2019

Three Essays on Market Strategy with Incomplete Information

Qichao Shi
Lehigh University, qichshi@gmail.com

Three Essays on Market Strategy with Incomplete Information

by

Qichao Shi

Presented to the Graduate and Research Committee
of Lehigh University
in Candidacy for the Degree of
Doctor of Philosophy

in

Business and Economics

Lehigh University

July 15, 2019

.

LEHIGH UNIVERSITY

College of Business and Economics

Approved and recommended for acceptance as a dissertation in partial fulfillment of the requirements of the degree of Doctor of Philosophy.

____
Candidate Name

____
Accepted Date

Committee Member Signatures:

____
Committee Chair: **James Dearden**

____
Committee Member: **Ernest Lai**

____
Committee Member: **Daisy Dai**

____
Committee Member: **David Goldbaum**

Acknowledgments

At this moment, as the last mile toward a Ph.D. is just under my feet, I look back to the long way which I have walked through after missing several crossings, knowing that I would never arrive here without the guidance and support from many individuals. I wish to express my gratitude to them. Let me go backward along this road.

First and foremost, I am especially indebted to my advisor James Dearden and my dissertation committee members Ernest Lai, Daisy Dai and David Goldbaum, who guided and encouraged me to accomplish this dissertation. From the day I started, they continuously provided me with comments and suggestions which relighted my life every time after I talked with them. I also learned a lot from James Dearden about game theory and industrial organization, from Ernest Lai about microeconomics and game theory, and from Daisy Dai about empirical industrial organization. Before taking their courses, I had never systematically studied game theory and industrial organization, which were so interesting to me that they have become my research focus for the rest many years of my life.

I am also grateful to other professors and administrators at the Department of Economics in Lehigh University. In particular, I thank program directors Chad Meyerhoefer and Muzhe Yang for their many helps to keep me on the right track, and I learned most econometrics from them as well. I thank Shin-yi Chou for the advice related to job market, and although the health economics she taught was not my research field, I learned a lot inspiring studies with useful techniques to read research papers. I thank Robert Thornton and Thomas Hyclak for the study of labor economics, and my third year paper was finished under the guidance of Thomas Hyclak. In several years, I was working as a teaching assistant for Irina Panovska, from whom I learned how to organize a course and make a

semester teaching smoothly. Her recommendation for my job market also indispensable. In addition, I really appreciate all other helps from my classmates, colleagues and friends at Lehigh University here or there during these years.

Backward along the PhD trip, I am indebted to my master advisor Inmee Baek who encouraged and guided me when I decided to apply for a PhD program. Following her advice, I worked on a reach paper and attended an academic conference, which not just strengthened my application but also was a valuable experience to lead me into the world academics.

At the beginning of this trip, my family was a solid support. I am extremely grateful to my parents who gave me complete freedom to do whatever I want. They must not image that their only son is still studying in school when they are already above 65. When I changed my major from medicine to economics, they did not stop me; when I decided to leave them behind and study 7000 miles away, they still did not stop me. All they have told me was encouragement. Finally, I am deeply indebted to my wife Shuangshuang Wang, who came to the U.S. earlier than me so that I could easily skip the most struggling days as living in a foreign country. Although she was also pursuing her own PhD during the same yeas, she still took most family responsibilities so that I was able to focus on my research. I would need a much longer time to accomplish this dissertation if I lived alone without her. Her love, support and understanding were everywhere along this road from the very beginning to all the end.

Abstract

The development of information technology is reshaping the market in which economic agents make decisions based on incomplete information which in the past may reduce market efficiency but nowadays is no longer a big problem when more information could be acquired. How would this change affect market behavior is still under explored. This dissertation explores two distinct topics that are related to information disclosure: third-party product rankings and return policies. Specifically, the first essay investigates the market strategy of a product expert who sells the product ranking to incompletely informed consumers. The model indicates that the expert may not always have incentive to rank products consistent with consumer preferences, especially when the ranking could influence consumer utility. We find evidence for this argument from a laboratory experiment. This type of third-party product ranking may influence not just consumers but also firms whose products have been reviewed. Under this circumstance, I model firms' advertising strategies in the second essay, and find that when their product quality has been disclosed by the third-party firms may rely more on persuasive advertising of certain attribute to influence consumers' preference. Essay 3 examines a retailer's optimal product return policy within an environment in which consumers differ in their product return propensities. The retailer may benefit from a discriminated return policy based on consumers' differential return behaviors which can be revealed by their purchase history.

Contents

List of Tables

List of Figures

Chapter 1: An Experimental Study of Selling Expert Advice

Abstract

This study explores the interaction between a product expert, who offers to sell a product ranking, and an incompletely informed consumer. The consumer considers acquiring the expert's product ranking not only because the expert has superior information about the quality of the products the consumer is considering and knows the consumer's utility function, but also because the expert can directly influence consumer utility of a product by the product's rank. There are multiple equilibria in this setting with strategic information transmission: ones in which the expert ranks products in a manner that is consistent with the consumer's pre-ranking utilities, which depend exclusively on the products themselves, and ones in which the expert does not. We design a laboratory experiment to investigate which equilibrium an expert and consumer play. Across the three treatments we examine, which vary by the consumer's possible pre-ranking utilities, we find evidence that product experts are likely to select a ranking methodology that involves considerable uncertainty about the final product ranking, even though doing so involves ranking products in a manner that is inconsistent with consumer pre-ranking utilities.

1.1 Introduction

Consumers prior to making purchase decisions frequently seek expert advice, which sometimes takes the form of product rankings. Familiar scenarios include students acquiring college and university ranking publications (e.g., *U.S. News & World Report Best Colleges Guidebook* and *The Wall Street Journal/Times Higher Education College Rankings*), car buyers acquiring auto rankings (e.g., *Kelley Blue Book New Car Buyer's Guide* and *Consumer Reports New Car Buying Guide*), and home cooks acquiring kitchen product rankings (e.g., *Cook's Illustrated*). A less familiar case is consumers acquiring a retailer's product ranking (e.g., TireRack.com's tire rankings).

These publications are experts because they measure, or collect information about, product attributes and they have information about consumer utility regarding those product attributes (e.g., the weights in their utility functions they attach to the various product attributes). In an extreme, an expert is completely informed about product attributes and knows the utility function, which maps product attributes and the ranking into product utilities, of the consumer seeking advice. Consumers are novices because while they may know their utility functions, they are incompletely informed about the qualities of the product attributes. Therefore, a completely informed expert has the opportunity to offer a consumer good advice in the sense of ranking products in a manner that reflects her product utilities.

Many rankings publishers are for-profit institutions, including some of the above, and the objective of a for-profit publisher is not necessarily to offer good advice to consumers by ranking products in a manner that is consistent with consumer utility of the products being considered, but rather it is to maximize sales of its ranking (or sales of detailed product information linked to the ranking), sales of its general publication, or sales of advertising associated with the ranking platform. We ask the question whether a completely-informed for-profit expert would rank products in a manner that is inconsistent with consumer utility of the products she is considering.

Popular press writers suggest that in the university ranking arena the answer to this question is no. Commercial college and university ranking publications receive frequent criticism about the methodologies they use to develop their rankings. In fact, the sentiment about university rankings in general and *U.S. News & World Report* in particular is so strong that an entire Wikipedia page is devoted solely to the criticism of college and university rankings.[1] One common criticism about such university rankings is that they "tinker" with their methods from year to year. Tierney (2013) in The Atlantic writes:

> U.S. News is always tinkering with the metrics they use, so meaningful comparisons from one year to the next are hard to make. Critics also allege that this is as much a marketing move as an attempt to improve the quality of the rankings: changes in the metrics yield slight changes in the rank orders, which induces people to buy the latest rankings to see what's changed.[2]

As Tierney (2013) suggests, these changes may have more to do with the marketing of these publications than with improving the quality of the ranking.

In the context of student decision making about universities to attend, Dearden, Grewal and Lilien (2019, forthcoming) evaluate whether a for-profit expert ranks universities (i.e., products) in a manner that is consistent with student (i.e., consumer) utilities. In doing so, they partition student utility into two components: the utility of the quality of a

[1]http://en.wikipedia.org/wiki/Criticism_of_college_and_university_rankings_(North_America); accessed October 2018.

[2]https://www.theatlantic.com/education/archive/2013/09/your-annual-reminder-to-ignore-the-em-us-news-world-report-em-college-rankings/279103/; accessed October 2018.

university's attributes (which are set by the university); and the utility of the university's rank (which is extrinsic to the university and determined by the expert). Their analysis demonstrates that the effect of a university' rank on a student's utility (e.g., a prestige effect of a university ranking) creates the incentive for a for-profit ranking publication to add uncertainty to its ranking methodology beyond that created by, when using the student's utility of university attributes to aggregate university attribute scores into overall scores, the randomness of university attribute scores. If the students evaluating universities are concerned with only the rank, and not with the quality of attributes, of the universities they attend, and therefore not with the information about university attributes provided (probabilistically) by the ranking, then the expert selects a ranking methodology where each possible ranking of the universities is equally likely – a uniform distribution over all possible rankings. However, if the students are concerned with only the quality of the attributes and accordingly with the information provided (probabilistically) by the ranking, and the university's rank does not affect their utilities of attending, then the expert ranks universities in a manner that is consistent with the student's utilities of the university attributes by using the students' utility function (which maps university attributes into utilities) as a ranking methodology.

The Dearden, Grewal and Lilien (2019, forthcoming) model has multiple equilibria in which the expert's ranking methodology, and the degree to which the expert adds uncertainty to the ranking outcome, varies across the equilibria. They examine only the equilibrium methodology that maximizes the expert's profit. Due to the influence of the expert's ranking on the students' utilities, this ranking that maximizes profit may not rank universities in a manner that is consistent with student utility. Hence, the question remains about whether an expert and student (more generally, consumer) play this equilibrium. On one hand, the expert selects this methodology because it provides consumers with the incentive to acquire the ranking. On the other hand, when the ranking influences the consumer's utility, then ranking may boost the utility of a product that, according to product attributes, is not the consumers' favorite. In this sense, the expert harms the consumers. If the expert is concerned about the consumers, then he may shy away from using the methodology that maximizes his profit.

In this paper, we experimentally investigate the key insight of Dearden, Grewal and Lilien (2019, forthcoming) by addressing two research questions. First, we explore under what circumstances the theoretical prediction that a product expert ranks products in a manner that is inconsistent with consumer utility of product attributes. Second, we examine whether the expert's willingness to use a ranking methodology that may not rank products according to consumer utility of product attributes varies with the extent to which the expert in maximizing profit must harm the consumer.

The results of our experiment suggest that product experts have a strong tendency to rank products in a manner that is inconsistent with a consumer's pre-ranking utilities of the products being considered. The expert does so to incent the consumer to acquire the expert's advice. Furthermore, by varying the possible pre-ranking utilities of the product across three treatments of our experiment, we find that the expert's selected product ranking methodology in the three treatments changes in a way that suggests competing incentives of the expert in selecting a ranking methodology.

The remainder of the paper proceeds as follows. Section 1.2 provides an overview of our analysis and surveys the related studies. Section 1.3 presents our experimental game of selling expert advice, followed by the experimental design in Section 1.4. We then report our experimental findings in Section 1.5, and summarize possible discussions in Section 1.6. Finally, Section 1.7 concludes.

1.2 Overview and Literature

To provide an overview of the strategic issues involved in an expert selling advice, in this section we develop consumer utility in a way that is relevant to the construction of our model, present related literature, and introduce a popular ranking methodology and how it relates to our model.

1.2.1 Consumer Utility

A product ranking not only affects a consumer's expected utility of selecting a product through the information it provides about the product, but also it influences the consumer's utility of consuming a product by the prestige or future economic value it may create for a product according to the product's position in the ranking.

Considering the prestige or economic effect of an expert's product ranking, a product's rank interacts in a consumer's utility function with the consumption of that product. With a bit of detail, the utility of consuming a product could be modeled a function of the quality of the product's attributes (which are designed by a product's manufacturer) and its rank in an expert's product ranking. In our model with one expert and one consumer, we refer to the consumer's utility of a product's attributes as the *pre-ranking utility* and the combined utility of these attributes and the product's rank in the expert's ranking as the *post-ranking utility.* The post-ranking utility less the pre-ranking utility is the *ranking utility.*

In our model, the quality of product attributes set by the product-selling firms is not modeled here; this pre-ranking utility is therefore exogenously given in our game. By contrast, the interaction between product expert's ranking and a consumer's utility – the ranking utility – is endogenously determined by the position the expert places the product in its ranking; the ranking utility is decreasing in the product's ranking position (with number 1 being the top-rank).

This added utility associated with a product's rank comes from the expert's power to create prestige or economic value for highly-ranked products. We interpret this power in two ways. First, a good product rank could provide consumers with future payoffs. Take the automobile market as one example. If a car is top-ranked in a popular auto ranking, its owner can resell it at a higher price. In another example, job recruiters may use a university's rank as salient information about quality of university graduates. If so, the best firms may visit the top-ranked universities, which benefits the students who attend them. That is, students and recruiters use a university ranking as a coordination device. Those potential returns result from the reputation generated by the product ranking. Second, a high ranking may contain a psychological benefit to consumers who feel better when the selected product is highly recommended. This is analogous to persuasive advertising that could influence consumers' preferences. In the context of university rankings, the ranking value is a prestige effect in Dearden, Grewal and Lilien (2019, forthcoming), meaning an increase in consumer utility when the products they choose improve their ranks. We adopt the same idea here.

Luca and Smith (2013) demonstrate that the actual college and university ranks in *U.S. News & World Report*, controlling for school attributes, influence application decisions.

Using the Chetty, Looney and Kroft (2009) definition of the salience of information as the simplicity of calculating the information, Luca and Smith (2013) suggest that this influence of the *U.S. News & World Report* ranking on student decisions is due to the salience of the information provided by a simple rank-order of colleges and universities. We contend that this effect of university ranks on student decisions could be due to this popular ranking's prestige conferred on colleges and universities, in which a student's utility of attending a college or university is decreasing in its position in the ranking (position 1 is best).

A product ranking has the opportunity to inform consumers about the attributes of the products being considered because either a consumer updates the distribution of product attributes values using Bayes' rule based on ranks (because higher-ranked products tend to have higher-quality attributes) or a consumer uses product ranks as salient information about the quality of th product's attributes. The case of Bayesian updating may apply to important purchases that significantly affect utility. A consumer may pay attention to an array of product attributes and, if the consumer acquires a ranking, she may carefully update expected attribute scores based on product ranks. The salience case may apply to smaller purchases of unfamiliar products. If a consumer views a ranking, she may not be as careful in contemplating attribute scores, and rather she may use the ranking as salient information about product quality (Chetty, Looney and Kroft, 2009). Whether a consumer rationally updates probabilistic beliefs about the quality of a product's attributes based on its rank or uses its rank as salient information about this quality, a product ranking affects a consumer's expected expected utilities of the considered products.

As product rankings are likely to be consumed in a manner similar to advertising (or any marketing communication), the literature on advertising is relevant. Product rankings provide information about products, and create ranking value for highly-ranked ones, similar to how advertisements provide information and create status or prestige. Those two perspectives, information and prestige, parallel those that dominate research in advertising, the first of which views advertising as an information source while the second views advertising as a persuasion tool (Bagwell, 2007). Hence, from an information perspective, similar to advertising in general, rankings inform consumers about product attributes (Butters, 1977; Nelson, 1974; Stigler, 1961). From a persuasion perspective, rankings influence the utilities of the consumers who select products (Ackerberg, 2001; Becker and Murphy, 1993; Galbraith, 1976).

1.2.2 The Expert's Ranking Methodology

In our model with one expert and one consumer, the expert's decision whether to offer good advice to the consumer by ranking products according to the consumer's pre-ranking product utilities involves its selection of the methodology it uses to rank products.

One popular ranking process – one used by all of the above institutions – to rank products is an attribute-and-aggregate process in which an expert identifies product attributes of importance, for each product measures its attribute scores or collects attribute score data, chooses weights for the attributes, for each product multiplies the weights and the attribute scores, and determines the aggregated scores. The publication then rank-orders products based on these aggregated scores to form a ranking of products. Each of the above steps in the ranking process can affect whether the expert ranks products in a manner that is consistent with the consumer's utilities of product attributes. Whether the expert does so is the focus of our research.

Our research simplifies the attribute-and-aggregate ranking process. The expert in our model knows the consumer's utility function. Within the process we examine, the expert's ranking methodology is a mapping from the consumer's pre-ranking utilities of the products being considered to ranks. To be consistent with practice in which rankings publications offer only overviews of their methodologies, we assume that the expert does not report its methodology to the consumer. After having selected its ranking methodology, the expert learns the consumer's pre-ranking utilities and ranks products according to its selected methodology. The consumer then decides whether to acquire the expert's ranking, either by expending the time cost required to do so or paying the monetary cost. If the consumer acquires the product ranking, then she learns the products' ranks prior to her product selection. If she does not acquire the ranking, then she learns the ranks only after she has selected a product. However, whether or not she acquires the ranking, it affects her utility of consuming her selected product.

1.2.3 Strategic Information Transmission

Our model involves strategic information transmission in which an expert (sender) sends a message in the form of an ordinal product ranking to a consumer (receiver). Unlike the canonical problem (see Crawford and Sobel, 1982; Argenziano, Severinov and Squintani, 2016) in our model: (i) the expert's message is an argument in the consumer's utility function, (ii) the expert is unconcerned with product selected by the consumer, and (iii) the consumer decides whether to acquire the expert's ranking. Furthermore in our model (iv) the expert selects a ranking methodology when he and the consumer hold the same prior probabilistic beliefs about the attributes of the products being ranked (and therefore the consumer's pre-ranking utilities of the products). The expert commits to using the methodology, which in terms of the strategic information transmission models is a message function from the space of possible pre-ranking product values to the set of possible ordinal product rankings. Also unlike the canonical model, our model (v) involves commitment because the expert selects the ranking methodology prior to learning the pre-ranking values, but ranks the products after having learned them.

A growing number of experiments have been conducted on strategic information transmission (see Crawford (1998) and Blume, Lai and Lim (2017) for surveys of experimental studies). Among those experimental studies, our research is similar to the setting with imperfect incentive alignment where Blume, Lai and Lim (2017) document that systematic over-communication is commonly evidenced in experiments with a single sender and a single receiver (e.g. Cai and Wang, 2006).[3] However, Blume, Lai and Lim (2017) argue that the evidence of equilibrium selection is limited.

This paper also relates to a special case of strategic information transmission named "persuasion," which is studied by Kamenica and Gentzkow (2011), Gentzkow and Kamenica (2014), and Hörner and Skrzypacz (2016). In these models, a sender strategically discloses his information to the receiver, aiming to influence the receiver's action. The studies characterize sender's decision whether to disclose complete information or only partial information (e.g. Rayo and Segal, 2010; Ostrovsky and Schwarz, 2010).[4]

This paper contributes to the literature of expert product rankings in the following

[3]The fully revealing equilibrium occurs more often in multiple-sender multidimensional cheap talk that is investigated by Lai, Lim and Wang (2015).

[4]When reviewing theoretical and empirical studies on quality disclosure, Dranove and Jin (2010) argue that even third-party certification can provide inaccurate and biased information due to either the data generating process or conflict of interest.

ways. First, it constructs a simple model to describe the strategic environment of an expert selling advice, which has implications that might be applied to many product markets. Second, we emphasize the expert's strategic behavior when he has the power to influence consumer utility, possible through affecting market values of the recommended products. This case in which the sender's (expert's) message (ranking) is an argument in the receiver's (consumer's) utility function adds a new angle to the literature on strategic information transmission.

1.2.4 Optimal Ranking Methodology and Randomness

Just as an expert may change its methodology from year to year in a dynamic setting, in our static model, the expert may select a methodology that introduces uncertainty into its product ranking. In our model, the expert generates the consumer's interest in his ranking by creating uncertainty about the top-ranked product (of the two in our model) and ultimately about the product that receives the expert's ranking value boost.

Dearden, Grewal and Lilien (2019, forthcoming) identify the ranking value created by the expert, and not the information provided by the expert, as the source of the expert adding uncertainty to its product ranking. Simply put, if the consumer learns the product ranking, then she selects the top-ranked product. However, if the consumer does not acquire the product ranking to learn the products' ranks, then she may make the *ex post* mistake of not selecting the top-ranked product. In adding uncertainty to the product ranking, the expert increases the probability that a consumer who does not acquire the ranking makes the mistake of not selecting the top-ranked product.

Similar to the way a product expert can create value for highly-ranked products, an influential fashion magazine can create prestige for the people who wear the items it singles as being stylish. As Kuksov and Wang (2013) suggest, fashion editors are "the single most important influencer of fashion" (p. 53). The authors note that since fashion editors carry great weight in determining fashion hits, they are key players in generating a fundamental property of the fashion marketplace, namely the seemingly random nature of the determination of a season's "it" (Kuksov and Wang, 2013). Fashion editors include that randomness to appeal to fashion-conscious consumers who are interested in wearing the "it" products, whether for their intrinsic pleasure or for the benefits associated with signaling their fashion senses.

In offering product recommendations, deception is a possible issue. One interesting point about the Kuksov and Wang (2013) analysis is that a fashion magazine influences consumer utility by its selection of "it" items without deceiving consumer because the designation does not indicate that the top product is better than the others. In designing our experiment, because an expert's product ranking methodology to motivate the consumer to acquire his product ranking might involve deception and lying, we want to measure the effect of the expert's influence on the consumer's utility – the ranking value – without muddying this influence by having the possibility of deception. In Section 1.6, we use the Sobel (2018) definition of deception in models of strategic information transmission to contend that our model does not involve deception.

1.3 The Experimental Game: Selling Product Advice

1.3.1 The Strategic Environment

There are two players, a product expert (he) and a consumer (she), and two products, A and B. The expert sells product advice to the consumer in the form of a report ranking Products A and B. The consumer, who is imperfectly informed about the values of the products, decides whether to acquire the expert's ranking report and which product to purchase.[5]

The value of a product to the consumer comprises two components, a *pre-ranking value* and a *ranking value*. The pre-ranking value, υ, refers to the value the consumer attaches to the product based on its quality and price, variables that are set by the product-selling firm not modeled here; this component of the value is therefore exogenously given in our game. By contrast, the ranking value of a product is endogenously determined by the expert; the consumer attaches an additional value, γ, to the product that is ranked first by the expert.

The pre-ranking value of Product A, $\bar{\upsilon}_A > 0$, is commonly known. The consumer is uncertain about the pre-ranking value of Product B, υ_B, which takes one of two values, 0 and $\bar{\upsilon}_B > 0$. The expert learns the realized value of υ_B, but he makes his only move in the game – commits to a ranking methodology – before observing the realization. The common prior is that $\upsilon_B = \bar{\upsilon}_B$ with probability $0 < p < 1$. The ranking value γ is also a common knowledge.

A *ranking methodology* is a mapping, $r : \{0, \bar{\upsilon}_B\} \to \{A, B\}$; it specifies for each possible value of υ_B a *ranking report* for the consumer, where report A (B) indicates that

[5] Because the utilities in our experiment are monetary rewards, we refer to utility as value.

Product A (B) is ranked first. Alternatively, the reports can be interpreted as non-binding recommendations to the consumer on which product to purchase.

We restrict attention to two ranking methodologies, an uninformative and an informative. The uninformative methodology, r_A, always ranks Product A first, i.e., $r_A(0) = r_A(\bar{\upsilon}_B) = A$, thus providing no information regarding the uncertain pre-ranking value of Product B. On the other hand, the informative methodology, $r_{A||B}$, which fully reveals the two possible pre-ranking values of Product B, generates ranking reports in the following manner:

$$r_{A||B}(\upsilon_B) = \begin{cases} A & \text{if } \upsilon_B = 0, \\ B & \text{if } \upsilon_B = \bar{\upsilon}_B. \end{cases}$$

The set of pure strategies of the expert is $\{r_A, r_{A||B}\}$.[6] After choosing a pure strategy, which is tantamount to committing to a ranking methodology, the expert observes the realized value of υ_B and offers a ranking report according to the adopted methodology. The use of a mixed strategy, for which we denote by $q_{r_{A||B}} \in [0, 1]$ the probability that the expert chooses $r_{A||B}$, is appropriately interpreted.

The consumer never observes the expert's strategy choice (the adopted ranking methodology). The consumer observes the output of the methodology – the ranking report – before selecting a product if and only if she acquires it at a cost $c > 0$. After making such an acquisition decision and seeing the ranking report, if any, the consumer chooses between Products A and B. A pure strategy of the consumer consists of two components: i) a report-acquisition decision $\alpha \in \{0, 1\}$, where 1 indicates acquiring the ranking report and 0

[6] In this study, we assume the expert might have incentive to misreport the ranking only when the difference between two products is not too big. This assumption is realistic since the expert may consider his reputation. We maintain this by assuming that the expert always ranks product A first if $\upsilon_B = 0$, assuring that a really bad product B will not be ranked first.

otherwise, and ii) a product-purchasing rule $\rho : \{\varnothing, A, B\} \to \{A, B\}$, where $\varnothing$ represents not seeing any report after acquisition decision $\alpha = 0$. For simplicity, we assume that the consumer does not pay for either product.

The expert earns a payoff of $\pi > 0$ if the consumer acquires his ranking report. The expert's choice of ranking methodology does not affect his payoff, and he takes no interest in the product choice of the consumer.[7] The consumer's payoff is the sum of the pre-ranking value and any ranking value of her chosen product, minus any report-acquisition cost. A consumer who chooses product $K \in \{A, B\}$ with pre-ranking value υ_K, makes report-acquisition decision $\alpha \in \{0, 1\}$, and faces an expert offering ranking report $r(\upsilon_B)$ receives a payoff of

$$u(K, \upsilon_K, \alpha, r(\upsilon_B)) = \upsilon_K + \gamma \mathbb{I}_{r(\upsilon_B)}(K) - \alpha c,$$

where $\mathbb{I}_{r(\upsilon_B)}(K)$ is an indicator function taking the value 1 if $K = r(\upsilon_B)$, i.e., Product K is ranked first given the expert's ranking methodology and the realized pre-ranking value of Product B, and 0 otherwise. Note that the consumer receives the ranking value γ for choosing the first-ranked product even if she does not acquire the ranking report ($\alpha = 0$).

The timing of the game is as follows. The expert privately chooses a ranking methodology. The pre-ranking value of Product B is realized and privately observed by the expert, who then issues a ranking report based on the adopted methodology. The consumer decides whether to acquire and see the report, after which she chooses a product.

Before proceeding to the equilibrium analysis, we further comment on the properties

[7] We assume that the expert earns 0 if the consumer does not acquire the ranking report. Notice that the expert's revenue is unaffected by which ranking methodology is used or whether the ranking orders products according to the consumer's pre-ranking values.

of the ranking methodologies. In addition to the dichotomy of informative versus uninformative, the methodologies can also be characterized as whether it is *consumer-optimal*. A consumer-optimal methodology *always* ranks higher the product based on the consumer's pre-ranking utilities. Only one of the ranking methodologies is consumer-optimal; and the one that is so depends on the parameter values.

If $\bar{v}_A \geqslant \bar{v}_B$, Product A is more valuable to the consumer irrespective of the realized pre-ranking value of Product B. Even though the informative $r_{A||B}$ resolves the uncertainty surrounding Product B, it is not consumer-optimal because it sometimes ranks Product B higher. On the other hand, r_A, though not informative, is consumer-optimal. If $\bar{v}_A < \bar{v}_B$, informativeness coincides with consumer-optimality: $r_{A||B}$ would be a consumer-optimal methodology while r_A would not be.

Informativeness and consumer-optimality are properties of the ranking methodologies related to the consumer's pre-ranking value. The ranking value of the first-ranked product also confers a *persuasive* property to the methodologies analogous to the effect of persuasive advertising, in which the provider not only offers information but can also directly influence the consumer's preferences. Our equilibrium analysis, as well as the experimental design, exploits the interplay of these three properties of the game.[8]

[8]Notice that the ranking value is independent of whether the expert's ranking reflects the pre-ranking values. That is, no matter which product is ranked first, the consumer obtains the same utility from the top ranking. This is reasonable from two viewpoints. First, at the point in time when a consumer makes a purchase decision, she evaluates the product values and not the source of the value, whether it is pre-ranking value or ranking value. Therefore, at this point, she is unconcerned with whether the expert had ranked products according to pre-ranking values. Second, sometimes a consumer cannot exactly tell whether the ranking reflects pre-ranking values, even after a transaction. In particular, some experience goods like cars are too complicated to detect whether the ranking reflects pre-ranking values, in part because those products have many attributes.

