasion value V^* from above as follows:

$$V^* = \max_{p \in [p_F, 1]} pV(p) \leqslant \max_{p \in [p_F, 1]} \begin{cases} \frac{1}{t} q_F & \text{if } q_F/t \leqslant p \\ p & \text{if } q_F/t \geqslant p \end{cases}$$
$$= \frac{q_F}{t}.$$

(i) Suppose that $t \leqslant \frac{q_F}{p_F}$. It is straightforward to verify that $(\frac{q_F}{t}, \delta_t)$ is feasible if $t \leqslant \frac{q_F}{p_F}$ and induces an expected payoff of $\frac{q_F}{t}$ which proves (i).

(ii) Suppose that $\frac{q_F}{p_F} \leqslant t$. Note that $(p_F, (1-\alpha)\delta_0 + \alpha\delta_t)$ is feasible if and only if $\alpha \in [0, 1]$ which is equivalent to $t \geqslant \frac{q_F}{p_F}$, and $(p_F, (1-\alpha)\delta_0 + \alpha\delta_t)$ achieves an expected payoff of $p_F\alpha = \frac{q_F}{t}$. □

References

Mark Armstrong and Jidong Zhou. Consumer information and the limits to competition. *American Economic Review*, 112(2):534–77, 2022. 3.8

Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated games with incomplete information.* MIT press, 1995. 3.8

Piotr Dworczak and Giorgio Martini. The simple economics of optimal persuasion. *Journal of Political Economy*, 127(5):1993–2048, 2019. 3.1, 3.3.4

J Elton and Theodore P Hill. Fusions of a probability distribution. *The Annals of Probability*, pages 421–454, 1992. 3.2, 3.A.3, 3.6, 3.A.6

Matthew Gentzkow and Emir Kamenica. A rothschild-stiglitz approach to bayesian persuasion. *American Economic Review*, 106(5):597–601, 2016. 3.1, 3.3.4

Shota Ichihashi. Online privacy and information disclosure by consumers. *American Economic Review*, 110(2):569–95, 2020. 3.8

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011. 4, 3.8, 3.A.6

Andreas Kleiner, Benny Moldovanu, and Philipp Strack. Extreme points and majorization: Economic applications. *Econometrica*, 89(4):1557–1593, 2021. 3.1, 3.3.4, 3.8

Anton Kolotilin. Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 13(2):607–635, 2018. 3.1, 3.3.4

Anton Kolotilin, Timofiy Mylovanov, and Andriy Zapechelnyuk. Censorship as optimal persuasion. *Theoretical Economics*, 17(2):561–585, 2022. 3.1

Changhwa Lee. Optimal recommender system design. *Unpublished manuscript, University of Pennsylvania*, 2021. 3.8

Alfred Müller and Dietrich Stoyan. *Comparison methods for stochastic models and risks*, volume 389. Wiley, 2002. 3.5, 3.A.6

Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of political Economy*, 118(5):949–987, 2010. 3.8

Moshe Shaked and J George Shanthikumar. *Stochastic orders.* Springer, 2007. 3.2

Volker Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, 1965. 3.2, 3.8

Wataru Tamura. Bayesian persuasion with quadratic preferences. *Available at SSRN 1987877*, 2018. 3.8

Kai Hao Yang and Alexander K Zentefis. Distributions of posterior quantiles and economic applications. 2022. 3.8