1.3.2 Equilibrium Analysis

We analyze the perfect Bayesian equilibria of the game, in which the expert best responds to the consumer, the consumer best responds to her beliefs, and beliefs are formed via Bayes' rule whenever possible. Note that the consumer has two types of belief: we denote by $\mu_{r_{A||B}}$ her *strategy belief* that the expert has chosen $r_{A||B}$ and $\mu_{\bar{v}_B}$ her *product belief* that the pre-ranking value of Product B is $\bar{v}_B$.[9]

We begin by characterizing the consumer's product-purchasing rule (all proofs are relegated to Appendix A.1):

Lemma 1. *For any given mixed strategy of the expert* $q_{r_{A||B}} \in [0,1]$*,*

1. *it is a best response for the consumer who has not acquired the ranking report (*$\alpha = 0$*) to purchase Product* A *(*$\rho(\varnothing) = A$*) if* $p\bar{v}_B \leqslant \bar{v}_A + (1-2p)\gamma$*; and*
2. *it is a best response for the consumer who has acquired the ranking report (*$\alpha = 1$*) to purchase the first-ranked product (*$\rho(K) = K$*,* $K = A, B$*) if* $\gamma \geqslant \max\{p\bar{v}_B - \bar{v}_A, \bar{v}_A - \bar{v}_B\}$*.*

Without seeing any report, the consumer's strategy belief $\mu_{r_{A||B}}$, which is relevant for the expected ranking value, coincides with the expert's strategy $q_{r_{A||B}}$, and her product belief $\mu_{\bar{v}_B}$, which is relevant for the expected pre-ranking value of Product B, coincides with the prior p. Her expected payoffs from Products A and B are respectively lower and higher *via the ranking value* when the probability that the expert chooses methodology $r_{A||B}$ is higher. For Product A to be more attractable under a frequent adoption of $r_{A||B}$,

[9] Although the expert's choice of ranking methodology and the consumer's report-acquisition decision are strategically simultaneous, a perfect Bayesian equilibrium explicitly specifies the consumer's strategy beliefs, in addition to her product beliefs.

its pre-ranking value needs to be sufficiently high and/or that of Product B needs to be sufficiently low. Part 1 of Lemma 1 provides an upper bound for the expected pre-ranking value of Product B, $p\bar{v}_B$, below which the consumer prefers Product A even if the expert always chooses methodology $r_{A||B}$.

When the consumer sees the ranking report, there is no uncertainty surrounding the ranking value; the only belief that is relevant to her product choice is the product belief. Since the report that ranks Product A first is an output of both methodologies, the consumer's updated product belief after seeing such a report, $\mu_{\bar{v}_B} = \frac{p-pq_{r_{A||B}}}{1-pq_{r_{A||B}}}$, depends on the expert's strategy and the prior over Product B. Note that the higher is the probability that the expert chooses $r_{A||B}$, the lower is the expected pre-ranking value of Product B. The ranking value from choosing the first-ranked Product A then needs to be sufficiently high for Product A to be preferred. Part 2 of Lemma 1 provides a lower bound, $p\bar{v}_B - \bar{v}_A$, for the ranking value, γ, for this to be the case even if the expert always chooses methodology $r_{A||B}$.

Unlike the above case, since the report that ranks Product B first is generated only under $r_{A||B}$, and only for pre-ranking value $\bar{v}_B$, the consumer's updated product belief is degenerate at $\mu_{\bar{v}_B} = 1$. The condition for the first-ranked Product B to be preferred is accordingly simpler, which is just that the ranking value is no less than $\bar{v}_A - \bar{v}_B$.

We next analyze the consumer's preferences with respect to the acquisition of the report, assuming that she chooses Product A without the report and the first-ranked product after seeing the report, the case that drives our experimental design:

Lemma 2. *Under product-purchasing rule $\rho(\varnothing) = \rho(A) = A$ and $\rho(B) = B$, the consumer is indifferent between acquiring the report ($\alpha = 1$) and not acquiring ($\alpha = 0$) if the ex-*

pert's mixed strategy $q_{r_{A||B}} = \frac{c}{p(\bar{v}_B+\gamma-\bar{v}_A)}$*; for* $q_{r_{A||B}}$ *higher (lower) than this threshold, the consumer strictly prefers* $\alpha = 1$ *(*$\alpha = 0$*).*

The intuition of Lemma 2 can be seen by the equivalent condition that the consumer prefers to acquire the report if $c \leqslant pq_{r_{A||B}}(\bar{v}_B+\gamma-\bar{v}_A)$. When the consumer would purchase Product A without seeing any report, acquiring the report makes a difference only when the report is B, which occurs when the pre-ranking value of Product B is $\bar{v}_B$ and the expert chooses $r_{A||B}$. By acquiring the report in this case, the consumer purchases the first-ranked Product B instead of Product A, receiving $\bar{v}_B + \gamma$ rather than $\bar{v}_A$. To justify paying for the report before learning what it is, the difference must present a probable benefit ($p > 0$ as is assumed and $\bar{v}_B + \gamma > \bar{v}_A$), the expert must choose $r_{A||B}$ at least some of the time ($q_{r_{A||B}} > 0$), and the resulting expected benefit must be no less than the cost $c > 0$.

We proceed to characterize the equilibria of the game, continuing to focus on the case where the consumer adopts the product-purchasing rule $\rho(\varnothing) = \rho(A) = A$ and $\rho(B) = B$, which is guaranteed to be a best response of the consumer under the parameter conditions in Lemma 1:[10]

Proposition 1. *Suppose that* $p\bar{v}_B \leqslant \bar{v}_A + (1-2p)\gamma$ *and* $\gamma \geqslant \max\{p\bar{v}_B - \bar{v}_A, \bar{v}_A - \bar{v}_B\}$ *and let* $\bar{q} = \frac{c}{p(\bar{v}_B+\gamma-\bar{v}_A)}$*. For* $0 < c \leqslant p(\bar{v}_B + \gamma - \bar{v}_A)$*, there exist three sets of perfect Bayesian equilibria:*

(a) the expert chooses $r_{A||B}$ *with probability* $q_{r_{A||B}} \in [0, \bar{q}]$*, and the consumer does not acquire the ranking report (*$\alpha = 0$*);*

(b) the expert chooses $r_{A||B}$ *with probability* $q_{r_{A||B}} \in [\bar{q}, 1]$*, and the consumer acquires the*

[10] Additional characterizations for product-purchasing rule $\rho(A) = A$ and $\rho(\varnothing) = \rho(B) = B$ can be found in Appendix A.2.

ranking report ($\alpha = 1$);

(c) *the expert chooses $r_{A||B}$ with probability $q_{r_{A||B}} = \bar{q}$, and the consumer randomizes between $\alpha = 0$ and $\alpha = 1$ with any probability.*

In all these equilibria, the consumer adopts the product-purchasing rule $\rho(\varnothing) = \rho(A) = A$ and $\rho(B) = B$.

Recall that the expert only cares about selling the report to earn $\pi > 0$. He is indifferent to the choice of ranking methodology (and the consumer's product choice) so long as the consumer acquires the report. Accordingly, the expert is always willing to randomize between the two methodologies, and the randomization probabilities are chosen in equilibrium to support either a fixed or a randomized acquisition decision of the consumer. On the other hand, when the consumer is indifferent between acquiring the report or not, any randomization on the part of the consumer constitutes an equilibrium – the consumer does not need to randomize in specific ways to generate indifference for the expert, who is already indifferent. Note that the characterized equilibria include two pure-strategy equilibria. We call the equilibrium in which the expert always chooses r_A [under Proposition 1(a)] the *uninformative-ranking equilibrium* and that in which he always chooses $r_{A||B}$ [under Proposition 1(b)] the *informative-ranking equilibrium.*

Proposition 1 provides several insights about the interaction between the expert and the consumer. First, to make the ranking report attractive so that the consumer acquires it, the ranking value has to be sufficiently high while the cost of the report sufficiently low. This may explain why profit-making ranking reports typically exist for expensive goods or important decisions, such as automobiles and colleges, while the rankings of other small goods either do not exist or are provided by non-profit institutes such as *Consumer Reports.*

Second, the expert can manage to induce the consumer to acquire the report by counter-intuitively choosing a methodology that makes the consumer worse off. Consider the pure-strategy equilibrium in which the expert always chooses $r_{A||B}$ and suppose that $\bar{v}_A > \bar{v}_B$. When Product A is always better than Product B, acquiring the report or not the consumer prefers methodology r_A: when acquiring, the consumer receives a payoff of $\bar{v}_A + \gamma - c$ from choosing the always first-ranked Product A under r_A and with probability p trades the higher $\bar{v}_A$ for the lower $\bar{v}_B$ in order to obtain γ when Product B is ranked first under $r_{A||B}$; when not acquiring, by choosing Product A she receives $\bar{v}_A + \gamma$ under r_A and with probability p receives γ less when Product B is ranked first under $r_{A||B}$.

While relative to r_A the adoption of $r_{A||B}$ harms the consumer, with the parameter values assumed in Proposition 1 it harms her more when she does not acquire the report, in which case she loses $p\gamma$, than when she acquires the report, in which case she loses $p(\bar{v}_A - \bar{v}_B)$. The consumer, however, is offered an opportunity to avoid the incremental harm $p(\bar{v}_B + \gamma - \bar{v}_A)$ by paying for the ranking report which costs less than the incremental harm. In a sense, the expert strategically creates a very unfavorable situation for the consumer if she does not acquire the report and offers a less unfavorable situation with a mutually beneficial price.

The last insight concerns the informativeness and consumer-optimality of the ranking methodologies. The discussion above shows that the consumer may be better off under the uninformative methodology r_A. However, since she makes the same product choice under r_A with or without the report, acquiring the report does not bring any marginal benefit yet it is costly. The consumer is only willing to pay for the report when there is sufficiently high probability that it is generated by the informative $r_{A||B}$. This is the case even when r_A

is consumer-optimal and $r_{A||B}$ is not (as when $\bar{v}_A > \bar{v}_B$): the consumer report-acquisition decision hinges on whether the underlying methodology is informative, not whether it is consumer-optimal (i.e., ranks products according to the consumer's pre-ranking values).

1.4 Experimental Design

1.4.1 Treatments

The theoretical analysis in Section 1.3 guides the design of our treatments. There are six parameters in our game, $\bar{v}_A$, $\bar{v}_B$, γ, p, c, and π. In choosing the treatment values for these parameters, we are guided by three considerations: (i) to satisfy the parameter conditions in Proposition 1 so that the characterized equilibria exist and serve as the theoretical benchmarks for evaluating observed behavior; (ii) to generate treatment variations so that the otherwise same equilibria have different implications regarding informativeness and consumer-optimality; and (iii) to provide salience within the confine of the first two considerations by making the payoff differences from different strategies reasonably large.

The outcome of these considerations is the following choices of parameters, which give us three treatments. We assign values for five parameters that are fixed across all three treatments: $\bar{v}_A = 100$, $\gamma = 250$, $p = 0.6$, $c = 70$, and $\pi = 5c = 350$. For treatment variations, we vary the value of $\bar{v}_B$ at 120, 80, and 60. We call the resulting treatments, *SUPERIOR-B*, *AVERAGE-B*, and *INFERIOR-B*. Table 1.1 summarizes the three treatments and presents the players' expected payoffs at the strategically simultaneous stage where the expert chooses the ranking methodology and the consumer makes the report-acquisition decision. The consumer's expected payoffs are calculated based on the equilibrium product-purchasing rule $\rho(\varnothing) = \rho(A) = A$ and $\rho(B) = B$.

Table 1.1: Experimental Treatments

Treatment *SUPERIOR-B*: $\bar{v}_B = 120$

Expert \ Consumer	$\alpha = 0$	$\alpha = 1$
r_A	**0, 350**	350, 280
$r_{A\|\|B}$	0, 200	**350, 292**

Treatment *AVERAGE-B*: $\bar{v}_B = 80$

Expert \ Consumer	$\alpha = 0$	$\alpha = 1$
r_A	**0, 350**	350, 280
$r_{A\|\|B}$	0, 200	**350, 268**

Treatment *INFERIOR-B*: $\bar{v}_B = 60$

Expert \ Consumer	$\alpha = 0$	$\alpha = 1$
r_A	**0, 350**	350, 280
$r_{A\|\|B}$	0, 200	**350, 256**

Note: r_A and $r_{A||B}$ are two ranking-methodology choices of the expert. $\alpha = 0$ and $\alpha = 1$ represent the respective cases where the consumer does not acquire and acquire the ranking report. In each cell, the first number is the expert's payoff and the second number the consumer's expected payoff for the given (r, α), based on the product-purchasing rule $\rho(\varnothing) = \rho(A) = A$ and $\rho(B) = B$. The bold pairs of numbers correspond to the pure-strategy equilibria of the game. In all three treatments, the parameters of the game take the values $\bar{v}_A = 100$, $\gamma = 250$, $p = 0.6$, $c = 70$.

The games induced in the three treatments have the same set of pure-strategy equilibria, while their mixed-strategy equilibria differ only by the randomization probabilities of the expert. For each treatment, the bold pairs of payoff numbers correspond to the two pure-strategy equilibria. Note that only in treatment *SUPERIOR-B* does the pure-strategy equilibrium in which the expert sells the report with $r_{A||B}$ rank products according to the consumer's pre-ranking values of the products (because $\bar{v}_B > \bar{v}_A$).

If we interpret choosing a consumer-sub-optimal methodology as a form of harm to the consumer, there is a sense that the expert in selecting $r_{A||B}$ harms the consumer more in *INFERIOR-B* than in *AVERAGE-B* and does not harm the consumer at all in *SUPERIOR-B*. Such a monotonicity in the extent of consumer harm corresponds with the increases in the consumer's expected payoffs from $(r, \alpha) = (r_{A||B}, 1)$ across the treatments, from 256 in *INFERIOR-B* to 268 in *AVERAGE-B* to 292 in *SUPERIOR-B*. Note that all the other expected payoff numbers are constant across the treatments; despite the existence of essentially the same sets of equilibria in all treatments, these distinctive changes in the consumer's expected payoffs provide us with a clean comparative statics.

Finally, note that in all treatments the expert and the consumer have conflicting preferences over the two pure-strategy equilibria. While the expert prefers the equilibrium in which he sells the report by choosing $r_{A||B}$, the consumer prefers the equilibrium in which she does not acquire the report generated by r_A. In their respective preferred equilibria, both the expert and the consumer earn the same (expected) payoffs of 350.

1.4.2 Procedures

Our experiment is conducted using oTree (Chen, Schonger and Wickens, 2016) at the Financial Services Laboratory of Lehigh University. A total of 200 subjects participate in the three treatments using a between-subject design. Subjects have no prior experience with the experiment and are recruited from the undergraduate and graduate population of the university.

Table 1.2 reports a summary of participation for three treatments. Four sessions are conducted for each treatment. On average, 17 subjects participate in a session, with half of them randomly assigned to the role of an expert and the other half to the role of a consumer. Roles remain fixed throughout a session. Experts and consumers in a session are *randomly matched* in each round to form groups of two to play 40 rounds of the game.

Upon arrival, subjects are instructed to sit at separate computer terminals with partitions. Each receives a copy of the experimental instructions. The instructions are read aloud using slide illustrations as an aid. A comprehension quiz and a practice round follow.

Subjects are told that there are two products, A and B. While the value of Product A is fixed at 100, that of Product B is random and varies between 0 and, using treatment *AVERAGE-B* as an example, 80, where there is a 60% chance that the computer selects 80.

The expert makes only one decision in each round, choosing between Ranking Method 1 (r_A) and Ranking Method 2 ($r_{A||B}$). The consumer makes two decisions, first choosing whether to acquire the expert's ranking report and then deciding which product to choose. If the consumer chooses to acquire the report, before her product choice is made she will see a report that is generated according to the randomly selected value of Product B and the ranking method chosen by the expert. The ranking report is framed as "Product K is ranked first." If the consumer chooses not to acquire the report, she will proceed directly to choose a product. A report will still be generated in this case to determine the ranking value.

Rewards in each round are determined based on, for the expert, the consumer's report-acquisition decision and, for the consumer, the generated report, her report-acquisition decision, and her product choice, all according to the parameters of the treatment. A more elaborate reward table than those in Table 1.1 is presented on each subject's screen throughout the experiment.[11]

A round is concluded by the provision of a feedback, which summarizes what has happened in the round, including the expert's chosen ranking method, the selected value of Product B, the generated ranking report, the consumer's product choice, and the subject's

[11]Refer to Appendix B for a sample of instructions.

reward for the round. A history of the happenings in all previous rounds is also available to subjects.

We randomly select three rounds for payments. The average reward a subject earns in the three selected rounds is converted into US Dollars at a fixed and known exchange rate of US$1 per 20 reward points. A show-up fee of US$5 is also provided. Subjects on average earn US$14.88 by participating in a session that lasts about an hour.

1.5 Findings

1.5.1 Pure-Strategy Equilibrium Plays

For the parameters adopted in our three treatments, there are two pure-strategy equilibria, what we call the informative-ranking and uninformative-ranking equilibria, and a continuum of mixed-strategy equilibria. We begin our analysis by examining the frequencies of plays of the two pure-strategy equilibria.

Finding 1. *In all three treatments, the frequencies of plays of either the informative-ranking or the uninformative-ranking equilibria are significantly higher than those that would occur under completely random behavior. There are no significant differences in the frequencies of these equilibrium plays between the three treatments.*

In the presence of multiple equilibria, subjects need to coordinate over which equilibrium to play, if at all. Our first finding provides evidence that subjects' behavior is purposefully shaped by the equilibrium incentives. Figure 1.1 presents the frequencies of plays of the two pure-strategy equilibria over rounds. In an instance of play of the game, there are 10 possible combinations of choices made by the expert and consumer. Among them, three

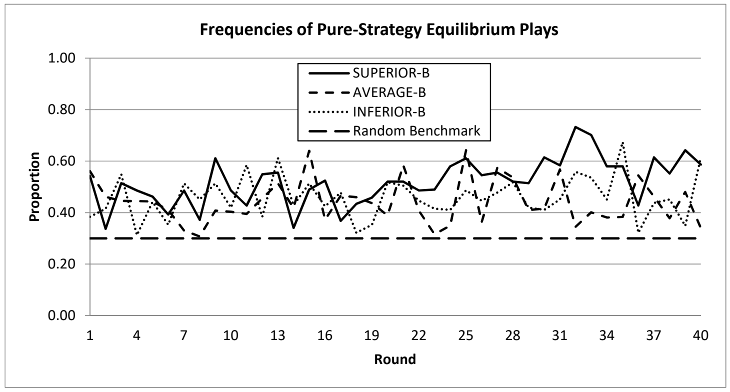

Figure 1.1: Pure-Strategy Equilibrium Plays

combinations are consistent with the pure-strategy equilibria, one with the uninformative-ranking equilibrium and two with the informative-ranking equilibrium. The frequency of plays of any one of these equilibria, aggregated across all sessions and all rounds of the three treatments, is 47.32%, significantly higher than the 30% benchmark when the equilibria are played as a result of completely random choices ($p < 0.001$, Wilcoxon signed-rank test).[12]

The frequency of equilibrium plays is higher in *SUPERIOR-B* (51.98%) than in *AVERAGE-B* (44.10%) and *INFERIOR-B* (45.87%), although the differences between the three treatments are not statistically significant (two-sided $p \geqslant 0.3429$ for any pairwise comparison, Mann-Whitney tests).[13]

We next examine the frequency of each of the two pure-strategy equilibria conditional on equilibrium behavior. We perform two comparisons, a within-treatment comparison and a between-treatment comparison.

Finding 2. *Conditional on equilibrium plays,*

[12]Figure 1.1 indicates that there is no systematic convergency in the frequencies of equilibrium plays. We accordingly use all-round data for statistical tests. Unless otherwise indicated, the reported p values are from one-sided tests.

[13]A Kruskal-Wallis test further confirms that the three frequencies have no statistical differences from one another ($p = 0.5367$).

(a) *the frequency of the informative-ranking equilibrium is significantly higher than that of the uninformative-ranking equilibrium in SUPERIOR-B and INFERIOR-B, while there is no significant difference between the two frequencies in AVERAGE-B; and*

(b) *the frequency of the informative-ranking equilibrium in SUPERIOR-B is higher than those in AVERAGE-B and INFERIOR-B, but only the difference with AVERAGE-B is statistically significant.*

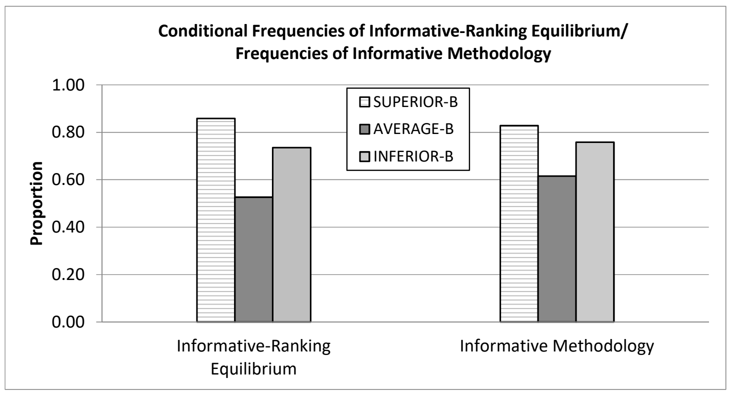

Figure 1.2: Informative-Ranking Equilibria and Informative Methodology

Finding 2(a) reports three within-treatment comparisons. The first set of bars in Figure 1.2 presents the conditional frequencies of the informative-ranking equilibrium. Among the equilibrium plays observed in *SUPERIOR-B* and *INFERIOR-B*, the informative-ranking equilibrium, which is played respectively 85.86% and 73.55% of the time, is the equilibrium that is observed more often with statistical significance ($p = 0.0625$, Wilcoxon signed-rank tests).[14] On the other hand, the informative-ranking and the uninformative-ranking equilibria are played almost as often as each other in *AVERAGE-B*, with no significant difference between their conditional frequencies, which are respectively 52.65% and 47.35% (two-sided $p = 0.625$, Wilcoxon signed-rank test).

[14] $p = 0.0625$ is the lowest possible p value from the Wilcoxon signed-rank test with four independent observations.

The between-treatment comparisons reported in Finding 2(b) are apparent from the figures above. The informative-ranking equilibrium is played most often in *SUPERIOR-B*, followed by *INFERIOR-B* and then *AVERAGE-B*. Statistically, the higher frequency in *SUPERIOR-B* over that in *AVERAGE-B* is significant ($p = 0.0286$, Mann-Whitney test), while the difference between *SUPERIOR-B* and *INFERIOR-B* is not ($p = 0.1227$, Mann-Whitney test). The higher frequency in *INFERIOR-B* over that in *AVERAGE-B* is likewise not significant ($p = 0.1$, Mann-Whitney test).

The frequencies reported above present a non-monotonic comparative statics. Recall that the only difference in the parameters of the three treatments is $\bar{v}_B$, the high realization of the pre-ranking value of Product B. The value assumed by $\bar{v}_B$ decreases monotonically from 120 in *SUPERIOR-B* to 80 in *AVERAGE-B* to 60 in *INFERIOR-B*. The conditional frequency of the informative-ranking equilibrium, however, decreases from 85.86% in *SUPERIOR-B* to 52.65% in *AVERAGE-B* and then increases again to 73.55% in *INFERIOR-B*. We proceed to examine the observed behavior of the expert-subjects and the consumer-subjects that comprise these non-monotonic frequencies of strategic outcomes.

1.5.2 Experts' and Consumers' Behavior

We begin our analysis of subjects' behavior by examining the ranking methodology choices of the experts. We again perform a within-treatment and a between-treatment comparisons.

Finding 3. *For the experts' choices of ranking methodology,*

(a) the frequency of the informative methodology is higher than that of the uninformative methodology in all three treatments, with statistical significance in SUPERIOR-B and INFERIOR-B; and

(b) *the frequency of the informative methodology in SUPERIOR-B is higher than those in AVERAGE-B and INFERIOR-B, but only the difference with AVERAGE-B is statistically significant.*

The second set of bars in Figure 1.2 presents the frequencies of the informative methodology $r_{A||B}$. The same non-monotonicity across treatments noted in the frequencies of the informative-ranking equilibrium is also observed in the experts' choices of the informative methodology. The experts choose $r_{A||B}$ 82.79% of the time in *SUPERIOR-B*, 61.48% of the time in *AVERAGE-B*, and 75.77% of the time in *INFERIOR-B*. Except for *AVERAGE-B*, these frequencies are significantly higher than the complementary frequencies of uninformative methodology ($p = 0.0625$ for *SUPERIOR-B* and *INFERIOR-B* and $p = 0.1875$ for *AVERAGE-B*, Wilcoxon signed-rank tests). For between-treatment comparisons, the higher frequency in *SUPERIOR-B* over that in *AVERAGE-B* is statistically significant ($p = 0.0571$, Mann-Whitney test), while the differences in the other two pairwise comparisons are not ($p \geqslant 0.1714$, Mann-Whitney tests).

The experts' strategies put an "upper bound" on what can be achieved as equilibrium outcomes. The cross-treatment non-monotonicity in the frequencies of the experts' ranking methodologies appears to be the main driving force for the similar non-monotonicity in the frequencies of the pure-strategy equilibria. To evaluate the extent of how consumers' behavior also contributes to shape the observed equilibrium outcomes, we proceed to examine the choices made by the consumers.

Our equilibrium characterization of the expert's choice of ranking methodology and the consumer's report-acquisition decision is based on consumer's adoption of the product-purchasing rule $\rho(\varnothing) = \rho(A) = A$ and $\rho(B) = B$, which is a best response given the parameter values of our treatments. We examine the consumers' report-acquisition decisions as well as the extent to which their plays are consistent with this rule.

Finding 4. *The consumers in AVERAGE-B acquire the ranking report significantly less often than they do not acquire, while there are no significant differences between acquiring and not acquiring in SUPERIOR-B and INFERIOR-B. For product choices,*

(a) conditional on not acquiring the report, the frequency of Product A is significantly higher than that of Product B in all three treatments; and

(b) conditional on acquiring the report, the consumers in all three treatments virtually always choose the first-ranked products.

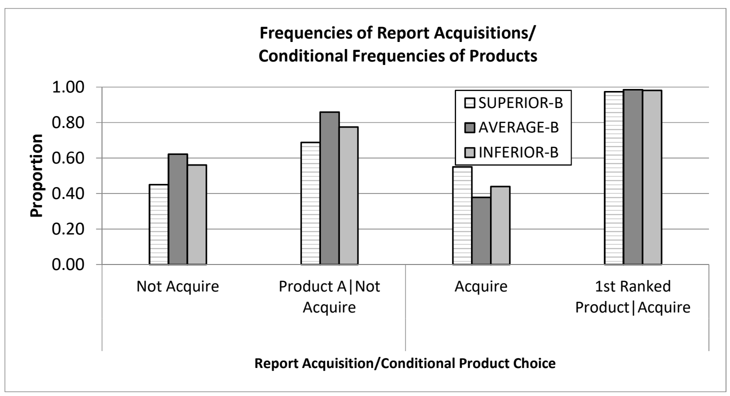

Figure 1.3: Consumers' Report Acquisitions and Product Choices

Figure 1.3 presents the frequencies of consumers' report-acquisition decisions and conditional product choices. Except in *AVERAGE-B*, consumers choose to acquire the ranking report roughly half of the time, with 55.01% in *SUPERIOR-B* and 43.94% in *INFERIOR-B*, which makes the frequencies of acquiring and not acquiring not significantly different from each other (two-sided $p \geqslant 0.375$, Wilcoxon signed-rank tests). The consumers in *AVERAGE-B*, however, chooses not to acquire the report 62.14% of the time, significantly more often than acquiring ($p = 0.0625$, Wilcoxon signed-rank test).

We also observe in *AVERAGE-B* the highest frequency of Product A conditional on not acquiring the report, which is recorded at 85.88%. This is followed by 77.55% in *INFERIOR-B* and then 68.33% in *SUPERIOR-B*. All these frequencies are significantly higher than the corresponding conditional frequencies of Product B ($p = 0.0625$, Wilcoxon signed-rank test). Conditional on acquiring the report, consumers choose the first-ranked products with frequencies above 97% in all three treatments.

The above analysis suggests that in *AVERAGE-B* consumers' report-acquisition decisions and product choices combine with experts' choices of ranking methodology to drive less frequent plays of the informative-ranking equilibrium (and more frequent plays of the uninformative-ranking equilibrium). In relative terms, the contribution of consumers' behavior to the observed equilibrium outcomes is less pronounced in the other two treatments, where the equilibrium outcomes are more strongly driven by the experts' choices of methodology. In section 1.6, we discuss what behavioral factors in addition to equilibrium incentives may contribute to these non-monotonic findings across the three treatments.

1.5.3 Dynamic Analysis

We adopt a panel setting to examine the individual behavior through 40 rounds in the experiment. For participant i in round t, we define dummy variables $Acquire_{it} = 1$ if the consumer acquires the ranking report, and $Informative_{it} = 1$ if the expert selects an informative ranking methodology $r_{A||B}$. The dynamic reactions are estimated separately by:

$$Acquire_{it} = \alpha_1 + \beta_1 \ \mathbf{L.Informative} + \gamma_1 \ \mathbf{L.Acquire} + \delta_1 \ \mathbf{X} + u_{it} \tag{1.1}$$

$$Informative_{it} = \alpha_2 + \beta_2 \ \mathbf{L.Acquire} + \gamma_2 \ \mathbf{L.Informative} + \delta_2 \ \mathbf{X} + \epsilon_{it} \tag{1.2}$$

where, **L.Informative** and **L.Acquire** are vectors with lag values of $Informative_{it}$ and $Acquire_{it}$ respectively, and **X** a vector including fixed effects of treatments and sessions.

We employ a random effect panel Logit regression to estimate above two models, with results presented in Table 1.3 and 1.4.

Finding 5. *For individual subjects through 40 rounds,*

(a) consumers are likely to change behavior responding to the experience from the previous rounds; and

(b) experts behave more consistently.

Table 1.3 presents the dynamic estimation of consumers' report-acquisition behavior. We find that consumers' report-acquisition behavior is significantly related to her paired experts' behavior in previous rounds. If a consumer was paired with an expert who selected ranking methodology $r_{A||B}$ in previous round, she is more likely to acquire the ranking report in current round, which suggests that the consumer adjusts her behavior responsively based on her experience. This finding is insightful and provides another explanation why we do not see a lot consumers acquiring the ranking. That is, if consumers saw ranking methodology r_A in the previous round, they have incentive to skip the ranking, which could make them better off under the condition of ranking methodology r_A.

On the other side, experts do not play responsively, which is shown by the regression results in Table 1.4. We do not see a statistically significant relation between the expert's behavior and the behavior of the paired consumer in the previous rounds. Instead, experts tend to have a self-consistent behavior: they are significantly more likely to make the same decision as they did in previous rounds. This result again demonstrates our main finding that experts could figure out an optimal strategy and resist during most of the time.

1.5.4 Experience-Weighted Attraction Learning

There is a growing literature on learning in laboratory experiments, with the concern that players usually need to learn to play equilibrium strategies in games where they do not have much experience playing. In this section, we present an econometric analysis of learning with two goals: (1) to examine whether players in our ranking games learn to play equilibrium strategies, and (2) to explain in more detail why not so many consumers choose to acquire the ranking report. We adopt the experience-weighted attraction (EWA) learning model developed by Camerer and Ho (1999), which integrates reinforcement and belief-based learning models.[15]

The EWA learning models a normal form game with N players and J strategies for each player. Let $s_i^j(t)$ be player i's strategy j in period t and $s_i(t)$ player i's chosen strategy in period t. Player i's payoff of choosing strategy j in period t is $\pi_i(s_i^j, s_{-i}(t))$. The core of EWA learning is two updated variables, experience weight $N(t)$ and attractions $A_i^j(t)$ with given initial values of $N(0)$ and $A_i^j(0)$ that represent pregame experience. Specifically, the experience weight is updated according to

$$N(t) = \phi(1-k)N(t-1) + 1 \tag{1.3}$$

[15]Studies like Salmon (2001), Hopkins (2002) and Camerer (2003) reviewed different learning models.

where $\phi \in [0,1]$ discounts past experience and $k \in [0,1]$ denotes the updating speed of an experience weight. The attraction of strategy j for player i is updated according to

$$A_i^j(t) = \frac{\phi N(t-1)A_i^j(t-1) + [\delta + (1-\delta)I(s_i^j, s_i(t))]\pi_i(s_i^j, s_{-i}(t))}{N(t)} \tag{1.4}$$

where the indicator function $I(s_i^j, s_i(t))$ is equal to one if $s_i^j = s_i(t)$, and $\delta \in [0,1]$ represents the weight put on foregone payoffs.

Based on attractions, the probability of player i choosing strategy j in period $t+1$ is computed by the following logistic function:

$$P_i^j(t+1) = \frac{exp(\lambda A_i^j(t))}{exp(\sum_{k=1}^{J} \lambda A_i^k(t))} \tag{1.5}$$

where $\lambda \in [0,+\infty]$ reflects the response sensitivity for mapping attractions into choice probabilities.

There are four key parameters in the EWA learning model: λ, ϕ, δ, k, of which each captures a behavioral principle of learning. It is a reinforcement learning when $\delta = 0$ and $k = 1$, while a belief based learning when $\delta = 1$ and $k = 0$. We estimate these parameters using the method of maximum likelihood with the log-likelihood function given by:

$$LL(\lambda, \phi, \delta, k) = \sum_{t=1}^{T}\sum_{i=1}^{N} \ln[\sum_{j=1}^{J} I(s_i^j, s_i(t))P_i^j(t-1)] \tag{1.6}$$

where we assume initial value of $N(0) = 1$ following previous studies,[16] and compute initial attractions $A_i^j(0)$ by the first-period data as described in Ho, Wang and Camerer (2007).

We examine the EWA learning process for consumers and experts separately in each

[16]Camerer (2003) suggests that this assumption has no substantial effects on the learning process.

treatment, and report the estimates for parameters λ, ϕ, δ, k in Table 1.5 and 1.6.

Overall, the estimates for λ, δ and k reveal the difference in learning between consumers and experts. Although both players have low decay of previous attractions (given by high ϕ), consumers are more likely to play according to belief based learning with positive δ and lower k, while experts are more likely to play according to reinforcement learning with insignificant δ and higher k.

Specifically, consumers are more sensitive to updated attractions when making decisions of whether to acquire the ranking (given by big λ), significantly care about the foregone payoffs (given by δ), and put higher weight on recent experience (given by lower k). This explains why not so many consumers choose to acquire ranking with high probabilities, since they pay attention to the hypothetical payoffs if they do not acquire the ranking, which increases the attraction of not acquiring.

In contrast, experts are less sensitive to updated attractions when choosing ranking methodologies and do not care about foregone payoffs. This is reasonable as experts are indifferent between two ranking methodologies. Comparing three treatments, experts in treatment *AVERAGE-B* seem to play differently by putting higher weights on recent experience, while in other treatments, experts consider each past period equally. This may help to explain why relatively fewer experts choose informative ranking methodology in treatment *AVERAGE-B*.

1.6 Discussion

We discuss three issues in this section: (i) the non-monotonic result that the expert is more likely to select the informative ranking methodology $r_{A||B}$ in *SUPERIOR-B* (when this ranking methodology is consumer optimal) and *INFERIOR-B* (when this methodology is the furthest from being consumer optimal) and less likely in *AVERAGE-B* (when this ranking methodology is consumer-sub-optimal, albeit closer to being consumer-optimal); (2) the consumer acquires the ranking report with a relatively low probability in *AVERAGE-B* and *INFERIOR-B*; and (3) whether lying or deception is relevant in our experiment.

1.6.1 Experts' Decision: Social Preferences vs Economic Incentives

Many studies have shown that people have social preferences. Bowles and Polania-Reyes (2012) define social preferences as motives such as altruism, reciprocity, intrinsic pleasure in helping others, inequity aversion, ethical commitments, and other motives that induce people to help others more than would an own-material-payoff maximizing individual. For instance, Engel (2011) conducts a meta study of 131 dictator game experiments with 616 treatments and finds solid evidence that the majority of dictators are generous, even though giving behavior varies across different social groups. In a survey of the extensive literature of measuring social preferences in experimental games, Levitt and List (2007) summarize various factors beyond monetary motivations that could influence subjects' behavior in the lab. Although we attempt to avoid the influence of most factors, it is impossible to completely eliminate the subjects' social preferences.

In our model, of the two ranking methodologies r_A and $r_{A||B}$ available to the product expert, if the expert selects r_A with sufficiently high probability, then the consumer does not acquire the expert's ranking. Alternatively, if the expert selects $r_{A||B}$ with sufficiently high probability, then the consumer does. Therefore, to motivate the consumer to acquire his ranking, the expert has an economic incentive to select $r_{A||B}$ with sufficiently high probability. However, social preferences might also influence the expert. Specifically, in all treatments an expert might select r_A because the consumer's greatest payoff is from the action profile in which the expert selects this ranking methodology and the consumer does not acquire the ranking (and selects product A). Therefore, an expert who is concerned about a consumer might select r_A.

Evaluating the consumer's and expert's payoffs across our three treatments, the only payoff that changes is consumer's from the expert selecting $r_{A||B}$ and the consumer acquiring the ranking. Therefore, in comparing the results of our experiment across the three treatments, we should focus on the consumer's payoff from this action profile. In Table 1.1, we see that this payoff is 292 in *SUPERIOR-B*, 268 in *AVERAGE-B*, and 256 in *INFERIOR-B*. Note that the consumer's payoff from this action profile changes across these treatments because $\bar{v}_B$ changes.

Comparing these treatments, the expert needs to provide the smallest economic incentive for the consumer to acquire the ranking by selecting $r_{A||B}$ in treatment *SUPERIOR-B* and the greatest incentive in treatment *INFERIOR-B*. With regard to social preferences, an expert who has concern for the consumer's utility might be most likely to select $r_{A||B}$ in treatment *SUPERIOR-B* and least likely to select this ranking methodology in treatment *INFERIOR-B*.

In sum, social preferences and economic incentives perform opposite trends from treatment *SUPERIOR-B* to treatment *INFERIOR-B*. Therefore, we cannot determine in each case the net effect of the economic incentives versus the social preferences effect. As a result, we could observe the non-monotonicity of the experts' selected ranking methodology, as displayed in Figure 1.2, moving from the *SUPERIOR-B* to the *AVERAGE-B* to the

INFERIOR-B treatment.

1.6.2 Consumers' Decision

A consumer in our experiment makes two decisions, whether to acquire the expert's product ranking and which product to select. Each of these decisions involves uncertainty about the expert's selected ranking methodology and about the value of product B. If the expert acquires the product ranking, then using Bayes rule, she can update her probabilistic belief about the value of product B. Some concerns might influence whether she applies Bayes rule correctly and whether she calculates expected payoffs correctly.

The first concern is cognitive cost that may discourage decision makers when information is overloaded (e.g. Chetty, Looney and Kroft, 2009; Kuksov and Villas-Boas, 2010; Ghose, Ipeirotis and Li, 2014). In our experimental results, we find that a certain number of consumers in all 40 rounds never acquire the product ranking report. They might, by doing so, be skipping the first decision and making their decision process easier by selecting only a product.

In addition, theories about salience contend that when dealing with uncertainty, decision makers may focus on in all 40 rounds partial information (e.g. Chetty, Looney and Kroft, 2009; Bordalo, Gennaioli and Shleifer, 2012). Therefore, salience may affect the consumer in treatments 2 and 3 where the pre-ranking value of product B is always lower than that of product A. If the consumer pays too much attention to this value rather than to the ranking value, she may think product A will give her a higher expected payoff.

Finally, the social preferences may also work on the consumer. For instance, if the consumer realizes that the expert is more likely to choose the ranking methodology $r_{A||B}$, which reduces the consumer's expected payoff compared to r_A, she may have incentive to fight back by refusing to acquire the ranking report. In particular, when refusing to acquire the ranking report, the expert would lose much more than the consumer.

1.6.3 Lying and Deception

In the context of our experiment, considering factors that could affect the expert's selection of $r_{A||B}$ in treatments *AVERAGE-B* and *INFERIOR-B*, we should evaluate the possibility that the expert might be deceiving, or lying to, the consumer. In our model, *ex ante* the expert and the consumer know that the pre-ranking value of product A is 100, and both are uncertain about the value of product B. Consider treatment *INFERIOR-B* in which the expert and consumer believe that the consumer's pre-ranking value of B could be either $\bar{v}_B = 60$ or 0. Despite being uncertain about the value of product B, both the expert and consumer therefore know that the pre-ranking value of product A is greater than either pre-ranking value of product B. Prior to ranking the products, but after selecting a ranking methodology, the expert learns the pre-ranking value of product B, whether it is 0 or 60.

Applying the Sobel (2018) definition of deception in settings with strategic information transmission to our model, an expert's ranking methodology (i.e., message function) is *not deceptive* if for each possible product ranking that results from using the methodology, the consumer who acquires the ranking accurately updates probabilistic beliefs about pre-ranking product values (so that her probabilistic beliefs are identical to the expert's). In our scenario, the ranking methodology $r_{A||B}$, despite possibly ranking product B first, is not deceptive.

In addition, our instructions for the experiment state that when using $r_{A||B}$, the expert is ranking products according to $\bar{v}_B$. Therefore, because the expert is not deceiving the consumer and is not lying to the consumer by using $r_{A||B}$ in treatments *AVERAGE-B* and *INFERIOR-B*, we believe that possible deception or even lying is not driving our results.[17]

However, we should note that by creating a ranking value in the consumer's utility of a product, in our two-product model, the expert's ranking boosts the value of the product he ranks first. Therefore, if the expert uses $r_{A||B}$ and ranks product B first, the expert boosts the consumer's value of product B. The expert does so despite the fact that the consumer's pre-ranking value of product B is less than her pre-ranking value of A.

If we consider an expert's new-car automobile ranking in which his creation of a ranking value is due to his influence on the eventual used-car market, then in a sense the expert is deceiving future used-car buyers. To do so, these used-car buyers must use the expert's ranking as salient information about pre-ranking value. (Otherwise, these used-car buyers would correctly update beliefs about pre-ranking value using Bayes rule.) In this scenario, the consumers in our experiment while benefiting from the expert's ranking value are not deceived by the expert's use of $r_{A||B}$.

In reality, the publisher of a product ranking has the opportunity to influence a consumer's choice of whether to acquire the ranking through ranking value or through deceiving the consumer. Our goal in the design of our experiment was to create an environment in which the expert by using $r_{A||B}$ influences the consumer's utility and therefore decision whether to acquire the ranking, without deceiving the consumer. That is, we wanted to isolate the effect of the expert's influence on the consumer's utility from the possible

[17] There is an ongoing debate in current studies about whether people have preferences for truth-telling, or equivalently an aversion to lying. For instance, Abeler, Nosenzo and Raymond (2016) conclude that people lie less than we expect; while Sánchez-Pagés and Vorsatz (2007) and Gneezy, Kajackaite and Sobel (2018) find evidence that lying occurs more often when the behavior involving a lie is not observed.

deception of the consumer.

1.7 Concluding Remarks

In our experimental analysis, we find evidence that product experts are likely to select a ranking methodology that involves considerable uncertainty about the final product ranking. In two of our three treatments, this methodology involves ranking products in a manner that is inconsistent with the consumers' pre-ranking values of the products being considered.

Dearden, Grewal and Lilien (2019, forthcoming) demonstrate in a theoretical model that a critical factor in this result is that a product's rank affects consumer utility. A consumer who acquires an expert's ranking learns the ranking value of each product being considered; and a consumer who does not acquire the ranking does not learn these ranking values. It is the consumer's uncertainty associated with the ranking values that drives the expert to select a ranking methodology that adds randomness to the ranking outcome, thereby motivating the consumer to learn these values by viewing an expert's ranking. The added uncertainty comes at the expense of ranking products in a manner consistent with a consumer's pre-ranking values.

The expert-consumer game that is the basis of our experiment is riddled with multiple equilibria. In one the expert selects a ranking methodology which involves no uncertainty about the product ranking and accordingly the consumer does not acquire the ranking. In another, the expert selects a ranking methodology which involves considerable uncertainty about the product ranking and accordingly the consumer acquires the ranking. In this game with multiple equilibria, we find the tendency in our experimental setting for product experts and consumers to play the equilibrium in which the expert selects a ranking methodology that does not necessarily rank products according to the consumer's pre-ranking utilities and which involves considerable uncertainty about the expert's ranking. This result indicates that product experts, for the purpose of motivating consumers to acquire their rankings, are willing to select ranking methodologies that are not best for consumers.

Tables

Table 1.2: Average Payment in Each Session

Treatment	Session	Number of Subjects	Average Payment
SUPERIOR-B	11	16	15.93
SUPERIOR-B	12	16	15.36
SUPERIOR-B	13	18	12.75
SUPERIOR-B	14	14	18.34
AVERAGE-B	21	20	15.22
AVERAGE-B	22	16	13.26
AVERAGE-B	23	16	15.29
AVERAGE-B	24	22	13.78
INFERIOR-B	31	14	14.53
INFERIOR-B	32	16	15.13
INFERIOR-B	33	16	14.61
INFERIOR-B	34	16	15.38
Mean		17	14.88

Table 1.3: Dynamic Analysis of Consumers' Behavior

	(1) VDL1	(2) VDL2	(3) VDL3
Acquire			
L.Informative	0.3490**	0.3697**	0.3893**
	(0.1204)	(0.1215)	(0.1265)
L2.Informative		0.1497	0.1602
		(0.1115)	(0.1188)
L3.Informative			0.0834
			(0.1266)
L.Acquire	0.0397	0.0066	−0.0611
	(0.2010)	(0.1816)	(0.1807)
L2.Acquire		0.5259**	0.4769**
		(0.1686)	(0.1648)
L3.Acquire			0.3760*
			(0.1836)
N	3900	3800	3700

Clustered standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.4: Dynamic Analysis of Experts' Behavior

	(1) RDL1	(2) RDL2	(3) RDL3
Informative			
L.Acquire	0.1850	0.2242*	0.1887
	(0.0984)	(0.0984)	(0.0999)
L2.Acquire		0.0876	0.0600
		(0.0983)	(0.1001)
L3.Acquire			0.0496
			(0.1090)
L.Informative	1.2344***	1.1039***	1.1074***
	(0.2039)	(0.1730)	(0.1781)
L2.Informative		0.5692***	0.4942**
		(0.1597)	(0.1581)
L3.Informative			0.3304*
			(0.1508)
N	3900	3800	3700

Clustered standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.5: Estimates of EWA Learning Model for Consumers

	(1) *SUPERIOR-B*	(2) *AVERAGE-B*	(3) *INFERIOR-B*
λ	1.3896***	1.1896***	0.9939***
	(0.2186)	(0.2003)	(0.1809)
ϕ	0.9263***	0.9030***	0.9434***
	(0.0186)	(0.0184)	(0.0190)
δ	0.2902***	0.1550*	0.1235*
	(0.0555)	(0.0817)	(0.0785)
k	0.0317	0.0307	0.0522*
	(0.0216)	(0.0254)	(0.0278)
N	1280	1480	1240

Standard errors in parentheses

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

Table 1.6: Estimates of EWA Learning Model for Experts

	(1) *SUPERIOR-B*	(2) *AVERAGE-B*	(3) *INFERIOR-B*
λ	0.4361	0.5736	0.3823
	(0.2580)	(1.1606)	(0.8409)
ϕ	0.8582***	0.8570***	0.9009***
	(0.0239)	(0.0178)	(0.0192)
δ	0.4629	0.1269	0.4862
	(0.3257)	(1.8068)	(1.1328)
k	1.0000***	0.4911	1.0000***
	(0.0006)	(0.3268)	(0.0005)
N	1280	1480	1240

Standard errors in parentheses

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

Chapter 2: Third-party Product Review and Advertising

Abstract

Third parties in reviewing products provide information about the products' attribute qualities, which can influence advertising strategies by the firms whose products have been reviewed. In our model, firms have the option to inform consumers about product attributes through third-party reviews as well as persuade consumers about the importance of particular product attributes. For example, if the information provided by third-party reviewer shows that a product is better in a particular attribute than the priors of consumers, then at first glance the product's manufacturer should guide consumers to read or view the third-party review. One result of our model – one that is counter to the above first-glance intuition – is that a firm may not inform consumers that a particular product attribute is better than expected. Doing so might reduce product differentiation and as a result stiffen price competition if the perceived quality of its product moves closer to its competitors'. This result extends the incomplete disclosure theory in the literature by finding that the optimal disclosure strategy is determined by each product's attribute advantages. However, considering persuasive advertising, the firm may benefit from persuading consumers about the importance of the attribute in which it has a comparative advantage (even in cases in which the quality of this attribute is lower than its competitors' quality of the attribute), especially for firms who cannot produce higher quality for any attribute.

2.1 Introduction

The growing popularity of third-party product reviews means that consumers are better informed when making purchase decisions, especially for experience goods for which consumers use reviews to gain information about product qualities. Most notably in the U.S., *Consumer Reports* ranks products within categories (e.g., refrigerators) and offers salient information about qualities of products for various important attributes. Those product views could influence consumer choice. Simonsohn (2011) among other studies demonstrates that the expert advice provided by *Consumer Reports* has a causal effect on consumer demand; and in particular, consumers who are heterogeneous in preferences for attributes armed with attribute scores provided by the ranking can find the best product for them. On the producer side, third-party product reviews could influence firms' marketing strategies. For instance, some firms directly use the supportive product review reports to disclose quality information,[1] or firms can advertise to offset the influence of unfavorable product reviews.

Chen and Xie (2005) in a model without product attributes develops firms' optimal responses to product reviews, focusing on advertising and pricing strategies. However, products have multiple attributes, and the advertising research, both informative and persuasive, in a setting with multi-attribute products is under explored.

In this paper, we consider products with multiple attributes and develop a new model to explore firms' advertising strategies in response to third-party reviews. In our model, firms have the option to inform consumers about product attributes in general, and about

[1]As an example, Figure A-1 shows a screenshot of Subaru's official website where it intensively highlights several product review reports.

third-party reviews in particular, as well as persuade consumers about the importance of particular product attributes. We address three specific research questions. First, what is the condition for one firm to advertise the result of third-party reviews? Second, which attribute should one firm choose to advertise? Third, with third-party reviews in place, which type of advertising is optimal – informative, persuasive, or both? The answers to these questions rest in part on a firm's attribute advantage of the products it produces.

Our model has two firms and each produces and sells a two-attribute product. A firm has an absolute advantage in a particular product attribute if the perceived quality of the attribute is greater than that of its competitor's. A firm has a comparative advantage in a attribute if the perceived quality of the attribute is comparatively greater than that of its other attribute. Our first finding is that the competitive environment as it relates to absolute and comparative attribute advantage is vital in determining the optimal advertising strategy, namely whether changes in the perceived quality following a third-party review reduces product differentiation or increases it.

Consider the following scenario. Suppose the third-party review reports that one firm's actual quality in one attribute is higher than the consumers' common expected value. There are three possible competition situations with a particular advertising strategy for each situation. (i) If the firm produces a higher quality in this attribute than its competitor, it should inform consumers about this review. However, (ii) if it produces a lower quality in this attribute (yet a higher quality in the other attribute) than its competitor, it should not inform consumers about this review, because consumer knowledge of the actual quality would reduce product differentiation and as a result stiffen price competition. Moreover, (iii) if it produces a lower quality in both attributes than its competitor, it should inform consumers about this review only when it has comparative advantage in this attribute. In the third case, the firm with lower quality in every attribute may lose the entire market and sell zero, especially when it cannot maintain a significant comparative advantage in one attribute.

In the literature on competition and the product quality reporting by competitors to consumers (Board (2009), Levin, Peck and Ye (2009), and Hotz and Xiao (2013)), there is one result called incomplete disclosure, in which a firm does not have the incentive to disclose the quality of its product, even though the actual quality is higher than consumers' prior expectation. We confirm and extend this argument, and make a new contribution by finding that the optimal disclosure strategy is determined by the attribute advantages of the products, in particular we demonstrate the importance of maintaining a significant comparative attribute advantage.

We then develop a model for persuasive advertising which can create consumer loyalty on advertised attributes. As our second finding, a firm's optimal persuasive advertising strategy is to persuade consumers of the importance of the attribute in which it has a comparative advantage, no matter whether it has an absolute attribute advantage. A firm's spending on persuasive advertising is increasing in the magnitude of its comparative advantage. Additionally, we find the firm without absolute advantage in any attribute can get extra benefit from such an optimal persuasive advertising as it can achieve a positive market share more easily.

Finally, to complement our analysis of the disclosure of third-party attribute scores, we examine the case for costly informative advertising. The finding consistently demonstrates that the market competition based on product differentiation plays an essential role in making advertising strategies. Moreover, consumers' purchase decision also affect firms' motivation of engaging in informative advertising, although it is just a necessary condition. In the policy arena, due to a free-rider problem among competitors whose products are being reviewed, we find potential industry-wide under-informative advertising about third-party product reviews.

2.2 Literature

This paper is related to three literature streams. The first examines the decision by firms in oligopolistic markets about whether to disclose product quality to consumers. A key finding in this literature is incomplete disclosure, that is, a firm may not have an incentive to disclose information about the quality of its product, even though its product is better than consumers perception. Board (2009), Levin, Peck and Ye (2009), and Hotz and Xiao (2013) all demonstrate this argument using different modeling techniques. The reasoning is that if revealing information about product quality heightens competition, then a firm might choose to let consumers believe that its product is of lower quality. In the Board (2009) model, consumers have heterogeneous preferences for product quality and the distribution in the consumer population is continuous. There are two firms that both know the quality of both products. He finds the firm with a higher-quality product discloses this quality to consumers, and the firm with the lower-quality product discloses quality if and only if this quality is neither too high nor too low. If the qualities of the two products in the market are close, then price competition is heightened by disclosure about the lower-quality product. Levin, Peck and Ye (2009) assume that consumers are heterogeneous in horizontal preferences but have the same preference for quality. In this two-firm model, the firms do not know the quality of each other's products. Their model characterizes a threshold for disclosure. Hotz and Xiao (2013) assume that consumers are heterogeneous in their preferences for both location and quality, but the heterogeneity in quality is binary – only two types of consumers. Their non-disclosure result requires correlation between consumers' vertical and horizontal preferences.

Our study confirms the incomplete disclosure findings as well as extends it to different competition situations with multiple attributes. By doing so, we add new conditions to the incomplete disclosure.

The second stream is the literature examining product reviews. Dranove and Jin (2010) review the literature of quality disclosure and certification, whereas Dilly (2014) surveys the literature of ratings. Studies such as Jin and Leslie (2003), Chen and Xie (2008), Zhu and Zhang (2010), Simonsohn (2011), Ghose, Ipeirotis and Li (2014) and Luca and Smith (2013) among others evaluate the impact of quality disclosure on consumers' or firms' behavior in various markets. However, to the best of our knowledge a small number of papers (in particular Chen and Xie (2005) and Chen, Liu and Zhang (2012)) in this literature examine how firms report product reviews in their advertising. We contribute to this literature by examining firms' responses in different competition situations to product reviews.

The third stream is the literature on advertising. This is a well-studied area (see Nelson (1974), Becker and Murphy (1993), Bloch and Manceau (1999), Ackerberg (2001), Bagwell (2007), Baye and Morgan (2009) among others). We focus on informative and persuasive advertising with the content related to third-party product reviews. Mayzlin and Shin (2011) find that advertising content can be used to invite consumers to engage in search.

Therefore, referring to third-party product review in the advertising could be one example, and we attempt to model the market outcome of this advertising content.

2.3 The Model

Two firms compete by setting prices and advertising in a two-stage model. In the first stage, two firms decide whether to disclose quality information by informing consumers about a third-party product review; in the second stage, they make advertising decisions and set prices.

Firm $i, i \in \{A, B\}$, offers a two-attribute product. We let (S_1^i, S_2^i) denote the actual quality of each product i's attributes. If a consumer is uninformed about these actual attribute scores, then she believes that the profile of the expected qualities is $((\tilde{S}_1^A, \tilde{S}_2^A), (\tilde{S}_1^B, \tilde{S}_2^B))$. In the unit mass of consumers, each one holds this same prior profile of expected qualities.

A third-party reviewer accurately measures the actual qualities of each of two attributes offered by the two firms $((S_1^A, S_2^A), (S_1^B, S_2^B))$ and publishes them. Either the third-party reviewer, firm A, firm B, or any of these institutions can inform consumers about these attribute scores. If a consumer receives a combination of any of the following messages – the third-party's publication, an informative advertising message from firm A with the third-party's review, or the same informative advertising message from firm B – then she learns the profile of actual attribute scores. If the consumer receives zero messages, then she continues to believe that the profile of the expected qualities is $((\tilde{S}_1^A, \tilde{S}_2^A), (\tilde{S}_1^B, \tilde{S}_2^B))$. A consumer's expected utility of purchasing product i is

$$U_i = \begin{cases} r + w \cdot S_1^i + (1 - w) \cdot S_2^i - P_i & \text{if the consumer is informed,} \\ r + w \cdot \tilde{S}_1^i + (1 - w) \cdot \tilde{S}_2^i - P_i & \text{if the consumer is uninformed} \end{cases} \tag{2.1}$$

where r is the consumer's reservation value, P_i is the price of product i, and w and $1-w$ is the weight the consumer attaches to attribute 1 and 2 respectively. We assume w is uniformly distributed in the consumer population, $w \sim \mathrm{U}[0,1]$.

We apply the usual assumption in models like ours that r is high enough so that each consumer in equilibrium purchases one unit of one product.

In this market with two competitors, consumers who make a purchase decision would consider not just the quality value of a single product but also the quality difference between two products. Therefore, we define the attribute differences in the following two attribute advantages:

Definition 1. *Attribute Advantages*

1. *If $S_j^i > S_j^{-i}$, then firm i has an* ***absolute advantage*** *in attribute j. Compared to firm -i, firm i's attribute j is higher quality;*
2. *If $S_j^i - S_{j'}^i > S_j^{-i} - S_{j'}^{-i}$, then firm i has a* ***comparative advantage*** *in attribute j. Compared to firm -i, firm i has either a bigger advantage or smaller disadvantage in attribute j than in attribute j'.*

Notice that one firm could have absolute advantage in both attributes, but it can have a comparative advantage in only one attribute.

We characterize equilibrium in terms of the differences between two products in the qualities of the two attributes, so we define $\Delta_1 \equiv S_1^A - S_1^B$, $\Delta_2 \equiv S_2^A - S_2^B$, and uninformed consumers' prior expectation of those differences as $\tilde{\Delta}_1 \equiv \tilde{S}_1^A - \tilde{S}_1^B$ and $\tilde{\Delta}_2 \equiv \tilde{S}_2^A - \tilde{S}_2^B$.

To simplify the analysis, without loss of generality, we assume that firm A has an absolute advantage in attribute 1. Moreover, if one firm offers a product that has higher quality in both attributes, without loss of generality, we assume firm A produces this product and that the quality difference between the two products is greatest for attribute 1. In addition, we assume consumers have correct prior expectations about absolute advantages, but are uncertain about the accurate quality differences. We summarize these assumptions as follows.

Assumption 1. *The following holds for two products:*

1. $\Delta_1 \geqslant 0$. *Firm A has absolute advantage in attribute 1.*
2. $\Delta_1 > \Delta_2$. *Firm A has comparative advantage in attribute 1, and firm B has comparative advantage in attribute 2.*[2]
3. $\tilde{\Delta}_1 \geqslant 0$ *iff* $\Delta_1 \geqslant 0$, $\tilde{\Delta}_1 \geqslant \tilde{\Delta}_2$ *iff* $\Delta_1 \geqslant \Delta_2$. *Consumer expectation about the absolute advantage for attribute 1 and the comparative advantages for the two attributes are correct.*

Note that we assume $\tilde{\Delta}_1 \neq \tilde{\Delta}_2$, because we consider $\tilde{\Delta}_1 = \tilde{\Delta}_2$ as a special case and present its analysis in C, which demonstrates that it does not change the result of our main analysis. Therefore, all the core analyses through this paper would be built on Assumption 1.

We begin our formal analysis with the consumers' product purchase decisions. As consumer heterogeneity is modeled by parameter w, an indifferent consumer, $\hat{w}$, enjoys the same expected utility whether purchasing product A or B. That is,

$$r + w\tilde{S}_1^A + (1-w)\tilde{S}_2^A - P_A = r + w\tilde{S}_1^B + (1-w)\tilde{S}_2^B - P_B$$
$$\Rightarrow \hat{w} = \frac{(P_A - P_B) - \tilde{\Delta}_2}{\tilde{\Delta}_1 - \tilde{\Delta}_2} \; (\tilde{\Delta}_1 \neq \tilde{\Delta}_2 \text{ under Assumption 1}). \tag{2.2}$$

[2]This does not exclude the case where firm B has absolute advantage in attribute 2, which is $\Delta_2 < 0$ and $|\Delta_1| < |\Delta_2|$.

This indifferent consumer splits the market demand in a way that consumers with w above (below) the cutoff $\hat{w}$ would purchase product A (B). With the assumption that $\tilde{\Delta}_1 > \tilde{\Delta}_2$, we have $0 \leqslant \hat{w} \leqslant 1$ if and only if $\tilde{\Delta}_2 \leqslant P_A - P_B \leqslant \tilde{\Delta}_1$. Therefore, the demand for each firm's product can be characterized as:

$$\begin{cases} D_A = 0, D_B = 1 & \text{if } P_A - P_B > \tilde{\Delta}_1; \\ D_A = 1 - \hat{w}, D_B = \hat{w} & \text{if } \tilde{\Delta}_2 < P_A - P_B \leqslant \tilde{\Delta}_1; \\ D_A = 1, D_B = 0 & \text{if } P_A - P_B \leqslant \tilde{\Delta}_2. \end{cases} \tag{2.3}$$

We demonstrate in C that the demand function (2.3) also includes the case of Bertrand competition when $\tilde{\Delta}_1 = \tilde{\Delta}_2$. Therefore, it inclusively characterizes all the market share possibilities. [3]

Whether the two firms have incentive to advertise the third-party reviews depends on how this disclosure would influence the consumers' beliefs so that affect their profits. We examine this question by solving firms' profit-maximization problem. In doing so, we normalize the marginal cost of producing the good to zero, and firms' profit functions are simply given by:

$$\pi_i = P_i \cdot D_i. \tag{2.4}$$

2.3.1 The Price Equilibrium

We start our analysis with the second-stage pricing equilibrium given consumers' prior beliefs of product qualities, followed by the first-stage quality disclosure equilibrium in the next subsection.

Given consumers' prior beliefs of product qualities, $\tilde{\Delta}_1$ and $\tilde{\Delta}_2$, the two firms set prices

[3] To further demonstrate that Assumption 1 does not lose the generality, we derive the results of the other case where $\Delta_1 < \Delta_2$ in D.

to maximize their own profits in equation (2.4). Assuming the demand for each firm's product is positive, the best-response functions of the two firms in the price-setting subgame are derived as:

$$
\begin{aligned}
P_A^{BR} &= \frac{1}{2}(P_B + \tilde{\Delta}_1), \\
P_B^{BR} &= \frac{1}{2}(P_A - \tilde{\Delta}_2).
\end{aligned}
\tag{2.5}
$$

It suggests that each firm directly responds to its competitor's price, with the concern of attribute heterogeneity. Then each firm's price is obtained by solving the best-response functions:

$$
\begin{aligned}
P_A^* &= \frac{2\tilde{\Delta}_1 - \tilde{\Delta}_2}{3}, \\
P_B^* &= \frac{\tilde{\Delta}_1 - 2\tilde{\Delta}_2}{3}.
\end{aligned}
$$

To verify the positive demands, we need to further examine the market share that is determined by the indifferent consumer given in equation (2.2). The value of $\hat{w}$ associated with the potential equilibrium prices P_A^* and P_B^* is:

$$
\hat{w}^* = \frac{\tilde{\Delta}_1 - 2\tilde{\Delta}_2}{3(\tilde{\Delta}_1 - \tilde{\Delta}_2)}.
$$

Recall that $\tilde{\Delta}_1 > \tilde{\Delta}_2$ under Assumption 1, then $\hat{w}^* < 1$ holds, which means firm A can always achieve a positive market share. However, firm B can achieve a positive market share if and only if $\hat{w}^* > 0$, which generates the following condition:

$$
\tilde{\Delta}_2 < \tilde{\Delta}_1/2. \tag{2.6}
$$

Otherwise, firm A obtains the entire market and firm B sells zero when $\tilde{\Delta}_1/2 \leqslant \tilde{\Delta}_2 < \tilde{\Delta}_1$.

To sum up, we characterize all the price equilibrium by the following proposition.[4]

[4]The price equilibrium for a Bertrand competition is presented in C, which is just a special case of this

Proposition 2. *When consumers have expectations $\tilde{\Delta}_1$ and $\tilde{\Delta}_2$ following Assumption 1, there exists a unique price equilibrium as:*

$$\begin{cases} P_A^* = (2\tilde{\Delta}_1 - \tilde{\Delta}_2)/3 \text{ and } P_B^* = (\tilde{\Delta}_1 - 2\tilde{\Delta}_2)/3 & \text{if } \tilde{\Delta}_2 < \tilde{\Delta}_1/2; \\ P_A^* = \tilde{\Delta}_2 \text{ and } P_B^* = 0 & \text{if } \tilde{\Delta}_1/2 \leqslant \tilde{\Delta}_2 < \tilde{\Delta}_1. \end{cases} \tag{2.7}$$

Proof. If $\tilde{\Delta}_2 < \tilde{\Delta}_1/2$, both firms can have positive market shares, and the equilibrium is derived in section 2.3.1. We then prove that $P_A^* = \tilde{\Delta}_2$ *and* $P_B^* = 0$ is an equilibrium if the demand for firm B's product is zero. We start from firm B. In order to obtain a positive market share, firm B needs to set a price so that $P_A - P_B > \tilde{\Delta}_2$, that is, $P_B < P_A - \tilde{\Delta}_2$, which is negative when $P_A = \tilde{\Delta}_2$. As firm B is worse off when obtaining a positive market share with a negative profit, it does not have incentive to lower price. Analyzing firm A, it is worse off by reducing its price. If it raises its price, firm A would lose some market share. And the optimal price for firm A when it has partial market is $(2\tilde{\Delta}_1 - \tilde{\Delta}_2)/3$, which is lower than $\tilde{\Delta}_2$ when $\tilde{\Delta}_2 < \tilde{\Delta}_1/2$, so firm A would lose profit when raising price. Therefore, firm A's best response to $P_B^* = 0$ is $P_A^* \tilde{\Delta}_2$. □

Proposition 2 has two relevant implications. First, a firm with an absolute advantage in at least one attribute is able to achieve a positive market share under the price competition. In our model for example, if $\tilde{\Delta}_1 > 0$ and $\tilde{\Delta}_2 < 0$, each firm has one absolute advantage and then each can earn a positive market share. This reflects the importance of maintaining an attribute advantage in a market with product differentiation. Second, a firm that does not have absolute advantage in any attribute may sell zero in the market. In our model, if $\tilde{\Delta}_1/2 \leqslant \tilde{\Delta}_2 < \tilde{\Delta}_1$, then firm B sells zero. Moreover, notice that under this condition firm B does not have any absolute advantage in producing either attribute and the disadvantages

general equilibrium.

in each attribute are close. In other words, firm A would like to occupy the entire market only when it is good at both attributes and its comparative advantage (here in attribute 1) is not too big. The intuition is that firm A can charge a higher price when it has a big comparative advantage, instead of competing for market share. Thus for an inferior firm, B here, if $\tilde{\Delta}_2 > 0$, to have one attribute that is substantially worse than firm A's attribute and the other that is close is better than having both attributes that are equally worse. Furthermore, since $\partial \hat{w}^*/\partial \tilde{\Delta}_1 > 0$ and $\partial \hat{w}^*/\partial \tilde{\Delta}_2 < 0$, firm A would reduce its market share when its advantage in attribute 1 gets bigger, while enlarge its market share when its advantage (disadvantage) in attribute 2 gets bigger (smaller).

Our primary interest is to evaluate the market in which both firms have positive market shares, so we will focus on this competition circumstance in the following analysis.

In the equilibrium where both firms have positive market shares, two firms' equilibrium profits are:

$$\pi_A^* = \frac{(2\tilde{\Delta}_1 - \tilde{\Delta}_2)^2}{9(\tilde{\Delta}_1 - \tilde{\Delta}_2)},$$
$$\pi_B^* = \frac{(\tilde{\Delta}_1 - 2\tilde{\Delta}_2)^2}{9(\tilde{\Delta}_1 - \tilde{\Delta}_2)}.$$

Note that, comparing these equilibrium profits with those of the Bertrand model, firms are better off through quality differentiation since both earn positive profits in this model. Furthermore, those profits are related to consumer expectation of quality differences. As a result, the firms disclose quality information to change consumer expectation, if doing so makes them better off. To analyze the effects on equilibrium profits when consumers change expectations, we turn to evaluate the following derivatives of equilibrium profits:

$$\partial\pi_A^*/\partial\tilde{\Delta}_1 > 0,$$

$$\text{sign}[\partial\pi_A^*/\partial\tilde{\Delta}_2] = \text{sign}[\tilde{\Delta}_2],$$

$$\text{sign}[\partial\pi_B^*/\partial\tilde{\Delta}_1] = \text{sign}[\tilde{\Delta}_1] > 0,$$

$$\partial\pi_B^*/\partial\tilde{\Delta}_2 < 0.$$

These results provide us with a baseline for the analysis of firms' quality disclosure behavior: (1) If one firm has comparative advantage in the attribute (attribute 1 for firm A, and attribute 2 for firm B), the firm would be better off by a larger comparative advantage; (2) If one firm does not have comparative advantage in the attribute (attribute 2 for firm A, and attribute 1 for firm B), the firm would be better off from a larger quality difference, no matter whether it has an absolute advantage or disadvantage in the attribute.

2.4 Quality Disclosure by the Firms

As we established in Proposition 2, the equilibrium prices are functions of the consumer expected quality of the firms' attributes. This works through the positions of the firms' best-response functions. Figure 2.1 illustrates the effect of change in consumer expected quality on the positions of the best-response functions and the equilibrium prices. This figure shows that a change in the consumer expected quality difference of attribute 1 (2) affects the second-stage price equilibrium by shifting firm A's (B's) best-response function. Recall that firm A (B) has comparative advantage in attribute 1 (2).

Interestingly, when either firm increases the size of its comparative attribute advantage, both firms respond by raising prices. Furthermore, when the quality difference in attribute 1 increases, the price of product A increases by more than that of product B; and when the quality difference in attribute 2 decreases, the price of product B increases by more than that of product A. Note that when the quality difference in attribute 2 decreases, the comparative advantage of firm B increases rather than decreases.

Figure 2.1: The Dynamic Price Equilibrium

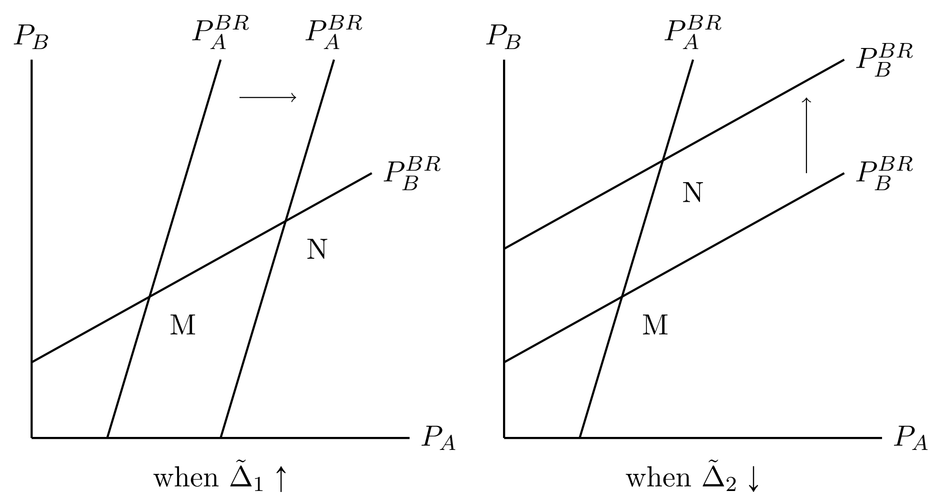

To interpret why the two firms respond to the changes of two attributes differently, we explore the impact on market shares and profits, which using two separate cases that are classified by whether each firm has an absolute attribute advantage are summarized in Table 2.1. Specifically, Case 1 represents the market behavior when each firm has an absolute attribute advantage (firm A in attribute 1 while firm B in attribute 2), whereas in Case 2 firm A has absolute advantage in both attributes so that firm B only has a comparative advantage in attribute 2. In the second case, we call the firm with absolute advantage in both attributes as a **superior firm**, while the firm without any absolute attribute advantage as an **inferior firm**. Comparing two cases, two firms respond in the same direction with respect to prices and demands, but not in profits. In particular, firm A's profit performs differently in two cases.

Table 2.1: Dynamic Equilibrium when Quality Differences Increase

	Price	Market Share	Profit
Case 1: $\tilde{\Delta}_1 > 0,\ \tilde{\Delta}_2 < 0$			
$\tilde{\Delta}_1 \uparrow$	$P_A \uparrow\uparrow$, $P_B \uparrow$	$Q_A \uparrow$, $Q_B \downarrow$	$\pi_A \uparrow\uparrow$, $\pi_B \uparrow$
$\tilde{\Delta}_2 \downarrow$	$P_A \uparrow$, $P_B \uparrow\uparrow$	$Q_A \downarrow$, $Q_B \uparrow$	$\pi_A \uparrow$, $\pi_B \uparrow\uparrow$
Case 2: $\tilde{\Delta}_1 > 0,\ \tilde{\Delta}_2 > 0$			
$\tilde{\Delta}_1 \uparrow$	$P_A \uparrow\uparrow$, $P_B \uparrow$	$Q_A \downarrow$, $Q_B \uparrow$	$\pi_A \uparrow$, $\pi_B \uparrow$
$\tilde{\Delta}_2 \uparrow$	$P_A \downarrow$, $P_B \downarrow\downarrow$	$Q_A \uparrow$, $Q_B \downarrow$	$\pi_A \uparrow$, $\pi_B \downarrow$

2.4.1 Case 1: $\tilde{\Delta}_1 > 0,\ \tilde{\Delta}_2 < 0$

Each firm has an absolute attribute advantage: firm A (B) has absolute advantage in attribute 1 (2); that is, $\tilde{S}_1^A > \tilde{S}_1^B$ and $\tilde{S}_2^A < \tilde{S}_2^B$. We frequently observe this case. For example, a car has many attributes such as driving comfort, safety, energy efficiency and so on, but usually one company does not produce a car that is best in all attributes. That is, one company produces a car that is most comfortable, while another manufactures a car that is safest. Yet another manufacturer's auto is most efficient. Within the context of our two-firm analysis, in this market setting the following analysis shows that both firms would be better off from disclosure of an attribute if it results in a higher quality difference. A firm is made better off whether or not it has an absolute advantage in the attribute.

This means that a firm with an absolute attribute disadvantage would not want to reveal to consumers that its attribute is better than expected.

We make three points about this result. First, since the equilibrium condition, as expressed in equation (2.6) is satisfied, both firms maintain positive market shares in the equilibrium. This is consistent with the reality that one firm can achieve some market demands when its product has higher quality in at least one attribute.

Second, a bigger quality difference in either attribute increases both firms' profits. Since $\tilde{\Delta}_1 > 0$ and $\tilde{\Delta}_2 < 0$, the quality difference in attribute 1 (2) enlarges when $\tilde{\Delta}_1$ increases ($\tilde{\Delta}_2$ decreases). As shown in Figure 2.1, both firms would respond by raising prices. As a result, the change of equilibrium profits can be shown by the derivatives as the following:

$$\partial \pi_A^* / \partial \tilde{\Delta}_1 > 0,\ \ \partial \pi_A^* / \partial \tilde{\Delta}_2 < 0;$$

$$\partial \pi_B^* / \partial \tilde{\Delta}_1 > 0,\ \ \partial \pi_B^* / \partial \tilde{\Delta}_2 < 0.$$

When the quality difference in either attribute increases, $\tilde{\Delta}_1$ increases or $\tilde{\Delta}_2$ decreases, both firms' profits increase.

Third, increasing the magnitude of a firm's absolute attribute advantage has a larger positive impact on its profit than increasing the absolute disadvantage of its other attribute. With an increase in absolute advantage or absolute disadvantage, a firm increases the price of its product. As shown in Figure 2.1, when a firm's absolute attribute advantage increases, its price increases by more than its competitor's. Additionally, a bigger absolute attribute advantage increases demand, whereas a bigger disadvantage reduces demand.

A firm producing a low-quality product might not improve its quality because doing so means intensifying the competition with the producer of a high-quality product. In the context of our model, for a similar reason, we demonstrated that a firm might not inform consumers that the quality of one of the attributes of its product is higher than expected.

2.4.2 Case 2: $\tilde{\Delta}_1 > \tilde{\Delta}_2 > 0$

Firm A has absolute advantages in both attributes, $S_1^A > S_1^B$, $S_2^A > S_2^B$; firm A produces the superior product and firm B an inferior one. For example, product A could be a luxury car, which is better in all attributes than product B, an economy car. Based on the condition expressed in equation (2.6), the inferior firm B sells its product only if $\tilde{\Delta}_2 < \tilde{\Delta}_1/2$. This inequality means that firm B has a substantial comparative advantage in attribute 2. The following analysis is based on this condition and reveals that two firms have different motivations to report the attribute qualities.

To evaluate whether firms choose to disclose quality, we start from the derivatives of profits with respect to quality differences:

$$\partial \pi_A^*/\partial \tilde{\Delta}_1 > 0,\ \ \partial \pi_A^*/\partial \tilde{\Delta}_2 > 0;$$

$$\partial \pi_B^*/\partial \tilde{\Delta}_1 > 0,\ \ \partial \pi_B^*/\partial \tilde{\Delta}_2 < 0.$$

These inequalities indicate that both firms earn higher profits when the quality difference in attribute 1 ($\tilde{\Delta}_1$) increases. However, only firm A earns higher profit when the quality difference in attribute 2 ($\tilde{\Delta}_2$) increases. That is, the superior firm is always better off with higher quality differences, no matter in which attribute; whereas the profit impact for the inferior firm is related to its comparative attribute advantage: when it does not have any absolute attribute advantage, a bigger comparative attribute advantage can be profitable.

Based on this reasoning, if the actual quality of the superior product is greater than the consumer expected quality, the firm that produces the superior product completely discloses its quality, whereas the firm producing the inferior product incompletely discloses. Specifically, the inferior firm reports a better quality only when doing so can increase its comparative attribute advantage. For firm A, if its actual quality is better than the consumers' prior belief, it discloses. For firm B, if its actual quality of attribute 1 is better than the consumers' prior belief (but still lower than firm A's), it will not disclose; but if its attribute 2 is better than the consumer prior belief (but still lower than firm A's), it will disclose. That is, the superior firm prefers to reveal to consumers that it is better than initially expected, but the inferior firm would prefer to reveal that it is better than initially believed in only the attribute in which it has a comparative advantage.

To further understand this result, we examine changes in the quality of each of the two attributes. If the quality difference in attribute 1 increases, both firms respond by raising prices and the superior firm A increases by more. However, firm A would lose some market demand in this process. As a result, firm A's profit would not increase very much. If the quality difference in attribute 2 increases, both firms respond by lowering prices, and firm B reduces its price by more. The intuition is that the only strategy firm B can use when its comparative attribute advantage shrinks is to lower price dramatically. As a response, firm A also lowers price but it can still benefit from a bigger attribute advantage through more market demand. In the end, firm A's profit would increase while firm B's profit would fall. This is a "always-losing" situation for firm B, who needs to give up price as well as market demand.

2.4.3 Summary of General Results

We observe the following implications about firms' motivations of reporting the results of a third-party product review.

If one firm produces an attribute with higher quality than consumer prior expectation, the firm does not always have incentive to disclose this higher quality. The optimal disclosure strategy depends on the attribute advantage of each firm. Specifically, there are three cases: (i) the firm has absolute advantage in this attribute, (ii) the firm does not have absolute advantage but have comparative advantage in this attribute, and (iii) the firm does not have either advantage in this attribute. We find that the firm would like to report the actual higher quality in the first two cases, but not to report in the third case. This finding extents the literature of incomplete disclosure by characterizing the market conditions for disclosure.

Moreover, we separately examine firms' disclosure strategy in two market circumstances. In the first circumstance, each firm has an absolute advantage in one attribute. Although a bigger absolute advantage for one firm means a bigger disadvantage for the other firm, two firms have consistent incentives: to advertise a bigger advantage or disadvantage. Specifically, (1) if one firm has absolute advantage in one attribute, it advertises this attribute when it's actual quality is higher than consumers' prior expectation; (2) if one firm has absolute disadvantage in one attribute, it advertises this attribute when it's actual quality is lower than consumers' prior expectation. In the second circumstance where one firm has absolute advantages in all attributes, two firms have conflict incentives: the superior firm advertises any attribute whose actual quality is higher than consumers' prior expectations, while the other inferior firm advertises a higher quality only for the attribute with comparative advantage. For the other attribute without comparative advantage, the inferior firm advertises it when it's actual quality is lower than consumers' prior expectation.[5]

In addition, if the product reviews show that the inferior product does not have significant comparative advantage, the demand for this product could fall to zero.

2.4.4 Numerical Analysis

To further illustrate the results from our model, we construct three numerical examples, with the parameter values reported in Table 2.2. We also report the equilibrium outcomes for those three examples in Table 2.3. Note that both firms always share the market in these examples. We examine the effect of changes in parameter values in Figures 2.2 through 2.4.

Table 2.2: Parameters in Numerical Examples

	S_1^A	S_2^A	S_1^B	S_2^B	Δ_1	Δ_2
Example 1	$[0,1]$	$1/4$	0	$3/4$	$[0,1]$	$-1/2$
Example 2	$[1/2,1]$	$1/4$	0	0	$[1/2,1]$	$1/4$
Example 3	$1/2$	$1/4$	0	$[0,1]$	$1/2$	$[1/4,-3/4]$

[5]However, it is rare in reality that one firm advertises the attribute with a lower actual quality. Firms may have other concerns when ignoring this opportunity to be more profitable from a deeper product differentiation. This points out a new direction for future research.

Table 2.3: Equilibrium Outcomes in Numerical Examples

	Prices		Demands		Profits	
	p_A	p_B	q_A	q_B	π_A	π_B
Example 1	$[1/6, 5/6]$	$[1/3, 2/3]$	$[1/3, 5/9]$	$[2/3, 4/9]$	$[0, 1/2]$	$[2/9, 2/7]$
Example 2	$[1/4, 3/5]$	$[0, 1/6]$	$[1, 7/9]$	$[0, 2/9]$	$[1/4, 4/9]$	$[0, 1/27]$
Example 3	$[1/4, 3/5]$	$[0, 2/3]$	$[1, 1/2]$	$[0, 1/2]$	$[1/4, 2/9, 1/4]$	$[0, 1/3]$

Examples 1 and 2 show the equilibrium outcomes if firm A increases the quality of attribute 1. In Example 1 each firm has absolute attribute advantage; in Example 2 firm A has an absolute advantage in both attributes. Figure 2.2 and 2.3 depicts the relationship between the equilibrium values and S_1^A. With an increase in S_1^A, both firms raise prices and the price of product A increases at a faster rate. However, because of the quality difference in attribute 2, the market demand of firm A (B) increases (decreases) when $\tilde{\Delta}_2 < 0$ while decreases (increases) when $\tilde{\Delta}_2 > 0$. As a result, although both firms could enjoy higher profits in two cases, firm A's profit goes up more than firm B's. Note that we generate the same outcomes if we hold S_1^A constant and decrease S_1^B, so we do not report those outcomes. These two examples show that firm A would like to disclose its quality of attribute 1 if it is higher than the consumer prior expectation, while firm B may not disclose its quality of attribute 1 if it is higher than the consumer prior.

Example 3 examines changes in the equilibrium outcomes if firm B increases the quality of attribute 2. Figure 2.4 depicts the results that with this change both firms raise prices and the price of product B is higher. As a result, the market demand of product A (B) increases (decreases). While firm A's profit increases, firm B's profit is decreasing if $S_2^B < S_2^A$ and increasing if $S_2^B > S_2^A$. Note that we generate the same outcomes if S_2^B is constant and S_2^A decreases, and gain we do not report those outcomes. This example reflects that firm B would like to disclose its quality of attribute 2 if it is higher than consumer's prior expectation, while firm A has incentive to disclose high quality of attribute 1 only if it has absolute advantage in this attribute.

In sum, these numerical outcomes reflect general results from our model that the optimal disclosure strategy depends on the absolute and comparative attribute advantages of each product.

2.5 Persuasive Advertising

In this section we turn to explore the effect of third-party product reviews on persuasive advertising.

A large number of studies on advertising have examined persuasive adverting (see Bagwell, 2007). In our analysis, a firm's persuasive advertising aims to influence consumer preferences by emphasizing the importance of certain attributes. For example in the car market, BMW's slogan is "the Ultimate Driving Machine" which is an attempt to push consumers toward driving attributes as opposed to ones involving comfort, whereas Mercedes pushes the opposite by the label of comfort and luxury over driving attributes. By doing so, those car companies not only differentiate their products from competitors' but also emphasize the importance of certain attributes.

Consumers who are persuaded by this type of advertising change preferences for the products' two attributes. In our model, inspired by the endogenous brand loyalty model developed by Baye and Morgan (2009), a consumer who is persuaded about the value of one attribute, changes her preferences in an extreme manner by placing a weight of 1 on that attribute (and zero on the other).

Each firm selects one attribute for which it attempts to persuade consumers. We let λ_{ij} denote the persuasive advertising expenditure of firm i on attribute j. So the total advertising on attribute j is $\lambda_j = \lambda_{Aj} + \lambda_{Bj}$. The advertising cost is $a \cdot \lambda_{ij}$, where a presents the marginal cost of advertising.

We assume that the proportion of consumers who are loyal to attribute j, denoted by β_j, is determined by the following response function:

$$\beta_j(\lambda_j, \lambda_{-j}) = \delta \frac{\lambda_j}{\lambda_j + \lambda_{-j}}, \tag{2.8}$$

where $\delta \in (0, 1)$ denotes the proportion of consumers who can be influenced by persuasive advertising. We adopt this response function from Baye and Morgan (2009). It reflects the reality that whether a consumer is loyal to one attribute depends on the relative intensity of advertising on the attribute. Furthermore, it is plausible to assume that if neither firm engages in persuasive advertising, then no consumer is loyal to either attribute, $\beta_j(0, 0) \equiv 0$.

For the proportion of consumers who are loyal to attribute j, their utility of purchasing product i is derived from equation (2.1) as

$$U_i = r + S_j^i - P_i. \tag{2.9}$$

These loyal consumers select the product which yields the greater utility.

Equation (2.1) expresses the utility function for consumers who are not loyal to one particular attribute. Therefore, the remaining $1 - \delta$ consumers who are not convinced by the ads maintain the original weights and will make choices based on the prior cutoff $\hat{w}(P_A, P_B) = \frac{(P_A - P_B) - \Delta_2}{\Delta_1 - \Delta_2}$, as stated in equation (2.2).

2.5.1 The Equilibrium

As we build our analysis of persuasive advertising on the basic model setting of Section 2.3, we continue to make Assumption 1.

We begin our equilibrium analysis with the purchase behavior of loyal consumers, which is characterized by the following lemma.

Lemma 3. *When two firms both achieve positive market shares by setting prices, consumers who are loyal to attribute 1 purchase product A and consumers who are loyal to attribute 2 purchase product B.*

Proof. Given the utility function in equation (2.9), the loyal consumer compares the utility from two products and makes purchase decisions based on the following rule:

$$\begin{cases} \text{Purchase } A, & \text{if loyal to attribute 1 and } P_A - P_B < \Delta_1; \\ \text{Purchase } B, & \text{if loyal to attribute 1 and } P_A - P_B > \Delta_1; \\ \text{Purchase } A, & \text{if loyal to attribute 2 and } P_A - P_B < \Delta_2; \\ \text{Purchase } B, & \text{if loyal to attribute 2 and } P_A - P_B > \Delta_2. \end{cases}$$

Since $0 < \hat{w} < 1 \Leftrightarrow \Delta_2 < P_A - P_B < \Delta_1$, in order to achieve a positive market share from either loyal consumers or the remaining consumers, both firms would set prices in a way so that consumers who are loyal to attribute 1 would purchase product A while consumers who are loyal to attribute 2 would purchase product B. $\square$

Since product A (B) has comparative advantage in attribute 1 (2), it implies that in the equilibrium loyal consumers purchase the product with comparative advantage in their loyal attribute, no matter whether the product has absolute advantage in that attribute. Comparative attribute advantage affects the loyalty created by persuasive advertising. Therefore, firms have the motivation to improve their comparative attribute advantages. Moreover, based on the purchase behavior described by Lemma 3, each firm chooses an attribute for persuasive advertising. The following lemma establishes necessary conditions for the equilibrium advertising.

Lemma 4. *The equilibrium expenditures for persuasive advertising, λ_{ij}, is chosen from the following set: $\{\lambda_{A1} \geqslant 0, \lambda_{B2} \geqslant 0, \lambda_{A2} = \lambda_{B1} = 0\}$. That is, if in equilibrium a firm engages in persuasive adverting, then firm A chooses attribute 1 while firm B chooses attribute 2 for the advertising.*

Proof. We can prove this lemma by contradiction. Given the equilibrium prices, if firm A chooses attribute 2 instead of attribute 1, it induces some consumers to be loyal to attribute 2. Because Lemma 3 shows that consumers who are loyal to attribute 2 always purchase product B, the more firm A spends on advertising, the less demand for product A. Therefore, the optimal advertising on attribute 2 is zero. Similarly, the optimal advertising for firm B on attribute 1 is zero. As a result, the possible positive advertising occurs only when firm A chooses attribute 1 while firm B chooses attribute 2. □

Lemma 4 provides a guideline for firms to choose attributes for advertising. We summarize this implication in the following corollary.

Corollary 1. *In equilibrium, if a firm engages in persuasive advertising, then it chooses the attribute with which it has a comparative advantage, regardless of whether it has absolute advantage in this attribute.*

Given the advertising choices and consumers purchase decisions characterized in Lemma

3 and 4, the demands for the two products are:

$$Q_A = \delta \frac{\lambda_{A1}}{\lambda_{A1} + \lambda_{B2}} + (1-\delta)(1-\hat{w}),$$
$$Q_B = \delta \frac{\lambda_{B2}}{\lambda_{A1} + \lambda_{B2}} + (1-\delta)\hat{w}.$$

As in the basic model, we normalize marginal production cost to be zero, then the profit function faced by each firm is given by sale revenue minus advertising cost, which are:

$$\pi_A = P_A \cdot Q_A - a \cdot \lambda_{A1},$$

$$\pi_B = P_B \cdot Q_B - a \cdot \lambda_{B2}.$$

In this setting, each firm simultaneously sets price and advertising expenditure to maximize its own profit. The first-order conditions of profit-maximization are computed as:

$$\begin{aligned}
\frac{\partial \pi_A}{\partial P_A} &= \frac{(1-\delta)(\Delta_1 - 2P_A + P_B)}{\Delta_1 - \Delta_2} + \delta \frac{\lambda_{A1}}{\lambda_{A1} + \lambda_{B2}} = 0 \\
\frac{\partial \pi_B}{\partial P_B} &= \frac{(\delta-1)(\Delta_2 - P_A + 2P_B)}{\Delta_1 - \Delta_2} + \delta \frac{\lambda_{B2}}{\lambda_{A1} + \lambda_{B2}} = 0 \\
\frac{\partial \pi_A}{\partial \lambda_{A1}} &= \delta \frac{P_A \lambda_{B2}}{(\lambda_{A1} + \lambda_{B2})^2} - a = 0 \\
\frac{\partial \pi_B}{\partial \lambda_{B2}} &= \delta \frac{P_B \lambda_{A1}}{(\lambda_{A1} + \lambda_{B2})^2} - a = 0
\end{aligned} \tag{2.10}$$

From equations (2.10), the unique equilibrium prices and advertising expenditures are solved as:

$$\begin{aligned}
P_A^* &= \frac{\Delta_1(2-\delta)-\Delta_2}{(1-\delta)(3-\delta)},\\
P_B^* &= \frac{\Delta_1-\Delta_2(2-\delta)}{(1-\delta)(3-\delta)},\\
\lambda_{A1}^* &= \frac{\delta[\Delta_1-\Delta_2(2-\delta)][\Delta_1(2-\delta)-\Delta_2]^2}{a(1-\delta)(3-\delta)^3(\Delta_1-\Delta_2)^2},\\
\lambda_{B2}^* &= \frac{\delta[\Delta_1(2-\delta)-\Delta_2][\Delta_1-\Delta_2(2-\delta)]^2}{a(1-\delta)(3-\delta)^3(\Delta_1-\Delta_2)^2}
\end{aligned} \tag{2.11}$$

In this equilibrium, loyal consumers choose products following the rule specified in Lemma 3, whereas the market demand from neutral consumers is determined by the cutoff given in equation (2.2). With prices P_A^* and P_B^*, the value of $\hat{w}$ is:

$$\hat{w}^* = \frac{1}{3-\delta}\left[1-\frac{(1-\delta)\Delta_2}{\Delta_1-\Delta_2}\right].$$

Recall that to guarantee both firms have positive market shares in this equilibrium, we need $0 < \hat{w}^* < 1 \Leftrightarrow \Delta_2 < P_A - P_B < \Delta_1$. Therefore, we can characterize the equilibrium condition for positive market shares as

$$\Delta_2 < \Delta_1/(2-\delta). \tag{2.12}$$

Otherwise firm B would sell to zero if $\Delta_1/(2-\delta) \leqslant \Delta_2 < \Delta_1$. The following proposition summarizes all the equilibrium of persuasive advertising.

Proposition 3. *With quality differences Δ_1 and Δ_2 that are characterized by Assumption 1, there exists a unique equilibrium as:*

$$\begin{cases}(P_A^*, P_B^*, \lambda_{A1}^*\lambda_{B2}^*) \textit{ given in (2.11)}, & \textit{if } \Delta_2 < \Delta_1/(2-\delta);\\ P_A^* = \Delta_2, P_B^* = 0, \lambda_{A1}^* = \lambda_{B2}^* = 0, & \textit{if } \Delta_1/(2-\delta) \leqslant \Delta_2 < \Delta_1.\end{cases}$$

Proof. If $\tilde{\Delta}_2 < \tilde{\Delta}_1/(2-\delta)$, both firms can have positive market shares, and the equilibrium is derived in above analysis. We then prove that if firm B sells zero to the unpersuaded consumers, then $P_A^* = \tilde{\Delta}_2, P_B^* = 0, \lambda_{A1}^* = \lambda_{B2}^* = 0$ is an equilibrium. In this case firm A can set $P_A^* = P_B^* + \Delta_2$ to sell to all of the unpersuaded consumers. Next, Lemma 3 demonstrates that if $P_A^* = \Delta_2, P_B^* = 0$, then consumers who are loyal to attribute 1 purchase product A while consumers who are loyal to attribute 2 are indifferent between two products. Considering persuasive advertising, with the the assumption that indifferent consumers purchase product A, the persuasive advertising of firm B is ineffective and costly. Therefore, firm B is better off from zero adverting. Then firm A can sell to the entire market even without advertising. Both firms have incentive to set zero advertising. $\square$

This equilibrium with persuasive advertising has several implications. First, we observe that under condition (2.12), each firm has a positive market share, all equilibrium prices and advertising expenditures are positive, suggesting that both firms have incentive to engage in persuasive advertising. Specifically, firm A advertises attribute 1 and firm B advertises attribute 2. Furthermore, condition (2.12) holds for either $\Delta_2 > 0$ or $\Delta_2 < 0$, which implies that it is profitable for firm B to persuade consumers about attribute 2 even if firm A has absolute advantage in both attributes. In other words, firm B persuades consumers about the importance of the attribute in which it has a comparative advantage, even if firm A's product is of higher quality in that attribute. For example, BMWs may drive better as well as more comfortable, but if Mercedes has a comparative advantage in comfort, it will persuade consumers about the value of comfort.

In addition, the persuasive advertising to some extent protects an inferior firm. When some consumers can be influenced by persuasive advertising, represented by δ, the probability that the inferior firm not participating in the market declines. At an extreme, if all consumers can be influenced, the inferior firm can always achieve a positive market share. We present this result in the following corollary.

Corollary 2. *As the proportion of consumers who are influenced by persuasive advertising, δ, increases, an inferior firm is more likely to participate in the market. In the limit, if $\delta = 1$, the inferior firm obtains positive sales under Assumption 1.*

Proof. If $\Delta_2 < \Delta_1/(2-\delta)$, then the inferior firm, firm B, can achieve a positive market share. Since $\Delta_1/(2-\delta)$ increases as δ increases, the constrain reduces. When $\delta = 1$, we have $\Delta_2 < \Delta_1$, which always holds under Assumption 1. □

2.5.2 Comparative Static Analysis

In this section we examine the effect of changes in the degree of product differentiation on the firms' advertising decisions. When either firm improves its comparative attribute advantage, both firms spend more on advertising, no matter whether one firm has absolute attribute advantage. Table 2.4 pictures this result emphasizing the importance of increasing comparative attribute advantage. This table also shows that with the greater comparative advantage, the firms raise prices. However, the changes in sales depend on particular market circumstances. Both the effect on prices and demands perform consistently as the basic model.

To emphasize this effect of increased product differentiation, we create three numerical examples with parameter values in Table 2.5. We set the attribute values to be the same as those for the basic model of pure price competition, reported in Table 2.2. In these three examples, both firms always obtain positive market demands. We present the equilibrium outcomes in Figure 2.5 to 2.7.

Table 2.4: Comparative Analysis of the Persuasive Advertising

	Price	Advertising	Market Share
$\Delta_1 \uparrow$	$P_A \uparrow, P_B \uparrow$	$\lambda_{A1} \uparrow, \lambda_{B2} \uparrow$	$Q_A \uparrow\downarrow, Q_B \downarrow\uparrow$*
$\Delta_2 \downarrow$	$P_A \uparrow, P_B \uparrow$	$\lambda_{A1} \uparrow, \lambda_{B2} \uparrow$	$Q_A \downarrow, Q_B \uparrow$

*Similar as the basic model, $Q_A \uparrow$, $Q_B \downarrow$ when $\tilde{\Delta}_1 > 0,\ \tilde{\Delta}_2 < 0$,

while $Q_A \downarrow$, $Q_B \uparrow$ when $\tilde{\Delta}_1 > 0,\ \tilde{\Delta}_2 > 0$.

Table 2.5: Parameters in Numerical Examples

	δ	a	S_1^A	S_2^A	S_1^B	S_2^B	Δ_1	Δ_2
Example 1	1/2	1/10	$[0, 1]$	1/4	0	3/4	$[0, 1]$	$-1/2$
Example 2	1/2	1/10	$[1/2, 1]$	1/4	0	0	$[1/2, 1]$	1/4
Example 3	1/2	1/10	1/2	1/4	0	$[0, 1]$	1/2	$[1/4, -3/4]$

In Examples 1 and 2 the quality of firm A's attribute 1 changes. In Example 1 each firm has an absolute attribute advantage; in Example 2 firm A has an absolute advantage in both attributes. Figures 2.5 and 2.6 depict the results in which both firms increase advertising expenditure in response to an increase in Firm A's quality of attribute 1. Note that firm A's response is greater than B's. Both firms raise prices. With regard to sales, if $\tilde{\Delta}_2 < 0$, then the demand of firm A (B) increases (decreases); if $\tilde{\Delta}_2 > 0$, then the demand for firm A (B) (increases). Note that if the value of S_1^A is constant and the value of S_1^B decreases, then we would observe the same changes in the market outcomes.

In Example 3, the quality of firm B's attribute 2 changes. Figure 2.7 depicts the result that if S_2^B increases, both firms increase advertising expenditure, with a greater increase in firm B's advertising. Meanwhile, both firms raise prices, and the market demand of firm A (B) decreases (increases). Note that if the value of S_2^B is constant and the value of S_2^A decreases, then we observe the same changes in the market outcomes.

In sum, our numerical analysis stresses our general result that if one firm's comparative attribute advantage goes up, whether or not one firm has an absolute advantage, then in equilibrium both firms would spend more on persuasive advertising.

2.6 Informative Advertising

In Section 2.4 we examined whether firms, even at a zero cost, are willing to report third-party product reviews to consumers. In this section, we continue with this analysis by examining costly informative advertising intended to inform consumers about third-party product reviews. In reality, we observe this type of informative advertising in all forms of media, in particular in online platforms. See for example Subaru's advertising shown in Figure A-1 of third-party reviews of two of its product line attributes.

In this section's model, we assume that the costly informative advertising would accurately report both firms' attribute quality, which would function like the third-party reviews by only informing consumers the existing product differentiation without influencing consumers' preference. With this setting, consumers who have accessed the product reviews or seen the ads would correctly update their beliefs from $\tilde{\Delta}_1$ and $\tilde{\Delta}_2$ to Δ_1 and Δ_2. Consumers who are not reached by either media just maintain priors. We denote that ρ fraction of total consumers would read the product reviews before shopping, and the advertising intensity from firm i is λ_i, which represents the fraction of consumers who are exposed by its ads. The cost function of advertising is $A(\lambda_i) = a \cdot \lambda_i$, where a presents the marginal cost of advertising.

With the existence of third-party product reviews, two firms need to compete for two groups of consumers: one group is informed before purchase, through product reviews or informative advertising; and the other group is not informed by either. The market demands from these two groups are divided by the indifferent consumer in each group, $\hat{w}^I$ and $\hat{w}^N$, for informed and uninformed consumers, respectively. The values of these two cutoffs are determined by consumers' beliefs of attributes, which are calculated based on equation (2.2) as:

$$\begin{aligned} \hat{w}^I &= \frac{(P_A - P_B) - \Delta_2}{\Delta_1 - \Delta_2}, \\ \hat{w}^N &= \frac{(P_A - P_B) - \tilde{\Delta}_2}{\tilde{\Delta}_1 - \tilde{\Delta}_2}. \end{aligned}$$

Before we present the formal equilibrium analysis, we can make two predictions from the outcome of quality report in section 2.4: (1) If each firm has an absolute attribute advantage, both firms would like to inform consumers if the actual product differentiation is bigger. Since two firms have common interest, there could be a free riding problem and the informative advertising works like a public good. (2) If one (superior) firm has absolute advantage in all attributes, it has incentive to inform consumers about bigger product differentiation, while the inferior firm would consider informative advertising only when it has a bigger comparative attribute advantage.

We begin our equilibrium analysis with consumer demands. Recall that this is a two-segment market with informed and uninformed consumers, and if we denote the fraction of total informed consumers as β, then its value is $\beta \equiv 1 - (1-\rho)\cdot(1-\lambda_A)\cdot(1-\lambda_B)$. As a result, we can summarize the proportion of two group consumers as:

$$\begin{cases} \beta: & \text{informed, cutoff of demand is } \hat{w}^I \\ 1-\beta: & \text{uninformed, cutoff of demand is } \hat{w}^N \end{cases}$$

Since consumers with w above (below) the cutoff would purchase product A (B), the total demand for each product can be calculated as:

$$Q_A = 1 - \beta\hat{w}^I - (1-\beta)\hat{w}^N,$$

$$Q_B = \beta\hat{w}^I + (1-\beta)\hat{w}^N.$$

Having derived the market demands, we can write down each firm's profit function. Consistently, we normalize marginal production cost to be zero, so that profit function is given by sale revenue minus advertising cost, which is

$$\pi_A = P_A \cdot Q_A - A(\lambda_A) = P_A \cdot (1 - \beta\hat{w}^I - (1-\beta)\hat{w}^N) - a \cdot \lambda_A, \tag{2.13}$$

$$\pi_B = P_B \cdot Q_B - A(\lambda_B) = P_B \cdot (\beta\hat{w}^I + (1-\beta)\hat{w}^N) - a \cdot \lambda_B. \tag{2.14}$$

We then evaluate firms' equilibrium behavior for informative advertising. In our setup, two firms simultaneously set price and advertising intensity to maximize their own profit, so the potential Nash equilibrium should contain the optimal price and advertising intensity each firm would choose given the competitor's action. To derive the equilibrium, the first-order conditions of profit-maximization are computed as:

$$\frac{\partial \pi_A}{\partial P_A} = -1 + (2P_A - P_B)(\beta M + (1-\beta)\tilde{M}) - \beta N - (1-\beta)\tilde{N} = 0 \tag{2.15}$$

$$\frac{\partial \pi_B}{\partial P_B} = (P_A - 2P_B)(\beta M + (1-\beta)\tilde{M}) - \beta N - (1-\beta)\tilde{N} = 0 \tag{2.16}$$

$$\frac{\partial \pi_A}{\partial \lambda_A} = P_A(1-\rho)(1-\lambda_B)(\hat{w}^N - \hat{w}^I) - a = 0 \tag{2.17}$$

$$\frac{\partial \pi_B}{\partial \lambda_B} = P_B(1-\rho)(1-\lambda_A)(\hat{w}^N - \hat{w}^I) - a = 0 \tag{2.18}$$

$$\text{where, } M \equiv \frac{1}{\Delta_1 - \Delta_2}, N \equiv \frac{\Delta_2}{\Delta_1 - \Delta_2}$$

In general, we cannot obtain an analytical solution from above first-order conditions, but they provide some features of the possible equilibria. To characterize the equilibria, we first investigate an extreme case of full advertising, which is illustrated in the following lemma.

Lemma 5. *If informative advertising is costly, neither firm would inform all consumers about a third-party product review. That is, if $a > 0$, in any equilibrium $\lambda_A, \lambda_B \neq 1$.*

Proof. We prove this lemma by contradiction. Suppose $\lambda_A^* = 1$ belongs to one equilibrium, then it must hold for all the first-order conditions. By equation (2.18), we should have $P_B(1-\rho)(1-\lambda_A^*)(\hat{w}^N - \hat{w}^I) - a = 0$, that is, $a = 0$. This is contradicted with the assumption of costly informative advertising. Therefore, $\lambda_A^* = 1$ does not belong to any equilibrium. Similarly, we can prove $\lambda_B^* = 1$ does not belong to any equilibrium. □

We have demonstrated in section 2.4 that a firm has incentive to disclose attribute quality only under some conditions, even when it is costless to disclose. Lemma 5 further reveals the fact that full disclosure is never optimal when it is costly to disclose. This potential industry-wide under-informative advertising could due to a free-rider problem among competitors whose products are being reviewed, since they can always be better off by reducing advertising if a lot consumers have been informed.

As our primary interest, we then turn to find the conditions under which firms have incentive to undertake the informative advertising. We start from the original price competition. Without the concern of informative advertising, $\lambda_A = 0$ and $\lambda_B = 0$, two firms just compete on prices, and there is one unique equilibrium that can be derived from equations (2.15) and (2.16):

$$P_A^o = \frac{2 + \rho N + (1-\rho)\tilde{N}}{3[\rho M + (1-\rho)\tilde{M}]}, \; P_B^o = \frac{1 - \rho N - (1-\rho)\tilde{N}}{3[\rho M + (1-\rho)\tilde{M}]}.$$

The equilibrium prices are determined by product differentiation. In particular, since $M > 0$ based on Assumption 1, $P_A^o > P_B^o \Leftrightarrow \rho N + (1-\rho)\tilde{N} > -1/2$. Given this price equilibrium (P_A^o, P_B^o), two firms' profits are computed through equations (2.13) and (2.14) as π_A^o and π_B^o, and the market cutoffs for uninformed and informed consumers are denoted as $\hat{w}_o^N$ and $\hat{w}_o^I$. Under this situation, firm i's marginal profit of starting informative advertising is

$$\frac{\partial \pi_i^o}{\partial \lambda_i} = P_i^o(1-\rho)(\hat{w}_o^N - \hat{w}_o^I) - a$$

Firm i has incentive to run informative advertising if $\partial \pi_i^o / \partial \lambda_i > 0$, which gives us the following condition:

$$P_i^o(1-\rho)(\hat{w}_o^N - \hat{w}_o^I) > a \tag{2.19}$$

Since P_i^o, $\hat{w}_o^N$ and $\hat{w}_o^I$ are all expressions of Δ_j and $\tilde{\Delta}_j$, this inequation is characterized by attribute quality differences with ρ and a as parameters. Therefore, the product differentiation determines whether each firm starts to run the informative advertising. We present its implication in the following proposition.

Proposition 4. *Suppose there are uninformed consumers, $\rho < 1$. If a firm engages in informative advertising, then $\hat{w}_o^N > \hat{w}_o^I$ and P_i^o is sufficiently large.*

Proof. Since $a > 0$, if condition (2.19) holds, we should have $\hat{w}_o^N - \hat{w}_o^I > 0$. Meanwhile, $P_i^o(1-\rho)(\hat{w}_o^N - \hat{w}_o^I)$ increases with respect to P_i^o, so the firm with higher price is more likely to satisfy the condition for running informative advertising. $\square$

Proposition 4 implies that the consumers' purchase decisions affect each firm's decision of whether to engage in informative advertising. Furthermore, in general the market competition based on product differentiation plays an essential role in making advertising strategies.

In addition, our analysis for costly informative advertising is a complement of the basic quality reporting analysis, and they have consistently demonstrated the importance of product differentiation, which is also a new finding that contributes to the literature of advertising.

2.7 Concluding Remarks

Our analysis of firms' advertising strategies involving third-party product reviews contributes to the literature of advertising and information disclosure in the following three ways.

First, we characterize conditions related to product attribute advantage under which a firm advertises the result of a third-party product review. One general result of our model – one that is consistent with incomplete disclosure in the literature – is that a firm may not inform consumers about a higher attribute quality if that would reduce product differentiation.

We distinguish scenarios based on whether or not only one firm in duopolistic market has an absolute advantage in each of the two product-attributes. If one firm has absolute advantages in each product attribute, then the superior firm advertises any attribute whose actual quality reported in a third-party review is higher than consumer prior expectation, whereas the other inferior firm advertises a higher quality only for the attribute with comparative advantage. Moreover, we find that the inferior firm exits the market following a third-party product review if it has a substantial comparative disadvantage in both product attributes.

Second, we find that a firm's persuasive advertising about the importance of a product attribute can help the firm in two ways: (1) it is profitable to run even if the third-party review reveals a lower actual quality than the consumer prior, and (2) it protects the inferior firm by making it more likely that the firm will remain in the market. In reality, we observe many practices of persuasive advertising about product attributes. For example in the automobile market, manufacturers attempt to convince consumers the importance of particular attributes.

Third, a firm's decision whether to engage in a costly informative advertising about the results of a third-party product review depends not only on consumer purchase decisions, but also on the advertising decisions of its competitors. This dependence is possibly extreme in the sense that the firm might decide to free ride off a competitor's advertising.

To the best of our knowledge, no empirical studies have estimated firms' advertising choice based on third-party reviews. Emphasizing one of our results, and that in the quality disclosure literature, about a firm's possible unwillingness to disclose that its product if higher quality that consumers expect raises the question about whether producers of low-quality products promote this low quality. Empirical analyses addressing these issues could provide evidence of our theoretical results.

Figures

Figure 2.2: Outcomes for Numerical Example 1: $\tilde{\Delta}_1 \in [0, 1]$, $\tilde{\Delta}_2 = -1/2$

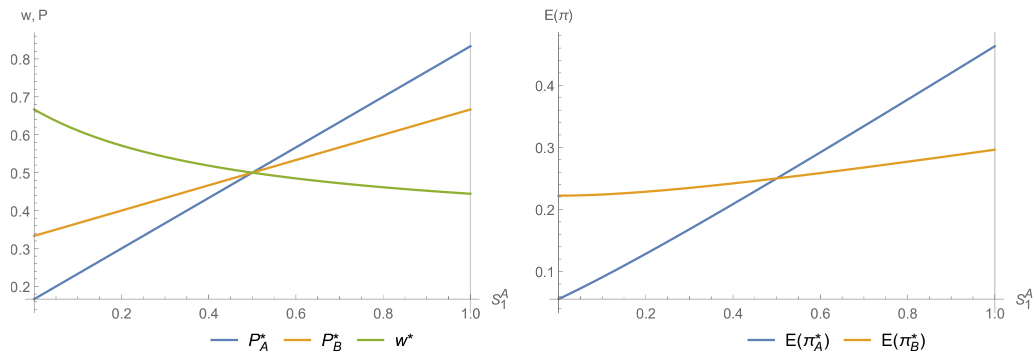

Figure 2.3: Outcomes for Numerical Example 2: $\tilde{\Delta}_1 \in [1/2, 1]$, $\tilde{\Delta}_2 = 1/4$

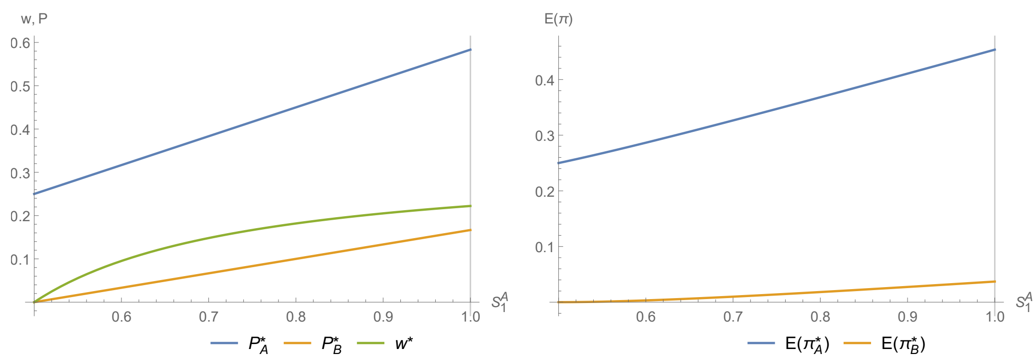

Figure 2.4: Outcomes for Numerical Example 3: $\tilde{\Delta}_1 = 1/2$, $\tilde{\Delta}_2 \in [-3/4, 1/4]$

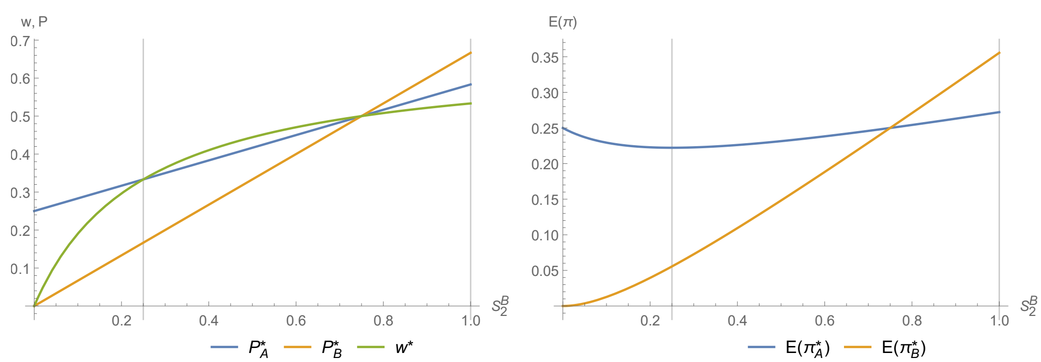

Figure 2.5: Equilibrium of Persuasive Advertising for Example 1

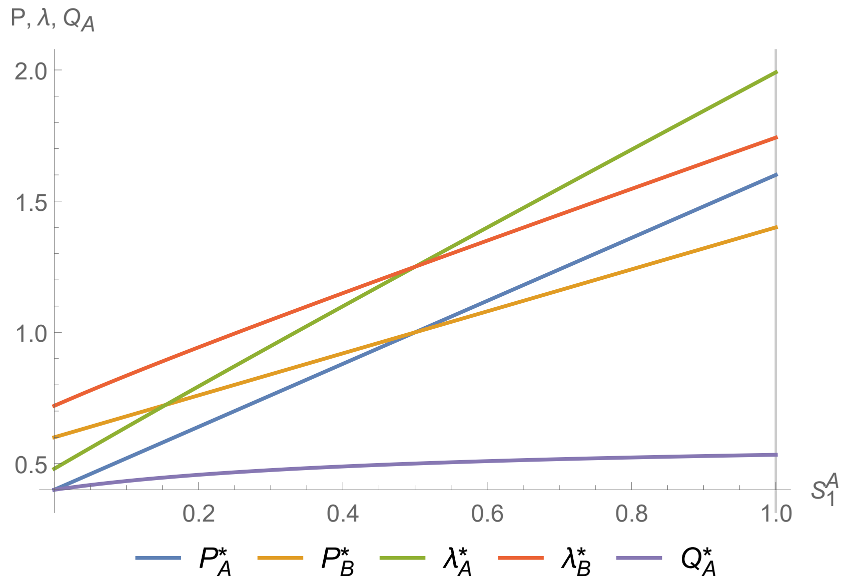

Figure 2.6: Equilibrium of Persuasive Advertising for Example 2

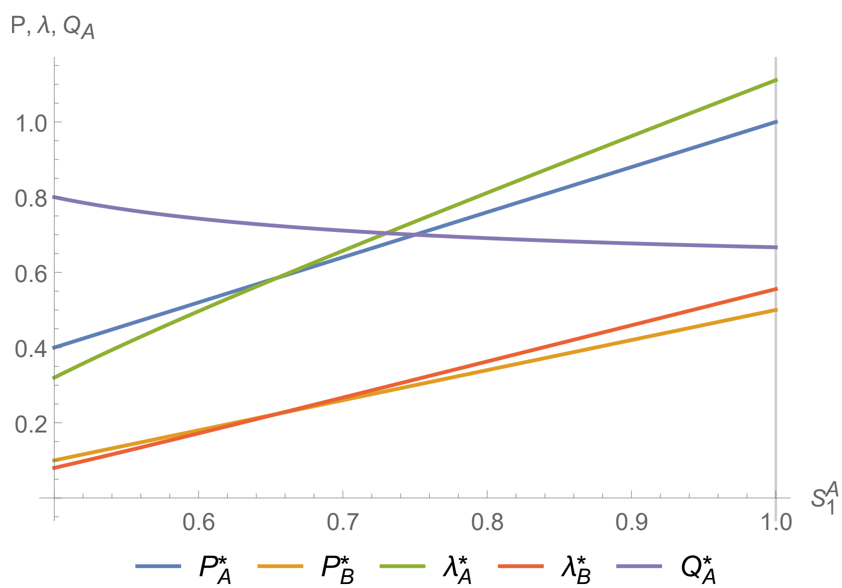

Figure 2.7: Equilibrium of Persuasive Advertising for Example 3

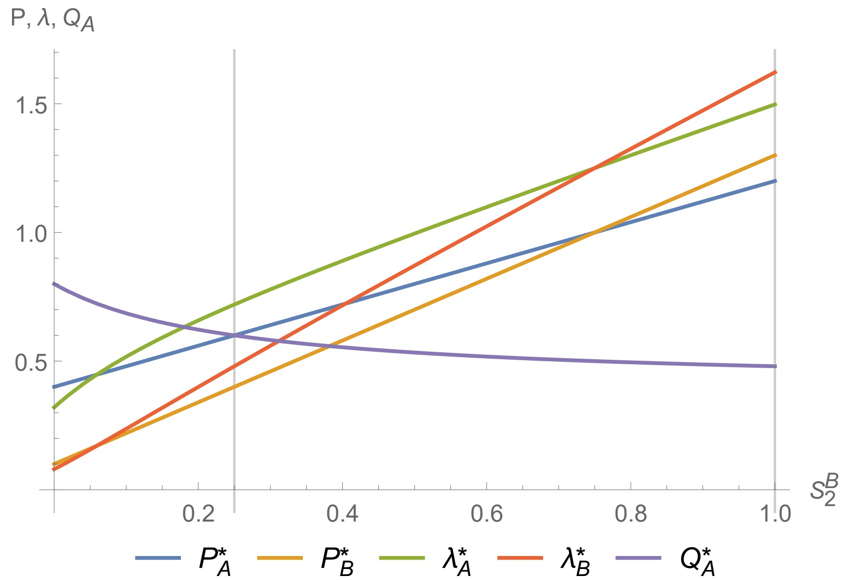

Chapter 3: Product Return Policies

Abstract

This paper examines a retailer's optimal product return policy within an environment in which consumers differ in their product return propensities. A liberal return policy is good for consumers because it increases consumer expected utility by lowering the risk involved in purchases that turn out to be of low value. However, it is bad for retailers because it leads to frequent returns. In addition to characterizing retailer-optimal product return policies, our primary goal is to identify the scenarios in which retailers: (i) offer free returns, no returns, or policies in between; (ii) ban some consumers from returning products, or more generally tailor return policies to consumer attributes; and (iii) effectively fire some consumers from purchasing products at all. To characterize each of these three scenarios, we include in our model cases in which consumers vary by their distributions over possible product values (and a retailer may or may not be able to identify consumers by these product return preferences), consumers differ by their hassle costs of returning products, and products returned by consumers involve an additional cost to retailers either in the form of restocking fees or losses in product value.

3.1 Introduction

A generous return policy is always welcome by consumers because after detailed product evaluation, following purchase, it allows them to return unwanted merchandise and keep the ones that they like. When making a purchase decision, particularly during online shopping in which consumers cannot physically evaluate product characteristics such as physical quality, texture, and fit, then a liberal return policy could increase their expected utility by reducing the risk involved in keeping low-value products. However, such a policy is costly for sellers because of product depreciation and management of returns (Anderson, Hansen and Simester, 2009). The costs created by returns have turned out to be so high that sellers have re-evaluated this policy. For example, two recent popular press articles, Safdar (2018) and Safdar and Stevens (2018), report that Best Buy and Amazon have been tracking customer shopping and product return behavior and have placed restrictions on consumers who have returned products at high rates and have even engaged in fraudulent returns.

Prior marketing research has investigated this trade-off involved in product return policies between the need to increase consumer expected utility of purchases and the cost of product returns. Padmanabhan and Png (1997) point out that a generous return policy could make a retailer more competitive; Chu, Gerstner and Hess (1998) and Bechwati and Siegal (2005) find it could increase purchases; and Petersen and Kumar (2009) find evidence that current customer product return behavior affects future buying decisions, so that the retailer could benefit in the long run from a generous return policy. On the contrary, opponents believe that the cost of a generous policy offsets its benefits. It might only increase the amount of orders but not profits (Anderson, Hansen and Simester, 2009), can lead to more returns (Davis, Hagerty and Gerstner, 1998), and can be abused (Rust, Zahorik and Keiningham, 1996). Additionally, Wood (2001) find that a retailer's return policy may affect consumers who live in geographically remote regions differently from those who live close to retailers in densely populated areas. In turn, a retailer's location may affect its return policy.

Apart from the influence of return policy on consumer behavior or retailer profit, this paper aims to explore return policy in a market environment where consumers are heterogeneous in return behavior. We ask three questions. (i) What are the conditions for the optimality of different return policies, namely free returns, no returns, or partial returns? (ii) What is the optimal behavioral return policies when the retailer is able to identify a consumer's purchase behavior and then tailor return policies to consumer attributes? and (iii) Does a retailer have incentive to fire consumers who are more likely to return products?

We examine return policies in a model in which consumers at the time of purchase are uncertain about product values. Consumers also bear costs of returning products, which we term as "hassle costs." Also in our model, products returned by consumers involve an additional cost to retailers either in the form of restocking fees or losses in product value.

The first finding in our model confirms that the optimal return policy is related to a retailer's cost of dealing with returns. In an optimal policy, the seller accounts for this cost and may even generate profit from returns. We also demonstrate that a retailer is made better off by detecting consumer return heterogeneity and treating consumers with different return propensities differently. After learning consumer types, a retailer might offer a better return policy, either a higher refund or a lower price, for consumers who are less likely to return products. Additionally, when it is too costly to deal with returns, a retailer may effectively restrict the purchases from consumers who are more likely to return products. In this situation, a "no return" policy would not be helpful, instead the retailer optimally chooses to "fire" consumers who have a high probability of returning products.

3.2 A Roadmap of the Models

In this market, a retailer attempts to sell one product to consumers, and the retailer establishes a price for his product and a return policy for buyers. To set the optimal price and return policy, the retailer needs to consider some market factors which include the cost of the product, the potential cost of managing returned items, the market competition, consumer heterogeneity, and information asymmetry. We investigate the role of those factors through a series of models following the same basic setting. This is the first study in this area to model return policies with all those factors. The models we have explored are listed in Table 3.1.

Table 3.1: Description of Cases

Models	Section	Market competition	Consumer heterogeneity	Identification
Baseline	3	monopoly	*ex ante* identical	No
Bertrand	4	price competition	*ex ante* identical	No
Return propensities	5	monopoly	by hassle costs	Yes
Return costs	7.1	monopoly	by restocking costs	Yes
Additional information	7.2	monopoly	by product values	No
Online vs Offline	E	Hotelling competition	*ex ante* identical	No
Return preference	F	monopoly	by preference for "free return"	No

We begin with a baseline model in which a monopoly retailer sells to consumers who are *ex ante* identical but with heterogeneous *ex post* product valuations. We find that the equilibrium return policy is determined by the cost to the retailer for managing returned products. The retailer would like to offer a generous return policy when its cost for dealing with returns is small, while offer restrictive return policy when that cost is big. Furthermore, "free return" policy could be the equilibrium strategy only when it is costless for the retailer to resell returned products and the product cost is big enough, and the retailer cannot earn positive profit under this circumstance. As this condition is too restrictive, we argue that a monopoly retailer does not have incentive to offer "free return" when they need to pay some cost to manage returned products. This outcome is consistent with existing studies built on similar model settings.

However, we do observe "free return" in reality, so it is possible that other factors motivate it. We develop three models with the concern of competition or consumers' preference for return policies. To examine the role of competition in determining optimal product return policies, we examine two cases: Bertrand competition and online-offline competition. In Bertrand competition, the retailers set the same return policy as does a monopolist, but with a lower equilibrium product price. In the online-offline competition model, two retailers offer different return policies, since consumers could physically examine the product before purchase in offline stores. However, in either case of competition, "free return" is still not an equilibrium policy when restocking cost exists, which implies that it is not competition that motivates a retailer to offer generous return policy. Then as a special case, we develop one model with the concern of consumers' preference for return policies and find that "free return" could be an optimal strategy in this circumstance.

To address the next question on the optimal return policy when consumers are heterogeneous in return behaviors, we model this using two factors: consumers' hassle cost of returning products and the retailer's cost of managing returns caused by consumers' return behaviors. Both models are developed based on the basic model with a monopoly retailer. In the model with two levels of hassle costs, high and low, we consider four cases classified by whether the retailer can identify consumer types and set discriminatory return policies.

In the first two of these four cases, the retailer cannot distinguish consumers by hassle costs. In one case, the retailer offers one price and product return policy. In the other, the retailer screens consumers by offering two policies, one intended for consumer types with low hassle costs and other for consumers with high hassle costs. The retailer can increase its profit by using two policies to screen consumers by using second-degree price-return policy discrimination. In cases three and four, the retailer can distinguish consumer types. In both cases the retailer practices third-degree (i.e., multi-market) price return-policy discrimination by offering two prices or product return policies, with each consumer type being offered only one policy. We find that the retailer can benefit from this price discrimination and offer a better return policy to consumers with higher hassle cost. In these cases, the retailer may offer free returns to one type and even no returns to the other, and even ban the purchase from consumers with low hassle cost.

Next, we add another dimension of consumer heterogeneity: return behavior that causes different return costs for the retailer. Finally, we model whether a retailer provides additional product information prior to purchase to reduce consumers' uncertainty of product value. In doing so, the retailer might reduce consumer *ex ante* about their values of the retailer's product, thereby reducing return propensities. However, the trade-off in this policy is that the retailer reduces initial product purchases.

3.3 The Basic Model

Following Che (1996), we develop our model with a monopoly retailer who faces a unit mass of consumers, each of whom desires at most one unit of a good. The retailer pays a product cost c_0 per unit of the good sold, which includes not only the payment to a supplier or the manufacturing cost, but also other transaction costs involved in selling it. Moreover, when permitting customers to return products, the retailer pays an additional cost c_r per unit of returned item, which includes depreciation, shipping, restocking, or rearrangement. We assume it measures not only the cost to the retailer for returned products, but also the loss to the product value caused by buyers. It is plausible to assume that $c_0 \geqslant 0$, $c_r \geqslant 0$. And in general, the retailer would accept returns only when $c_r < c_0$, because otherwise the retailer would prefer to acquire or manufacture new product to process a return.[1]

The retailer charges a price, p, for per unit of the good sold, and sets a return policy for buyers. We characterize the retailer's return policy by the amount of refund, r, that the retailer pays consumers for each returned item. Effectively, the return charge to a consumer is $p - r$. Two common return policies are noteworthy: a "free return" policy, $r = p$; and a "no return" policy, $r = 0$. It is reasonable to assume that the refund is non-negative and no greater than the purchase price, that is, $r \in [0, p]$.[2]

We model the transaction between the retailer and the consumers as a three-stage process. In stage 0, each consumer is uncertain about her value $v \in [0, V] \subset \mathbb{R}$ of the good,

[1] Amazon did this for some items like cheap books. It offers a full cash refund but does not ask customers to ship the returned books back.

[2] A negative refund means buyers need to pay extra money to return an item, then no one would like to return and the outcome is the same as "no return". If refund is greater than the purchase price, buyers can make some money from a "buy-and-return" deal, which is not a typical transaction.

and she and the retailer believe that it is distributed according to $G(v)$. The distribution function $G(v)$ is also the distribution of values for the unit mass of consumers, which the retailer learns. In stage 1, the retailer sets price p and return policy r. After learning (p, r), the consumers decide whether to purchase the product. In stage 2, after receiving and evaluating the product, each consumer learns her value v and then decides whether to keep or return the product. Note that consumers are *ex ante* identical, and that their *ex post* realized valuations are idiosyncratic. This three-stage process can be commonly applied to many online shopping in which consumers cannot physically evaluate product characteristics such as physical quality, texture, and fit, so that they cannot form an accurate willingness to pay for the listed product. Because of this, when making the purchase decision, each consumer is uncertain about her value of the product.

To return an item, a consumer incurs a "hassle cost," h, which includes the travel cost to return a product and the delayed purchase or the search cost for a replacement. Note that this hassle cost is beyond the monetary return fee charged by the retailer (which is represented by $p - r$). Even if the retailer sets "free return" policy, it is not costless for buyers to return purchased items, especially for those consumers who really need the product at the moment. To some extent, it represents consumer heterogeneity other than product value. Because of this, the retailer also encounters an uncertainty of the probability that one buyer would return the product in the end.

We analyze the model beginning with the consumers' third-stage decisions whether or not to return the purchased products. Given price p and realized value v, a consumer enjoys utility $v - p$ if she keeps the item, and $r - h - p$ if she returns it. Thus she keeps the purchased product if and only if $v - p \geqslant r - h - p$, that is, $v \geqslant r - h$. Therefore, one consumer would keep the product if its perceived value is above the threshold which is determined by refund and hassle cost. And consumers are more likely to keep the product when refund is lower or hassle cost is higher.

Based on consumers' return decisions at the third stage, the measure of the final product demand is $1 - G(r - h)$ and the measure of returns is $G(r - h)$. Note that the final consumer demand is determined by the retailer's return policy. Therefore, we characterize the retailer's reasonable return policy in the following lemma.

Lemma 6. *If r is an equilibrium refund, then $r \in [h, V + h]$. Furthermore, if $r = h$ is an optimal refund, then any $r \in [0, h]$, including the no-return policy $r = 0$ is equally optimal.*

Proof. We can derive it from consumer demand $1-G(r-h)$. If $r-h > V$, then $G(r-h) = 1$, in which case 100% of the consumers return the product. If so, then a consumer purchases only when more than full refund is offered, $r > p$, which is unprofitable for the retailer. So the retailer would consider $r \leqslant V + h$. Moreover, if $r - h < 0$, then $G(r - h) = 0$, in which case 100% of the consumers keep the product. If so, the exact value of r does not affect either consumers' purchase decision or the retailer's profit. In other words, both consumers and the retailer are indifferent with any $r \in [0, h]$. Therefore, a "no return" policy can be offered without any change when the optimal refund is $r = h$. □

Lemma 6 suggests that to obtain a positive consumer demand, the retailer cannot set a too high refund. In particular, the refund should be lower than $V + h$. We will apply this restriction to our analysis.

Next we consider the stage-1 decision in which the retailer sets a price and return policy that attract consumers to purchase. Since each consumer is uncertain about her value at this stage, she makes a purchase decision based on her expected utility of purchase, which is denoted as $E[U(p,r;v)]$. We assume a consumer's utility is zero if she does not purchase. Therefore, she purchases if and only if her expected utility is non-negative:

$$E[U(p,r;v)] \geqslant 0. \tag{3.1}$$

Recall that in the end a consumer enjoys utility $v-p$ if she keeps the item, and $r-h-p$ if she returns it. Therefore, a consumer's expected utility of purchase can be calculated as the following:

$$\begin{aligned} E[U(p,r;v)] &= E_v[\max\{v, r-h\}] - p \\ &= (r-h)\cdot G(r-h) + \int_{r-h}^{V} v\cdot g(v)dv - p \\ &= V - \int_{r-h}^{V} G(v)dv - p. \end{aligned} \tag{3.2}$$

The retailer's stage-1 goal is to maximize its profit. We assume that if consumers do not purchase, then the retailer's profit is zero. When some consumers choose to purchase in stage 1, the retailer's profit is determined by how many those buyers would keep the product. Therefore, the retailer's expected profit contains two parts: (1) the profit from buyers who keep the product, $p-c_0$, and (2) the profit or loss from dealing with returns, $p-r-c_r$. Note that we assume the retailer could resell returned goods after some simple treatments like repackaging, with restocking cost c_r for each returned item. Thus he only needs to maintain an inventory of $q = 1-G(r-h)$, which determines the total product cost.[3] As a result, the seller's expected profit function is:

$$\begin{aligned} \pi(p,r) &= (p - c_0)[1 - G(r-h)] + (p - r - c_r)G(r-h) \\ &= p - c_0 + (c_0 - c_r - r)G(r-h) \end{aligned} \tag{3.3}$$

Since Lemma 6 implies that we only need to focus on refund $r \in [h, V + h]$, both the expected consumer utility in equation (3.2) and the expected profit in equation (3.3) are continuous and differentiable with respect to p and r in this region.

3.3.1 The Equilibrium Return Policy

In this section, we characterize the equilibrium by examining the retailer's potential optimal price and return policy, which are obtained by maximizing his expected profit in equation (3.3), subject to the representative consumer's participation constraint, as expressed in equation (3.1) and (3.2). The potential optimal price and return policy $(\tilde{p}, \tilde{r})$ are presented in the following proposition, with proof in the Appendix.

Proposition 5. *In a market where a monopoly retailer sells a product to a unit mass of consumers, the retailer's potential optimal price and return policies are:*

$$(\tilde{p}, \tilde{r}) = \begin{cases} (V - \int_0^V G(v)dv,\ 0) & \text{if } c_r \geqslant c_0 - h, \\ (V - \int_{c_0 - c_r - h}^V G(v)dv,\ c_0 - c_r) & \text{if } c_r \in [c_0 - V - h,\ c_0 - h), \\ (V,\ V + h) & \text{if } c_r < c_0 - V - h. \end{cases} \tag{3.4}$$

[3]This assumption is consistent with that used by Che (1996), and we think it is plausible to assume that one retailer could sell out all the inventory in the end. Although studies like Conlon and Mortimer (2013) examine the impact of product availability on demand, our model does not have this concern.

Given $(\tilde{p}, \tilde{r})$, *the retailer's profit is:*

$$\tilde{\pi} = \begin{cases} V - c_0 - \int_0^V G(v)dv & \text{if } c_r \geqslant c_0 - h, \\ V - c_0 - \int_{c_0 - c_r - h}^V G(v)dv & \text{if } c_r \in [c_0 - V - h,\ c_0 - h), \\ V - c_0 & \text{if } c_r < c_0 - V - h. \end{cases} \tag{3.5}$$

Proposition 5 suggests that the retailer's optimal strategies depend on the restocking cost c_r, and he may offer no refund if the restocking cost is high while more generous refund when it is small. This result reveals the potential impact of consumer decisions on the retailer's choice of return policies. We can observe the evidence from the real world, such as big online retailers like Amazon and Best Buy are considering to modify their free return policy (which is reported in (Safdar, 2018; Safdar and Stevens, 2018). In addition, the retailer's optimal profit is decreasing in the restocking cost. This is consistent with the arguments that many retailers have attempted to reduce restocking costs (Drake, 2014).

However, some of the potential optimal strategies $(\tilde{p}, \tilde{r})$ expressed in Equation (3.4) will not be chosen by the retailer in the equilibrium, because his profit is actually negative. Taking into account the retailer's decision to opt out of the market, the equilibrium is in the following proposition.

Proposition 6. *If* $c_0 \geqslant h$ *and* $V - c_0 - \int_{c_0 - h}^V G(v)dv \geqslant 0$, *then there exits a unique equilibrium in which the retailer sells to all consumers, but only those with perceived product value* v *greater than* $r^* - h$ *keep it. The equilibrium price and refund are*

$$(p^*, r^*) = \begin{cases} (V - \int_0^V G(v)dv,\ 0) & \text{if } c_r \geqslant c_0 - h, \\ (V - \int_{c_0 - c_r - h}^V G(v)dv,\ c_0 - c_r) & \text{if } 0 \leqslant c_r < c_0 - h. \end{cases} \tag{3.6}$$

The retailer's equilibrium profit is:

$$\pi^* = \begin{cases} V - c_0 - \int_0^V G(v)dv & \text{if } c_r \geqslant c_0 - h, \\ V - c_0 - \int_{c_0 - c_r - h}^V G(v)dv & \text{if } 0 \leqslant c_r < c_0 - h. \end{cases} \tag{3.7}$$

Proof. Recall that all parameters in our model, c_0, c_r, V and h, are non-negative. We first prove by contradiction that the third line $(\tilde{p}, \tilde{r}) = (V, V+h)$ in Equation (3.4) is not an equilibrium. If the seller offers $(p, r) = (V, V+h)$, we have $c_0 - V - h > 0$ to guarantee non-negative values for c_r, that is $V - c_0 < -h \leqslant 0$. The optimal profit then is $V - c_0 < 0$. Therefore, the retailer is better off not selling the product. For the first two lines of Equation (3.4), we need $c_0 - h \geqslant 0$ so that a non-negative region for possible values of c_r exists. If so, then these two cases exist. Because $\tilde{\pi}$ is decreasing in c_r, if $\tilde{\pi}|_{c_r=0} = V - c_0 - \int_{c_0-h}^{V} G(v)dv \geqslant 0$, then the retailer's profit is positive for these two cases.

Considering consumers, they purchase because $E[U(p^*, r^*; v)] = 0$, and they keep the product if and only if $v \geqslant r^* - h$. □

The equilibrium has the following implications. First, the retailer's restocking cost plays an essential role in determining return policy. If restocking cost approaches or is larger than the product cost, then the retailer does not accept returns. When the retailer allows returns, the amount of refund depends on two costs: c_0 and c_r. Since c_0 is either, or equivalent to, the product's manufacturing cost, while c_r is the loss or depreciation of this value, in some studies the difference between them is called salvage value (e.g. Su, 2009). A higher salvage value means less damage to returns. We find that the optimal refund offered by the retailer equals the returned product's salvage value.

Second, the retailer's restocking cost c_r contributes to determining his equilibrium profit. Since the retailer reduces equilibrium refund r^* when restocking cost goes up, he lowers product price to maintain purchases. As a result, although fewer consumers return the product, the retailer's profit declines. That is, when restocking cost increases, the retailer cannot maintain high profit by lowering price. To raise profit, the retailer needs to reduce restocking cost.

Third, the consumer propensity to return purchased products, which is shown by the distribution function $G(v)$ in our model, also contributes to determining the retailer's equilibrium profit. If $G(v)$ shifts from $G_1(v)$ to $G_2(v)$, where $G_1(v) < G_2(v)$ for each $v \in (0, V)$, so that buyers are more likely to return their purchased products, then the retailer would respond by reducing p^*, but not changing r^*. Although it is less costly for consumers to return an item, fewer consumers keep the product so that the retailer's profit goes down. Therefore, a retailer would be better off from a population that have high values for the product so that they are more likely to keep it.

Fourth, when the retailer offers refund for returns, he sets it such that $\pi^* = p^* - c_0 = p^* - r^* - c_r$. Because $p^* - c_0$ represents the profit from each unit sold, and $p^* - r^* - c_r$ from each returned item, the retailer earns the same profit, whether or not the consumer returns the product. This to some extent could explain why so many retailers would like to allow returns: it will not just encourage consumers to purchase but also provides another (identical) source of profits.

Fifth, the cost that a consumer needs to pay to return an item is $p^* - r^* = V - (c_0 - c_r) - \int_{c_0 - c_r - h}^{V} G(v)dv$, which is decreasing in c_0 and increasing in c_r. Interestingly, the return policy is more generous for products that are either expensive or easy to resell, while be restrictive for products that are either cheap or easy to be damaged or depreciated. As an extreme, if the restocking cost is so high that $c_r \geqslant c_0 - h$, the retailer optimally constructs a "no return" policy.

Finally, we examine the condition for "free turn" policy, $r = p$, which is presented in the following corollary.

Corollary 3. *In equilibrium, if the retailer offers a "free return" policy,* $r^* = p^*$*, then* $c_r = 0$ *and* $\int_{c_0-h}^{V} G(v)dv + c_0 = V$*, and* $\pi^* = 0$.

Proof. Suppose "free turn" is the equilibrium, $p^* = r^*$, then the equilibrium profit is $\pi^* = p^* - r^* - c_r = -c_r$. In the equilibrium, we should have $\pi \geqslant 0$. Therefore, we must have $c_r = 0$. Then we apply the expressions of p^* and r^* to replace $p^* = r^*$ to derive $V - \int_{c_0-c_r-h}^{V} G(v)dv = c_0 - c_r$, which is $\int_{c_0-h}^{V} G(v)dv + c_0 = V$ when $c_r = 0$. □

The interpretation is that the retailer has incentive to offer "free return" when restocking cost is zero and product cost is sufficiently large.[4] However, it is common in reality that the retailer needs to spend some cost to deal with returned products, and sometimes this cost is substantial. Therefore, in the context of the model in this section, a "free return" policy is not an equilibrium.

3.4 Competition

In this section, we add market competition to the basic model so as to investigate the effect of competition on optimal pricing and return policies. Although previous studies such as Che (1996) have discussed competition as an potential explanation for return policies, no study has demonstrated the conditions under which competition leads to free return policies. We contributes to this investigation by extending our basic model to a market with J retailers that compete by selling non-differentiated products. To isolate the effect of competition on optimal pricing and return policies, our model of the consumers is identical to that of the previous section. Recall that the distribution of the product values for the

[4] With a uniform distribution on interval $[0, 1]$, that is, $G(v) = v$, the condition for a "free return" policy is $c_r = 0$ and $c_0 = h - \sqrt{2h} + 1$. Since $h \geqslant 0$, so $c_0 \geqslant 1/2$ when "free return" is optimal to provide.

unit mass of consumers is $G(v)$ on $[0,\ V]$. And a retailer j sets his product price p_j and refund r_j to maximize his expected profit.

The effect of competition on optimal pricing and return policies is illustrated in the following lemma.

Lemma 7. *Shifting from monopoly to competition, a retailer lowers the equilibrium product price p^* but does not change the equilibrium refund r^*. In the competitive equilibrium, if the retailers select a "free return" policy, then his profit is zero.*

Proof. Suppose there are J retailers and the maximum expected utility that a representative consumer could obtain from retailers other than j is $\bar{u}$. Retailer j needs to provide at least $\bar{u}$ in order to sell its product. That is, $E[U_j(p_j, r_j; v)] \geqslant \bar{u}$. Thus the retailer's optimization program is

$$\max_{p_j, r_j} \pi_j = (p_j - c_0)[1 - G(r_j - h)] + (p_j - r_j - c_r)G(r_j - h),$$

$$\text{subject to } E_v[\max\{v, r_j - h\}] - p_j \geqslant \bar{u}.$$

Following the same process as we derive Proposition 9, we obtain the potential equilibrium (p_j^*, r_j^*) as:

$$(p_j^*, r_j^*) = \begin{cases} (V - \bar{u} - \int_0^V G(v)dv, 0) & \text{if } c_r \geqslant c_0 - h, \\ (V - \bar{u} - \int_{c_0 - c_r - h}^V G(v)dv, c_0 - c_r) & \text{if } 0 \leqslant c_r < c_0 - h. \end{cases}$$

Respectively, the retailer's equilibrium profit is:

$$\pi^* = \begin{cases} V - \bar{u} - c_0 - \int_0^V G(v)dv & \text{if } c_r \geqslant c_0 - h, \\ V - \bar{u} - c_0 - \int_{c_0 - c_r - h}^V G(v)dv & \text{if } 0 \leqslant c_r < c_0 - h. \end{cases}$$

The conditions for those equilibrium to exist are $c_0 \geqslant h$ and $V - \bar{u} - c_0 - \int_{c_0 - h}^V G(v)dv \geqslant 0$.

Comparing this result with the equilibrium in Proposition 9, the equilibrium refunds are identical but the price decreases by $\bar{u}$.

Again, we have $\pi^* = p^* - r^* - c_r$, so that "free return" in the equilibrium cannot bring positive profit for the retailer. The conditions for "free return" as the equilibrium is $c_r = 0$ and $c_0 + \int_{c_0-h}^{V} G(v)dv = V - \bar{u}$.[5] □

In the context of our model, competition has not effect on a retailer's optimal product return policy. Instead competition causes the retailers to lower their prices. From consumer side, although the refund stays the same, a lower equilibrium price increases consumers' welfare.

3.4.1 Bertrand Competition

As a special case, we investigate Bertrand competition to further demonstrate the equilibrium with competition. Suppose there are $J \geqslant 2$ retailers who sell identical products. Each firm j sets price and refund (p_j, r_j) to maximize his expected profit. A consumer's expected utility from selecting firm j is $E[U_j(p_j, r_j; v)]$, which is calculated based on equation (3.2). Unlike traditional Bertrand competition in which consumers purchase from the seller with lower price, in our model a consumer purchases from the retailer that offers her the greatest expected utility. As a result, the market demand for retailer j is:

$$D_j = \begin{cases} 1, & \text{if } EU_j > \max[EU_{-j}] \\ \alpha_m, & \text{if } EU_j = \max[EU_{-j}] \\ 0, & \text{if } EU_j < \max[EU_{-j}] \end{cases}$$

[5]Under a uniform distribution on interval $[0, 1]$, that is, $G(v) = v$, the condition for "free return" is $c_r = 0$ and $c_0 = h - \sqrt{2(h + \bar{u})} + 1$. Since $h - \sqrt{2(h + \bar{u})} + 1 \leqslant h - \sqrt{2h} + 1$, the requirement for c_0 is relaxed compared to that in the basic model.

where EU_{-j} is a consumer's expected utility from selecting a firm other than firm j, and $\alpha_m \in (0,1)$ denotes the fraction of consumers that select firm j when m firms offer the highest expected utility to each consumer.

If retailer j sells one unit of his product, his expected profit π_j is calculated by equation (3.3). Considering the market demand, the retailer's expected total profit is:

$$\pi_j = \begin{cases} p_j - c_0 + (c_0 - c_r - r_j)G(r_j - h), & \text{if } EU_j > \max[EU_{-j}] \\ \alpha_m[p_j - c_0 + (c_0 - c_r - r_j)G(r_j - h)], & \text{if } EU_j = \max[EU_{-j}] \\ 0, & \text{if } EU_j < \max[EU_{-j}] \end{cases}$$

3.4.2 The Bertrand Equilibrium

In our model of Bertrand competition, only retailers who can provide highest expected utility to consumers sell their products. Therefore, instead of maximizing profit subject to consumer participant constraint, a retailer's objective is to maximize consumers' expected utility subject to a non-negative profit constraint. Specifically, the retailer's optimization program is:

$$\max_{p_j, r_j} E[U(p_j, r_j; v)] = V - \int_{r_j - h}^{V} G(v)dv - p_j,$$

$$\text{s.t. } \pi_j \geqslant 0.$$

Graphically, we can illustrate the retailer's motivation in the Bertrand competition by Figure 3.1 which depicts consumers' indifferent curves with red lines and the retailer's isoprofit curves with blue lines. A retailer's objective is to select (p_j, r_j) to maximize profit subject to the representative consumer's utility constraint.

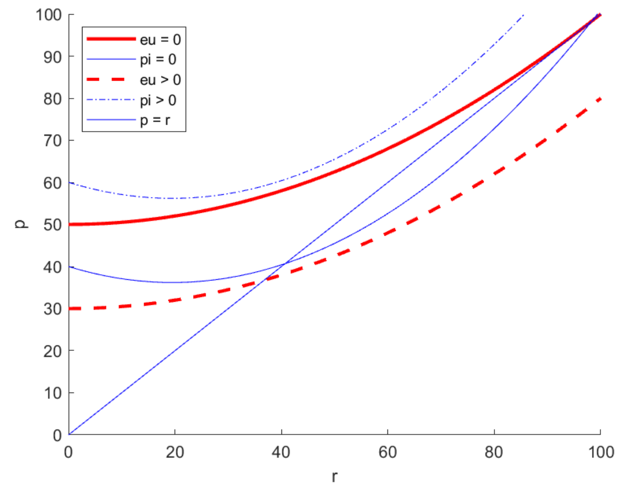

Figure 3.1: Expected Utility and Profit under Bertrand Competition

We characterize the market equilibrium of the Bertrand competition in the following proposition.

Proposition 7. *If $c_0 \geqslant h$ and $V - c_0 - \int_{c_0-h}^{V} G(v)dv \geqslant 0$, then there exits a unique equilibrium in which retailer j sets*

$$(p_j^*, r_j^*) = \begin{cases} (c_0,\ 0) & \text{if } c_r \geqslant c_0 - h, \\ (c_0,\ c_0 - c_r) & \text{if } 0 \leqslant c_r < c_0 - h. \end{cases} \tag{3.8}$$

In each case retailer j's equilibrium profit is $\pi_j^* = 0$ and each consumer purchases one product but only those with perceived product value above $r^* - h$ keep it.

Proof. The retailer's optimization program in selecting r_j is equivalent to

$$\max_{r_j} EU_j = V - \int_{r_j-h}^{V} G(v)dv - c_0 + (c_0 - c_r - r_j)G(r_j - h),$$

$$\text{with } p_j = c_0 - (c_0 - c_r - r_j)G(r_j - h).$$

The derivative of expected utility with respect to r_j is

$$\frac{\partial EU_j}{\partial r_j} = (c_0 - c_r - r)g(r - h).$$

Since $g(r-h) > 0$ for any $r \in [h, V+h]$, then if $h \leqslant c_0 - c_r \leqslant h + V$, we can obtain $r^* = c_0 - c_r$; if $c_0 - c_r \leqslant h$, we have $\frac{\partial EU_j}{\partial r_j} < 0$, which gives us a corner solution $r^* = h$; and if $c_0 - c_r \geqslant V + h$, we have $\frac{\partial EU_j}{\partial r_j} > 0$ which gives us another corner solution $r^* = V + h$. However, $EU_j < 0$ when $r^* = V + h$, so it is not an equilibrium. And the first two solutions could be equilibrium if $\max[EU_j] \geqslant 0$, that is, $V - \int_{c_0-h}^{V} G(v)dv - c_0 \geqslant 0$. In addition, Lemma 6 shows that $r^* = h$ is equivalent to $r^* = 0$. Thus, we have proved all the cases. Since the retailer's expected profit is zero in all cases, he does not have incentive to stop the business. Therefore, all cases could be equilibrium outcomes.

For consumers, as their expected utility $E[U(p^*, r^*; v)] > 0$, they would purchase, and keep it if $v - p^* \geqslant r^* - h - p^*$, that is $v \geqslant r^* - h$. Other buyers with lower perceived value return the product. □

Comparing this equilibrium with that of the basic model in Proposition 9, we find the only variation is equilibrium price and profit. The market competition causes retailers to reduce product price and earn zero profit, but it does not influence either the condition for the equilibrium or refund. Again, the retailer would not offer any refund if the restocking cost is sufficiently high while offer more generous refund for lower restocking costs.

We examine the condition for "free turn" to be the equilibrium, which is presented in the following corollary.

Corollary 4. *In the market with Bertrand competition, if there is no restocking cost,*

$c_r = 0$, the retailer chooses a "free return" policy.

Proof. Based on Proposition 7, if $c_r = 0$, we always have $p^* = r^*$, which is "free return". □

Although the condition for a "free return" policy is relaxed compared to the case of monopoly seller, it is still determined by the restocking cost, and in reality the retailer usually needs to pay some cost to manage returned items.

In sum, Bertrand competition results in zero profit, and each retailer offers the same refund as in the monopoly market, but with lower price. A "free return" policy would be offered only if the restocking cost is zero. In Appendix E, we examine the competition between online and offline retailers, and do not observe "free return" as the equilibrium for the online retailer.

As one of our primary interests in this paper, we still want to ask what could make a "free return" policy optimal as well as profitable with the concern of restocking cost. To address this question, in F we explore another model with consumer preference for "free return" policy, which provides an answer to the question.

3.5 Heterogeneous Return Propensities

In the basic model, although we assume consumers may have different *ex post* product values, they are *ex ante* identical with the same probability of returning a product. Now we relax this assumption with the concern that consumers may differ in return propensities even before their purchase decisions. For example, some consumers have a high opportunity cost of time or delay purchase, in which case they may be unlikely to return products due to the high cost of hassle involved. Other consumers, who have a low opportunity cost of time or delay purchase, may be likely to return products because there are unfazed by the hassle of doing so. We model consumers' return propensities by introducing heterogeneous hassle costs of returning products.[6]

Let $\theta \in \{H, L\}$ denote a consumer's type, where H denotes a high type and L a low type. Two types differ in the hassle cost of returning products h_θ. We further assume $h_H > h_L$ so that with the same perceived value high type consumers are more likely to keep a purchased product, whereas low type consumers are more likely to return a product. The proportion of high type consumers is λ and the proportion of low types is $1 - \lambda$.

In this section, we simplify the consumer value to be uniformly distributed on interval $[0, 1]$, that is, $G(v) = v$. If possible, the retailer could offer a price and refund profile (p_θ, r_θ) for type θ who then calculates her own expected utility $E[U_\theta(p_\theta, r_\theta; v)]$ based on equation (3.2), which is:

$$\begin{aligned} E[U_\theta(p_\theta, r_\theta; v)] &= E_v[\max\{v, r_\theta - h_\theta\}] - p \\ &= [(r_\theta - h_\theta)^2 - 2p_\theta + 1]/2. \end{aligned} \tag{3.9}$$

When a consumer with type θ purchases the product, the expected profit that the retailer can make from this consumer is calculated based on equation (3.3) as

$$\pi_\theta = p_\theta - c_0 + (c_0 - c_r - r_\theta) \cdot (r_\theta - h_\theta). \tag{3.10}$$

Notice from consumers' expected utility function that if the retailer sets the same price and refund (p, r) for each consumer type, we have $E[U_H(p, r; v)] < E[U_L(p, r; v)]$, which

[6]Hassle costs represent other factors that could cause consumers to have different propensities to return. For example, some consumers might need a product immediately and are willing to accept lower-quality ones, and other buyers might be willing to exchange products to try others.

means high type consumers have lower expected utility given the same price and refund. Because of this, high type consumers would be more careful when they make the purchase decision based on the information of the product on the retailer's website. In other words, high type consumers are less likely to buy online, but when buying online they are also less likely to return the purchased product.

Furthermore, the retailer's expected profit function implies that if $r < c_0 - c_r$, then $\pi_H(p, r) < \pi_L(p, r)$; and if $r > c_0 - c_r$, then $\pi_H(p, r) > \pi_L(p, r)$. It suggests that which type can bring more profit for the retailer depends on his return policy. The explanation is that the retailer could earn positive profit from returned items when refunds are low, while experience a loss for returned items when refunds are high, so it is more profitable to sell to low type consumers when refunds are low, while more profitable to sell to high type consumers when refunds are high. In addition, if $r = c_0 - c_r$, then $\pi_H(p, r) = \pi_L(p, r)$. That is, the retailer would be indifferent to two types if he sets a refund that is equal to the salvage value of returned items.

We examine two scenarios in the following based on whether the retailer could identify consumer types before purchase.

3.5.1 The Retailer Cannot Identify Consumers by Types

If the retailer cannot identify consumer types, he may consider one of the two strategies: either set one universal price and refund for all purchases, or set two profiles to screen two consumer types. We explore both strategies in the following.

One Universal Price and Return Policy

Based on expected utility expressed in equation (3.9), recall that if the retailer offers a university price and refund (p, r), of the two types of consumers, high type consumers would have lower expected utility. As a result, if the retailer sells to both types, he only needs to assure a non-negative expected utility for high type consumers, and his profit-maximization program is:

$$\max_{p,r} \pi = \lambda \cdot \pi_H + (1 - \lambda) \cdot \pi_L$$

$$\text{subject to } E[U_H] = [(r - h_H)^2 - 2p + 1]/2 \geqslant 0$$

where π_H, π_L are expected profit from each consumer type that are calculated based on equation (3.10).

We denote the optimal universal price, refund and related profit as (p^u, r^u, π^u), which are obtained as:[7]

$$p^u = [c_0 - c_r - (2 - \lambda)h_H + (1 - \lambda)h_L]^2/2 + 1/2,$$

$$r^u = c_0 - c_r - (1 - \lambda)(h_H - h_L),$$

$$\pi^u = [c_0 - c_r - (2 - \lambda)h_H + (1 - \lambda)h_L]^2/2 - c_0 + 1/2.$$

Comparing this result to the one with homogeneous consumers in the basic model (where $\lambda = 1$), the retailer lowers refund because of the existence of low type consumers. As a result, both price and profit decline. This is consistent with our observation from the

[7]For simplicity, we focus on the results with partial refunds and do not report the corner solutions with zero refund.

basic model that the retailer would be worse off from a population with a high probability of returning products.

Menu Pricing

Although the retailer cannot identify consumer types, in a menu pricing strategy, he can set up two different options (p_H, r_H) and (p_L, r_L) to screen consumers. In the real world business, this can be done by two websites or deals. To attract different types of consumers choose different options, the participation and incentive compatibility for menu pricing are as follows:

$$\begin{aligned} &\text{PCH} \quad E[U_H(p_H, r_H)] \geqslant 0, \\ &\text{ICH} \quad E[U_H(p_H, r_H)] \geqslant E[U_H(p_L, r_L)], \\ &\text{PCL} \quad E[U_L(p_L, r_L)] \geqslant 0, \text{ and} \\ &\text{ICL} \quad E[U_L(p_L, r_L)] \geqslant E[U_L(p_H, r_H)]. \end{aligned} \tag{3.11}$$

We first can prove the existence of menu pricing. We have shown that $E[U_H(p, r; v)] < E[U_L(p, r; v)]$ given the same price and refund (p, r). Therefore, if PCH is satisfied, then PCL is as well. Therefore, we need to consider only three constraints: PCH, ICH, and ICL. We depict three indifferent curves for $EU_H = 0, EU_L = 0$ and $EU_L(p_L, r_L) = EU_L(p_H, r_H) > 0$ in Figure 3.2, which shows the possibility that all constraints are satisfied. For instance, point H and L respectively represent the profile for (p_H, r_H) and (p_L, r_L), which are both on the curve of $EU_L > 0$, so it could be one candidate for menu pricing. We can tell that $r_H < r_L$ and $p_H < p_L$.

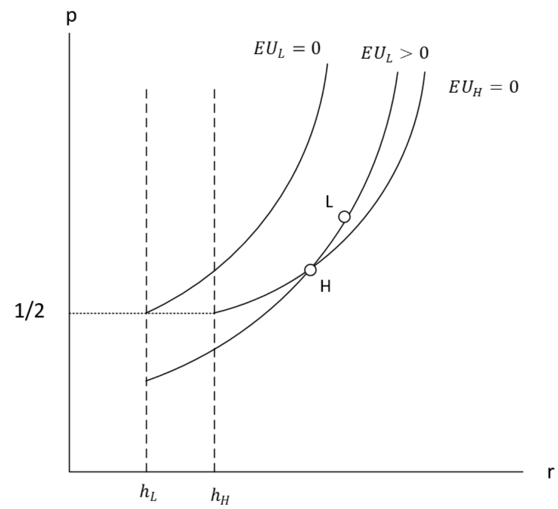

Figure 3.2: The Menu Pricing

Then we characterize the retailer's optimal menu price and refund mathematically. Given the possible menu pricing options, the retailer's profit-optimization program is:

$$\max_{p_H, r_H, p_L, r_L} \pi = \lambda \cdot \pi_H + (1 - \lambda) \cdot \pi_L$$

subject to PCH and ICL.

One menu pricing solution $(p_H^m, r_H^m, p_L^m, r_L^m)$ is obtained as:

$$r_H^m = c_0 - c_r - (h_H - h_L)(1 - \lambda)/\lambda,$$

$$r_L^m = c_0 - c_r,$$

$$p_H^m = [(r_H^m - h_H)^2 + 1]/2,$$

$$p_L^m = [(r_H^m - h_H)^2 + (r_L^m - h_L)^2 - (r_H^m - h_L)^2 + 1]/2.$$

The retailer sets a low price and low refund for high type consumers, while high price and high refund for low type consumers. This is consistent with the case in Figure 3.2.

What is the motivation to provide lower refund for high type consumers? One explanation is that low type consumers are more likely to return the product so that they are more sensitive to a refund, whereas high type consumers are more likely to keep the product so that they are more sensitive to price. Therefore, they have different preferences for price and refund policies. The retailer could use those preferences to screen the two consumer types.

Comparing the two practices when the retailer cannot identify consumer types, we observe that $\pi^m - \pi^u = (h_H - h_L)^2(1-\lambda)^3/2\lambda \geqslant 0$, which indicates the retailer could be better off from the menu pricing practice compared to uniform pricing. This provides an evidence to multiple-platform business, in which the retailer offers different price and return policies.

3.5.2 The Retailer Can Identify Consumers by Types

With the ability of retailers to track shopping behavior, retailers are able to identify consumer types based on their purchase histories. For example, Retail Equation, hired by some big retailers like Best Buy, is one such company that scores customer shopping behavior (Safdar, 2018). When the retailer can observe consumer purchase history, he may have the incentive to treat consumers differently by types.[8] In examining discriminatory pricing and return policies, we consider three retailer options: (1) sell to both types with the same price but different return policies, (2) sell to both types with different prices and return policies, and (3) restrict low-type consumers from purchasing.[9] We examine the first two

[8]This is similar to a "behavior-based price discrimination". Fudenberg and Villas-Boas (2006) survey the literature on this area.

[9]All practices exist in reality. For example, Amazon has banned some accounts, and some retailers do not offer free return for some consumers who have returned too many products in the past. More real world

cases separately in this section and the last option in a separate section.

Same Price, Different Return Policies

The first option is to set one price and two refunds, that is (p, r_H, r_L). This is easy to undertake when the retailer has collected the information to distinguish consumers. And we also observe that some retailers have adopted this strategy in the real world. Under this scenario, the retailer's profit-optimization program is:

$$\max_{p, r_H, r_L} \pi = \lambda \cdot \pi_H + (1 - \lambda) \cdot \pi_L$$

$$\text{subject to } EU_H \geqslant 0 \text{ and } EU_L \geqslant 0$$

We derive the optimal price, refund and related profit as:

$$p^d = [(c_0 - c_r - H)^2 + 1]/2,$$

$$r_H^d = c_0 - c_r + (1 - \lambda)(h_H - h_L),$$

$$r_L^d = c_0 - c_r - \lambda(h_H - h_L),$$

$$\pi^d = [(c_0 - c_r - H)^2 - 2c_0 + 1]/2, \qquad \text{where } H \equiv \lambda \cdot h_H + (1 - \lambda)h_L.$$

We observe $r_H^d - r_L^d = h_H - h_L > 0$, indicating that the retailer provides a higher refund for high type consumers. And surprisingly, the difference between two refunds has no relation with the proportion of high type consumers, but rather is determined by the retailer's costs and the consumers' hassle costs. The difference between two refunds is increasing in the difference between the hassle costs, which means a bigger consumer

examples are discussed in articles like Drake (2014), Safdar (2018) and Safdar and Stevens (2018).

heterogeneity.

Recall that when the retailer cannot identify consumer types, he would like to offer a lower refund for high type consumers in order to screen them from low types. It suggests that the retailer would use a different mechanism when he can identify consumer types: he would like to discourage low-type consumers from returning items with a lower refund while encouraging high-type consumers to purchase with higher refund.

Comparing this case to that the retailer sets a universal price and refund, we have $r_H^d - r^u = 2(1-\lambda)(h_H - h_L) \geqslant 0$, which implies that the retailer would increase the refund for high-type consumers when he can identify consumer types and set two different refunds. For low-type consumers, as $r_L^d - r^u = (1-2\lambda)(h_H - h_L)$, we would observe that if $\lambda \leqslant 1/2$, then $r_L^d \geqslant r^u$; and if $\lambda > 1/2$, then $r_L^d < r^u$. That is, the retailer would raise the refund for low-type consumers if there is a significantly small proportion of high-type consumers, while lower the refund for low-type consumers if more than one half of the consumers are high type. Moreover, as $\pi^d - \pi^u = (1-\lambda)(h_H - h_L)[c_0 - c_r - h_H] \geqslant 0$, the retailer is better off from this refund discrimination, even though he sells at the same price.

Finally, we investigate whether the retailer offers a "free return" policy in this scenario. The retailer may offer "free return" to only high-type consumers. Interestingly, the condition for the retailer to offer a "free return" for high-type consumers is not restrictive. The retailer may offer it even there is restocking cost,[10] which is not the case in the basic model. In addition, when it is optimal to offer "free return" policy, the retailer makes a positive profit from high-type consumers.[11] We have demonstrated in the basic model and

[10]The condition for $r_H^d = p^d$ is only $c_0 - c_r = \lambda h_H - \sqrt{2h_H} + (1-\lambda)h_L + 1$.

[11]The condition for a positive profit contributed by high-type consumers under "free return" is $\lambda \cdot (h_H - h_L)(1 - \sqrt{2h_H}) > c_0 - (h_H - h_L)\sqrt{2h_H} - h_L - 1$.

competition model that "free return" as an equilibrium is not profitable. This is the only scenario under which "free return" policy can make a positive profit. To confirm our result, we evaluate one numerical example: $c_0 = 3/5, h_H = 1/8, h_L = 0, \lambda < 3/5$, in which we can calculate that $\pi^d > \pi^u, r_H^d = p^d, \pi_H^d > 0$, which is a "free return" equilibrium for high-type consumers with a positive profit.

Different Prices and Return Policies

The second option that the retailer may consider is to set two prices and two refunds as (p_H, p_L, r_H, r_L). By doing so, the retailer completely separates two types and treats them as two market segments. Therefore, the retailer faces the following profit-maximization program:

$$\max_{p_H, p_L, r_H, r_L} \pi = \lambda \cdot \pi_H + (1 - \lambda) \cdot \pi_L$$

$$\text{s.t. } EU_H \geqslant 0 \text{ and } EU_L \geqslant 0$$

The optimal price, refund and related profit are:

$$\begin{aligned}
r_H^{dd} &= r_L^{dd} = c_0 - c_r, \\
p_H^{dd} &= [(c_0 - c_r - h_H)^2 + 1]/2, \\
p_L^{dd} &= \frac{1}{2}[(c_0 - c_r - h_L)^2 + 1], \\
\pi^{dd} &= [\lambda(c_0 - c_r - h_H)^2 + (1 - \lambda)(c_0 - c_r - h_L)^2 - 2c_0 + 1]/2.
\end{aligned}$$

Since each type is treated as a separate market, the optimal price and refund for each type is exactly the same as those in the basic model. Notice that it is optimal for the retailer to set same refund but different prices for two types. As $p_H^{dd} < p_L^{dd}$, the retailer would offer lower price for high type consumers, so that the return cost for high type consumers is lower. And no surprise that the condition for "free return" policy as the equilibrium is also the same as that in the basic model, but only for high-type consumers.

Moreover, it is verifiable that $\pi^{dd} > \pi^d$, so same refund with different prices is more profitable than same price with different return policies. Specifically, $r_H^{dd} - r_H^d = (\lambda - 1)(h_H - h_L) \leqslant 0$, $r_L^{dd} - r_L^d = \lambda(h_H - h_L) \geqslant 0$, which means when the retailer charges different prices instead of setting different refunds, he would reduce the refund for high-type consumers while increase the refund for low type.

3.5.3 Summary

The above analyses with heterogeneous return propensities provide us with the following implications. First, the retailer benefits from detecting consumer types and treating two types differently. Even if he cannot identify consumer types, he still has incentive to undertake a menu pricing to screen two types.

Second, when the retailer can identify consumer types, he chooses one of the two options: (1) one price with two return policies or (2) two prices with one return policy. If it is possible for him to engage in price and refund discrimination, then he chooses only price discrimination. However, when the retailer chooses the first option, it may be optimal and profitable for him to offer "free returns" to only the high type consumers.

Third, when the retailer can treat two types differently, he always prefers to offer a better return policy to high-type consumers by either a higher refund or a lower price. However, he chooses to do the opposite in a menu pricing design. Therefore, high type consumers may benefit from disclosing their types, whereas low type consumers may change their return behavior strategically.

3.6 "Fire" Low-Type Consumers

As we have discussed in the previous section, the retailer has another option to choose: restricting the purchase by low type consumers. A few studies have analyzed this business operation. For example, Shin, Sudhir and Yoon (2012) in a different context model the case to selectively "fire" high-cost customers by raising price to them, which could make the seller more profitable. In the popular press, Safdar and Stevens (2018) report accounts in which Amazon bans consumers based on their product-return histories from purchasing products. In this section, we explore whether the retailer could benefit from restricting low-type consumers. In our model, the retailer may selectively restrict the purchase by low type consumers who are more likely to return items. We investigate two operations the retailer may consider: (1) sell to both types but use a "no return" policy for low-type consumers, and (2) only sell to high-type consumers.

3.6.1 "No Return" for Low-Type Consumers

We start from the case in which the retailer sells to both types but does not allow low type consumers to return. He can achieve this by setting $r_L \leqslant h_L$. If doing so, the low type consumers would purchase only when the price is not above the expected product value, that is, $p_L \leqslant 1/2$. We still assume that the retailer can charge either one universal price or two differential prices. We investigate both cases in the following.

With One Price

If the retailer sets one universal price, the price should not below 1/2. Otherwise the low type consumers would not purchase. When $p \leqslant 1/2$, we calculate that $EU_H = [(r_H - h_H)^2 - 2p + 1]/2 \geqslant 0$, so the high-type consumers would always purchase the product. As a result, the retailer faces the following profit-maximization program:

$$\max_{p,r_H} \pi = p - c_0 + \lambda(c_0 - c_r - r_H)(r_H - h_H)$$

$$\text{subject to } p \leqslant 1/2$$

The optimal price, refund and related profit are:

$$r_H^{nr1} = (c_0 - c_r + h_H)/2,$$

$$p^{nr1} = 1/2,$$

$$\pi^{nr1} = \lambda(c_0 - c_r - h_H)^2/4 - c_0 + 1/2.$$

To evaluate whether this option is profitable, we compare this to the one that permits returns from low-type consumers. Since $\pi^{nr1} \leqslant \pi^d$, the retailer cannot be made better off from this "no return" policy. Therefore, one universal price is not an reasonable option for the retailer to ban returns from low-type consumers.

With Multiple Prices

We then examine the case in which the retailer charges two differential prices. As we have shown, the price for low type consumers should be no less than 1/2, so profit-maximization program for the retailer is:

$$\max_{p_H, p_L, r_H} \pi = \lambda \cdot \pi_H + (1 - \lambda) \cdot \pi_L$$

$$\text{subject to } EU_H \geqslant 0 \text{ and } p_L \leqslant 1/2$$

The optimal price, refund and related profit can be solved as:

$$r_H^{nr2} = c_0 - c_r,$$

$$p_H^{nr2} = [(c_0 - c_r - h_H)^2 + 1]/2,$$

$$p_L^{nr2} = 1/2,$$

$$\pi^{nr2} = [\lambda(c_0 - c_r - h_H)^2 - 2c_0 + 1]/2.$$

Again, we compare this option to the one that allows for low type consumers. We derive the result that $\pi^{nr2} \leqslant \pi^{dd}$, which suggests that "no returns" from low-type consumers with differential prices also cannot make the retailer better off.

Above two cases demonstrate that "no returns" from low-type consumers is not an optimal policy.

3.6.2 Sell to Only High-Type Consumers

If the retailer sells to only high-type consumers, then his objective is to maximize π_H, subject to the constraint that $EU_H \geqslant 0$. This is the case we have analyzed in the basic model, and the optimal price, refund and related profit are:

$$p^h = \frac{1}{2}[(c_0 - c_r - h_H)^2 + 1],$$

$$r^h = c_0 - c_r,$$

$$\pi^h = \frac{\lambda}{2}[(c_0 - c_r - h_H)^2 - 2c_0 + 1].$$

Notice that low-type consumers do not have incentive to purchase at this price, because $p^h \geqslant 1/2$. Therefore, the retailer only needs to set this price to screen low type consumers out.

In the case in which the retailer would set one price for both the high and low types, and the case in which the retailer sets two prices, he may want to "fire" low-type consumers from purchasing. Specifically,

1. In the scenario where the retailer sets up one price and two return policies, he would like to "fire" low type consumers if and only if $\pi^h(p^h, r^h) \geqslant \pi^d(p^d, r^d_H, r^d_L)$, which gives us $c_0 \geqslant [1 - \lambda(h_H - h_L)^2]/2$ and $c_0 - h_L - \sqrt{\lambda(h_H - h_L)^2 + 2c_0 - 1} \leqslant c_r \leqslant c_0 - h_L + \sqrt{\lambda(h_H - h_L)^2 + 2c_0 - 1}$;
2. In the scenario where the retailer sets up two prices and one return policy, he would like to "fire" low type consumers if and only if $\pi^h(p^h, r^h) \geqslant \pi^{dd}(p^{dd}_H, p^{dd}_L, r^{dd}_H, r^{dd}_L)$, which gives us $c_0 \geqslant 1/2$ and $c_0 - h_L - \sqrt{2c_0 - 1} \leqslant c_r \leqslant c_0 - h_L + \sqrt{2c_0 - 1}$.

The following proposition summarizes the conditions for the retailer to "fire" low-type consumers.

Proposition 8. *If c_0 and c_r are sufficiently large, the retailer has incentive to ban low-type consumers from purchasing.*

Proof. In Lemma 6 of the basic model, we find the equilibrium constraint for c_r as $0 \leqslant c_r \leqslant c_0 - h$. Therefore, we can simplify the conditions in above two scenarios to $c_r \geqslant c_0 - h_L - \sqrt{\lambda(h_H - h_L)^2 + 2c_0 - 1}$ and $c_r \geqslant c_0 - h_L - \sqrt{2c_0 - 1}$. In either case, we need a lower bound for c_r. □

Proposition 8 suggests that if the retailer's cost of acquiring or manufacturing the product and the cost of managing returns are sufficiently large, then the retailer chooses to "fire" low-type consumers.

3.7 Extensions

In this section we consider two extensions: heterogeneous restocking costs and whether the retailer should provide additional product information prior to consumer purchase decisions.

3.7.1 Heterogeneous Restocking Costs

In above analysis, we model consumers' heterogeneous return propensities through hassle cost, which affects the probability one consumer would return the purchased product. We extend this model by including consumer heterogeneity in other return behavior that could cause different restocking cost for the retailer. For instance, some fraudulent returns could be used or damaged so that the product value has dropped a lot. Those returns affect the retailer's profit so that he has incentive to tailor his return policy for consumers who are more likely to make fraudulent returns. Based on the same model setting with two consumer types, we denote the restocking cost caused by type θ as c_θ. We further assume

$c_L > c_H$ so that high type consumers cause lower restocking cost if they return the product, while low type consumers could cause fraudulent returns that more costly for the retailer to manage and resell returned products. All other settings maintain the same. And we still examine the retailer's optimal strategies when he can identify consumer types versus he cannot.

Because consumers' expected utility before purchase is not affected by the restocking cost, its expression stays the same as equation (3.9). However, when purchase occurs, the retailer's expected profit from each consumer that is determined by both hassle cost and restocking cost changes to

$$\pi_\theta = p_\theta - c_0 + (c_0 - c_\theta - r_\theta) \cdot (r_\theta - h_\theta) \tag{3.12}$$

The retailer cannot identify consumer types

For comparison, we first examine the case in which the retailer cannot identify consumer types and sets a universal price and refund. The optimal strategy and related profit can be obtained as:

$$\begin{aligned}
p^u &= \frac{1}{2}\{[c_0 - h_H + \lambda(h_H - c_H) + (1-\lambda)(h_L - c_L)]^2 + 1\}, \\
r^u &= c_0 - h_H + \lambda(h_H - c_H) + (1-\lambda)(h_L - c_L), \\
\pi^u &= \frac{1}{2}\{[c_0 - h_H + \lambda(h_H - c_H) + (1-\lambda)(h_L - c_L)]^2 - 2c_0 + 1\}.
\end{aligned}$$

All the findings in the basic model as well as heterogeneous return propensity model stay robust, and notice that the restocking cost plays a similar role as hassle cost. The retailer is likely to further lower refund when low type consumers cause higher restocking cost.

The retailer can identify consumer types

Then we examine the effect on price discrimination when the retailer can identify consumer types. Similarly, there are two options available, he can either set up one price with different refunds, which can be derived as:

$$
\begin{aligned}
p^d &= \frac{1}{2}[(c_0 - W)^2 + 1], \\
r_H^d &= c_0 + h_H - W, \\
r_L^d &= c_0 + h_L - W, \\
\pi^d &= \frac{1}{2}[(c_0 - W)^2 - 2c_0 + 1].
\end{aligned}
$$

$$\text{where } W \equiv \lambda(h_H + c_H) + (1-\lambda)(h_L + c_L);$$

Or the retailer can set up different prices with different refunds as:

$$
\begin{aligned}
p_H^{dd} &= \frac{1}{2}[(c_0 - h_H - c_H)^2 + 1], \\
r_H^{dd} &= c_0 - c_H, \\
p_L^{dd} &= \frac{1}{2}[(c_0 - h_L - c_L)^2 + 1], \\
r_L^{dd} &= c_0 - c_L, \\
\pi^{dd} &= \frac{1}{2}[\lambda(c_0 - h_H - c_H)^2 + (1-\lambda)(c_0 - h_L - c_L)^2 - 2c_0 + 1].
\end{aligned}
$$

Again, we observe the similar results as those in the model only with heterogeneous return propensities. It consistently suggests that the retailer would offer higher refund for high type consumers. Furthermore, the heterogeneity of restocking cost amplifies the return policy discrimination by offering a even higher (lower) refund for high (low) type consumers.

3.7.2 Providing More Product Information

When consumers are uncertain about the product value, a generous return policy effectively attracts consumers to purchase. However, this may not be the only option. An online retailer could provide more product information from which consumers get more accurate expectation about product valuation. As another extension, we explore whether it is optimal for the retailer to provide this information. The following analysis is based on the setting of the basic model in section 3.3.

Without this additional information, we assume a consumer's product value to be uniformly distributed on interval $[0, 2]$, that is, $G(v) = v/2$. However, if a consumer views this additional information, then she knows whether her product value is below or above the median. Specifically, consumers will be divided into two segments: those whose value is distributed on $[0, 1]$ and those whose value is distributed on $[1, 2]$. We denote these two groups as L and H, respectively for low and high value consumers. Those two groups would have different distributions: $v_L \sim \text{U}[0, 1]$ with $G_L(v) = v$, while $v_H \sim \text{U}[1, 2]$ with $G_H(v) = v - 1$. We assume that the prior belief of uniform distribution is correct, so consumers are equally divided into two groups.

Given this additional information, we investigate the retailer's optimal return policy through two available options: (1) with a high refund that $r - h > 1$, or (2) with a low refund that $r - h \leqslant 1$.

High refund: $r - h > 1$

If the retailer chooses the first option with high refund, the low value consumers would always return the product if they purchase it. So the retailer would not sell to the low value consumers but only to high value consumers. Therefore, the retailer faces the following profit-maximization program:

$$\max_{p,r} \pi = [p - c_0 + (c_0 - c_r - r)(r - h - 1)]/2$$

$$\text{subject to } EU_H = [(r - h - 1)^2 - 2p + 3]/2 \geqslant 0$$

We solve for the optimal price, return policy and profit as:

$$p^{hr} = (c_0 - c_r - h + 1)^2/2,$$

$$r^{hr} = c_0 - c_r,$$

$$\pi^{hr} = (c_0 - c_r - h + 1)^2/4 - c_0/2.$$

And the condition for this outcome to be the equilibrium is $0 \leqslant c_r < c_0 - h - 1$. And the retailer only sells to high value consumers. Since the equilibrium profit in the basic model is $\pi^* = 1 - c_0 + (c_0 - c_r - h)^2/4$, then $\pi^{hr} > \pi^*$ if $0 \leqslant c_r < 2c_0 - h - 3/2$. In other words, the retailer is better off by providing additional information when the restocking cost is sufficiently small.

Low refund: $r - h \leqslant 1$

If the retailer chooses the second option with low refund, all consumers may purchase, but high value consumers always keep the product. Given the same price and refund, the expected utility of high value consumers, $3/2 - p$, would be higher than that of low value consumers, $[(r - h)^2 - 2p + 1]/2$. As a result, the retailer would sell to both types with the following profit-maximization program:

$$\max_{p,r} \pi = (p - c_0)/2 + [p - c_0 + (c_0 - c_r - r)(r - h)]/2$$

$$\text{subject to } EU_L = [(r - h)^2 - 2p + 1]/2 \geqslant 0$$

The optimal price and return policy can be obtained as a corner solution, which is:

$$p^{lr} = 1/2, r^{lr} = 0, \pi^{lr} = 1/2 - c_0.$$

And the condition for this outcome to be the equilibrium is $c_0 \leqslant 1/2$ and $c_r \geqslant c_0 - h$. In this equilibrium, the retailer offers a "no return" policy. However, since the high-value consumers never return with a low refund, this policy only constrains low-value consumers. Comparing to the basic model, since $\pi^{lr} < \pi^*$, when the product cost is low and restocking cost is high, the retailer does not have incentive to provide this additional information.

Based on above analysis of the two options, providing additional information is beneficial for the retailer only when the restocking cost is low, and the retailer sets a high refund for high value consumers.

3.8 Concluding Remarks

Although a generous return policy encourages consumers to purchase, it reduces a retailer's total profit by dealing with more returns. We examine two factors that could affect the profitability of return policy: return propensity and cost of returns. In a scenario in which consumers are *ex ante* identical, we find that a partial return policy is favored by the retailer, which implies that a "free return" policy is suboptimal in both the monopoly and Bertrand competition cases. The optimal return policy is more liberal for expensive products that are easy to maintain the value in returns.

With consumer heterogeneity, the retailer is better off from detecting consumers' return behavior and treating them differently based on return propensities or restocking costs. Consumers who are less likely to return products or abuse returns are more likely to be provided with a better return policy with either a higher refund or a lower price. As high-type consumers may benefit from disclosing their type, doing so may have a strategic implication for buyers to return fewer products.

As we observe in practice, a retailer may want to "fire" consumers who are more likely to return products and who can cause huge depreciation for returned items. Our results indicate that this policy is optimal for expensive products in which retailers bear substantial costs for product returns.

This paper is an attempt to model return policy within a new market circumstance, but the implication is not limited to only return policies. Actually, any other markets with uncertainty on both sides of the market, in which buyers at the time of purchase are uncertain about product value and firms are uncertain about consumers' purchase behavior, may have similar mechanisms. Although we have investigated a number of model settings, there is still much work to be done. We noticed that further research should focus on the strategic interaction between the retailer and buyers. With the development of purchase-tracking technologies, either the retailer's return policy or consumers' return behavior could have a signaling effect.

Bibliography

Abeler, Johannes, Daniele Nosenzo, and Collin Raymond. 2016. "Preferences for truth-telling." Unpublished Paper.

Ackerberg, Daniel A. 2001. "Empirically distinguishing informative and prestige effects of advertising." *RAND Journal of Economics*, 316–333.

Anderson, Eric T, Karsten Hansen, and Duncan Simester. 2009. "The option value of returns: Theory and empirical evidence." *Marketing Science*, 28(3): 405–423.

Argenziano, Rossella, Sergei Severinov, and Francesco Squintani. 2016. "Strategic information acquisition and transmission." *American Economic Journal: Microeconomics*, 8(3): 119–55.

Bagwell, Kyle. 2007. "The economic analysis of advertising." *Handbook of Industrial Organization*, 3: 1701–1844.

Baye, Michael R, and John Morgan. 2009. "Brand and price advertising in online markets." *Management Science*, 55(7): 1139–1151.

Bechwati, Nada Nasr, and Wendy Schneier Siegal. 2005. "The impact of the prechoice process on product returns." *Journal of Marketing Research*, 42(3): 358–367.

Becker, Gary S, and Kevin M Murphy. 1993. "A simple theory of advertising as a good or bad." *The Quarterly Journal of Economics*, 108(4): 941–964.

Bloch, Francis, and Delphine Manceau. 1999. "Persuasive advertising in Hotelling's model of product differentiation." *International Journal of Industrial Organization*, 17(4): 557–574.

Blume, Andreas, Ernest K Lai, and Wooyoung Lim. 2017. "Strategic information transmission: A survey of experiments and theoretical foundations." Unpublished Paper.

Board, Oliver. 2009. "Competition and disclosure." *The Journal of Industrial Economics*, 57(1): 197–213.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience theory of choice under risk." *Quarterly Journal of Economics*, 127(3): 1243–1285.

Bowles, Samuel, and Sandra Polania-Reyes. 2012. "Economic incentives and social preferences: substitutes or complements?" *Journal of Economic Literature*, 50(2): 368–425.

Butters, Gerard R. 1977. "Equilibrium distributions of sales and advertising prices." *Review of Economic Studies*, 44(3): 465–491.

Cai, Hongbin, and Joseph Tao-Yi Wang. 2006. "Overcommunication in strategic information transmission games." *Games and Economic Behavior*, 56(1): 7–36.

Camerer, Colin, and Teck Hua Ho. 1999. "Experience-weighted attraction learning in normal form games." *Econometrica*, 67(4): 827–874.

Camerer, Colin F. 2003. *Behavioral game theory: Experiments in strategic interaction.* Princeton University Press.

Chen, Daniel L, Martin Schonger, and Chris Wickens. 2016. "oTree—An open-source platform for laboratory, online, and field experiments." *Journal of Behavioral and Experimental Finance*, 9: 88–97.

Chen, Yubo, and Jinhong Xie. 2005. "Third-party product review and firm marketing strategy." *Marketing Science*, 24(2): 218–240.

Chen, Yubo, and Jinhong Xie. 2008. "Online consumer review: Word-of-mouth as a new element of marketing communication mix." *Management Science*, 54(3): 477–491.

Chen, Yubo, Yong Liu, and Jurui Zhang. 2012. "When do third-party product reviews affect firm value and what can firms do? The case of media critics and professional movie reviews." *Journal of Marketing*, 76(2): 116–134.

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and taxation: Theory and evidence." *American Economic Review*, 99(4): 1145–77.

Che, Yeon-Koo. 1996. "Customer return policies for experience goods." *The Journal of Industrial Economics*, 44(1): 0021–1821.

Chu, Wujin, Eitan Gerstner, and James D Hess. 1998. "Managing dissatisfaction: How to decrease customer opportunism by partial refunds." *Journal of Service Research*, 1(2): 140–155.

Conlon, Christopher T, and Julie Holland Mortimer. 2013. "Demand estimation under incomplete product availability." *American Economic Journal: Microeconomics*, 5(4): 1–30.

Crawford, Vincent. 1998. "A survey of experiments on communication via cheap talk." *Journal of Economic Theory*, 78(2): 286–298.

Crawford, Vincent P, and Joel Sobel. 1982. "Strategic information transmission." *Econometrica*, 1431–1451.

Davis, Scott, Michael Hagerty, and Eitan Gerstner. 1998. "Return policies and the optimal level of "hassle"." *Journal of Economics and Business*, 50(5): 445–460.

Dearden, James A, Rajdeep Grewal, and Gary L Lilien. 2019, forthcoming. "Strategic manipulation of university rankings, the prestige effect, and student university choice." *Journal of Marketing Research.*

Dilly, Mark. 2014. "Rating Quality: A Literature Review." Unpublished Paper.

Drake, Samantha. 2014. "Should you offer free return shipping? What to consider." *Enterepreneur*, available at https://www.entrepreneur.com/article/238701.

Dranove, David, and Ginger Zhe Jin. 2010. "Quality disclosure and certification: Theory and practice." *Journal of Economic Literature*, 48(4): 935–63.

Engel, Christoph. 2011. "Dictator games: A meta study." *Experimental Economics*, 14(4): 583–610.

Fudenberg, Drew, and J Miguel Villas-Boas. 2006. "Behavior-based price discrimination and customer recognition." *Handbook on economics and information systems*, 1: 377–436.

Galbraith, John Kenneth. 1976. "The Affluent Society (New York: New American Library, 1970)." *The New Industrial State.*

Gentzkow, Matthew, and Emir Kamenica. 2014. "Costly persuasion." *American Economic Review*, 104(5): 457–62.

Ghose, Anindya, Panagiotis G Ipeirotis, and Beibei Li. 2014. "Examining the impact of ranking on consumer behavior and search engine revenue." *Management Science*, 60(7): 1632–1654.

Gneezy, Uri, Agne Kajackaite, and Joel Sobel. 2018. "Lying Aversion and the Size of the Lie." *American Economic Review*, 108(2): 419–53.

Hopkins, Ed. 2002. "Two competing models of how people learn in games." *Econometrica*, 70(6): 2141–2166.

Hörner, Johannes, and Andrzej Skrzypacz. 2016. "Selling information." *Journal of Political Economy*, 124(6): 1515–1562.

Ho, Teck H, Xin Wang, and Colin F Camerer. 2007. "Individual differences in EWA learning with partial payoff information." *The Economic Journal*, 118(525): 37–59.

Hotz, V Joseph, and Mo Xiao. 2013. "Strategic information disclosure: The case of multiattribute products with heterogeneous consumers." *Economic Inquiry*, 51(1): 865–881.

Jin, Ginger Zhe, and Phillip Leslie. 2003. "The effect of information on product quality: Evidence from restaurant hygiene grade cards." *Quarterly Journal of Economics*, 118(2): 409–451.

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

Kuksov, Dmitri, and J Miguel Villas-Boas. 2010. "When more alternatives lead to less choice." *Marketing Science*, 29(3): 507–524.

Kuksov, Dmitri, and Kangkang Wang. 2013. "A model of the "it" products in fashion." *Marketing Science*, 32(1): 51–69.

Lai, Ernest K, Wooyoung Lim, and Joseph Tao-yi Wang. 2015. "An experimental analysis of multidimensional cheap talk." *Games and Economic Behavior*, 91: 114–144.

Levin, Dan, James Peck, and Lixin Ye. 2009. "Quality disclosure and competition." *The Journal of Industrial Economics*, 57(1): 167–196.

Levitt, Steven D, and John A List. 2007. "What do laboratory experiments measuring social preferences reveal about the real world?" *Journal of Economic Perspectives*, 21(2): 153–174.

Luca, Michael, and Jonathan Smith. 2013. "Salience in quality disclosure: evidence from the US News college rankings." *Journal of Economics & Management Strategy*, 22(1): 58–77.

Mayzlin, Dina, and Jiwoong Shin. 2011. "Uninformative advertising as an invitation to search." *Marketing Science*, 30(4): 666–685.

Nelson, Phillip. 1974. "Advertising as information." *Journal of Political Economy*, 82(4): 729–754.

Ostrovsky, Michael, and Michael Schwarz. 2010. "Information disclosure and unraveling in matching markets." *American Economic Journal: Microeconomics*, 2(2): 34–63.

Padmanabhan, Venkata, and Ivan PL Png. 1997. "Manufacturer's return policies and retail competition." *Marketing Science*, 16(1): 81–94.

Petersen, J Andrew, and V Kumar. 2009. "Are product returns a necessary evil? Antecedents and consequences." *Journal of Marketing*, 73(3): 35–51.

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

Rust, Roland T, Anthony J Zahorik, and Timothy L Keiningham. 1996. *Service marketing.* HarperCollins.

Safdar, Khadeeja. 2018. "How Your Returns Are Used Against You at Best Buy, Other Retailers." *The Wall Street Journall*, available at https://www.wsj.com/articles/how–your–returns–are–used–against–you–at–best–buy–other–retailers–1520933400.

Safdar, Khadeeja, and Laura Stevens. 2018. "Banned From Amazon: The Shoppers Who Make Too Many Returns." *The Wall Street Journal*, available at https://www.wsj.com/articles/banned–from–amazon–the–shoppers–who–make–too–many–returns–1526981401.

Salmon, Timothy C. 2001. "An evaluation of econometric models of adaptive learning." *Econometrica*, 69(6): 1597–1628.

Sánchez-Pagés, Santiago, and Marc Vorsatz. 2007. "An experimental study of truth-telling in a sender-receiver game." *Games and Economic Behavior*, 61(1): 86–112.

Shin, Jiwoong, K Sudhir, and Dae-Hee Yoon. 2012. "When to "fire" customers: Customer cost-based pricing." *Management Science*, 58(5): 932–947.

Simonsohn, Uri. 2011. "Lessons from an "oops" at Consumer Reports: Consumers follow experts and ignore invalid information." *Journal of Marketing Research*, 48(1): 1–12.

Sobel, Joel. 2018. "Lying and deception in games." *University of California-San Diego, Working Paper.*

Stigler, George J. 1961. "The economics of information." *Journal of Political Economy*, 69(3): 213–225.

Su, Xuanming. 2009. "Consumer returns policies and supply chain performance." *Manufacturing & Service Operations Management*, 11(4): 595–612.

Tierney, John. 2013. "Your annual reminder to ignore the US News & World Report college rankings." *The Atlantic*.

Wood, Stacy L. 2001. "Remote purchase environments: The influence of return policy leniency on two-stage decision processes." *Journal of Marketing Research*, 38(2): 157–169.

Zhu, Feng, and Xiaoquan Zhang. 2010. "Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics." *Journal of Marketing*, 74(2): 133–148.

Chapter 4: Appendix

A Proofs and Additional Characterizations

A.1 Proofs

Proof of Lemma 1. For Part 1, it is straightforward that when $\alpha = 0$, $\mu_{r_{A||B}} = q_{r_{A||B}}$ and $\mu_{\bar{v}_B} = p$. Given these beliefs, the consumer's expected payoffs from purchasing Products A and B are, respectively, $\bar{v}_A + [q_{r_{A||B}}(1-p) + (1-q_{r_{A||B}})]\gamma$ and $p(\bar{v}_B + q_{r_{A||B}}\gamma)$. If $p\bar{v}_B \leqslant \bar{v}_A + (1-2p)\gamma$, it is a best response for the consumer to purchase Product A for all $q_{r_{A||B}} \in [0,1]$.

For Part 2, consider first the case where the consumer sees that Product A is ranked first. Bayes' rule implies that $\mu_{\bar{v}_B} = \frac{p - pq_{r_{A||B}}}{1 - pq_{r_{A||B}}}$. Given this belief, the consumer's expected payoffs from purchasing Products A and B are, respectively, $\bar{v}_A + \gamma - c$ and $\left(\frac{p - pq_{r_{A||B}}}{1 - pq_{r_{A||B}}}\right)\bar{v}_B - c$. If $\bar{v}_A + \gamma \geqslant p\bar{v}_B$, it is a best response for the consumer to purchase Product A for all $q_{r_{A||B}} \in [0,1]$. Consider next the case where the consumer sees that Product B is ranked first. Bayes' rule implies that $\mu_{\bar{v}_B} = 1$ for $q_{r_{A||B}} > 0$. For $q_{r_{A||B}} = 0$, we assign an off-path product belief that $\mu_{\bar{v}_B} = 1$. Given these beliefs, the consumer's expected payoffs from purchasing Products A and B are, respectively, $\bar{v}_A - c$ and $\bar{v}_B + \gamma - c$. If $\bar{v}_B + \gamma \geqslant \bar{v}_A$, it is a best response for the consumer to purchase Product B for all $q_{r_{A||B}} \in [0,1]$. Combining $\bar{v}_A + \gamma \geqslant p\bar{v}_B$ and $\bar{v}_B + \gamma \geqslant \bar{v}_A$ gives $\gamma \geqslant \max\{p\bar{v}_B - \bar{v}_A, \bar{v}_A - \bar{v}_B\}$.

□

Proof of Lemma 2. It is straightforward to verify that, given that $\rho(\varnothing) = \rho(A) = A$ and $\rho(B) = B$, the consumer's expected payoff from choosing $\alpha = 1$ minus that from choosing

$\alpha = 0$, $\{q_{r_{A||B}}[p\bar{v}_B+(1-p)\bar{v}_A]+(1-q_{r_{A||B}})\upsilon_A+\gamma-c]\}-\{\bar{v}_A+[q_{r_{A||B}}(1-p)+(1-q_{r_{A||B}})]\gamma\} \gtreqless 0$ if and only if $q_{r_{A||B}} \gtreqless \frac{c}{p(\bar{v}_B+\gamma-\bar{v}_A)}$.

□

Proof of Proposition 1. We construct the equilibria. Consider first the parameter values such that $\gamma \geqslant \max\{p\bar{v}_B - \bar{v}_A, \bar{v}_A - \bar{v}_B\}$ and $p\bar{v}_B \leqslant \bar{v}_A + (1-2p)\gamma$. Suppose that $q_{r_{A||B}} = 1$. According to Lemma 1, it is a best response for the consumer to adopt $\rho(\varnothing) = \rho(A) = A$ and $\rho(B) = B$. It is easy to show that $\bar{v}_A + \gamma > p\bar{v}_B$. Together with $\bar{v}_A + \gamma > \bar{v}_A + \gamma - c$, we have ($r_A$, (Not Acquire, A)) as one equilibrium. Given Lemma 2, we have (r_A, (Acquire, First)) as another equilibrium.

□

A.2 Additional Equilibrium Characterizations

The intuition for Part 2 is similar except that now acquiring the report makes a difference only when the report is A, which occurs when the pre-ranking value of Product B is 0 or when it is $\bar{v}_B$ and the expert chooses r_A. By acquiring the report in these cases, the consumer purchases the first-ranked Product A instead of Product B, receiving $\bar{v}_A + \gamma$ rather than, in the first case, 0 or, in the second case, $\bar{v}_B$. The requirement that the cost of the report is no larger than the expected benefit gives rise to the condition in the lemma.

The consumer prefers to acquire the ranking report ($\alpha = 1$) if and only if

1. $c \leqslant pq_{r_{A||B}}(\bar{v}_B + \gamma - \bar{v}_A)$ under product-purchasing rule ; or
2. $c \leqslant (1-p)(\bar{v}_A+\gamma)+p(1-q_{r_{A||B}})(\bar{v}_A+\gamma-\bar{v}_B)$ under product-purchasing rule $\rho(A) = A$ and $\rho(\varnothing) = \rho(B) = B$.

Consider $\rho(A) = A$ and $\rho(\varnothing) = \rho(B) = B$ for Part 2. Given $q_{r_{A||B}} \in [0,1]$, the consumer's expected payoff from choosing $\alpha = 0$, $p(\bar{v}_B + q_{r_{A||B}}\gamma)$, is no larger than that from choosing $\alpha = 1$, $q_{r_{A||B}}[p\bar{v}_B + (1-p)\bar{v}_A] + (1-q_{r_{A||B}})\upsilon_A + \gamma - c$, if and only if $c \leqslant (1-p)(\bar{v}_A + \gamma) + p(1-q_{r_{A||B}})(\bar{v}_A + \gamma - \bar{v}_B)$.

B Experimental Instructions for Treatment *AVERAGE-B*

INSTRUCTIONS

Welcome to this experiment, which studies decision making between two individuals. The experiment will last approximately 1 hour. There will be 40 rounds of decisions. Please read these instructions carefully. The cash payment you receive at the end of the experiment depends on your decisions.

Your Role and Decision Group

There are 20 participants in today's session. 10 participants will be randomly assigned the role of **Product Expert**, and the other 10 the role of **Consumer**. Your role will remain fixed throughout the experiment. In each and every round, one Product Expert is matched with one Consumer. Participants will be randomly rematched after each round. You will not learn the identity of the participant you are matched with in any round, nor will that participant learn your identity.

Your Decision in Each Round

Overview. There are two products, A and B. They differ by the values to the Consumer. The value of Product A is fixed at 100.

The value of Product B is either 0 or 80. In each and every round, the computer will randomly select one of these two values for Product B according to the following:

(a) The chance that 0 will be selected is 40%.

(b) The chance that 80 will be selected is 60%.

The Product Expert makes one decision: how to rank the two products. The Consumer makes two decisions: whether to purchase to view the Product Expert's ranking report and which product to choose.

Product Expert's Decision: Choosing a Ranking Method. In each and every round, you choose a **ranking method**. There are two choices:

(a) Method 1: Rank Product A first regardless of the value of Product B.

(b) Method 2:

  (i) Rank Product A first if the value of Product B is selected to be 0.

  (ii) Rank Product B first if the value of Product B is selected to be 80.

Your decision therefore comes down to which product to rank first in the case that the value of Product B is selected to be 80.

You choose a ranking method before the computer selects the value for Product B, i.e., you don't know the selected value of Product B when you choose. Your task for the round is completed after you make the choice.

Depending on the selected value of Product B, a **ranking report** is generated according to your chosen ranking method. If the Consumer chooses to purchase to view your ranking report, the generated report, either "Product A is ranked 1st" or "Product B is ranked 1st," will be revealed to the Consumer.

Note that during the round the Consumer will never see the selected value of Product B. Note also that even if the Consumer chooses not to purchase your ranking report, the product ranking will still affect him/her. This will be further explained below.

Consumer's Decisions: Whether to Purchase the Ranking Report and Choosing a Product. In each and every round, your first decision is whether to purchase to view the ranking report (Purchase or Pass). You need to pay for the report if you choose "Purchase," which will be further explained below.

Irrespective of your first decision, you will make a second decision on which product, A or B, to choose. While you will receive the value of the product you choose, both products are free—you don't need to pay for it.

If you purchase the ranking report, you will make your product choice after seeing the generated report (you will only see the ranking report, not the ranking method that generates it). If you pass, you will make your product choice right after you choose "Pass." In either case, you will make your product choice without seeing directly the selected value of Product B. Your task for the round is completed after you make your product choice.

Your Reward in Each Round

Your reward in each round is expressed in "experimental currency unit" (ECU). How your earned ECU converts into cash payment will be explained below.

Product Expert's Reward. You will receive the amount the Consumer pays for the report, 70 ECU, multiplied by 5. Thus, if the Consumer chooses to purchase your ranking report, you will receive a total of 350 ECU for the round.

Table A-1 summarizes your potential reward in a round as a Product Expert, which depends on the Consumer's choice of *Purchase* or *Pass*.

Consumer's Decision on Purchasing Ranking Report	
Purchase	*Pass*
350	0

Table A-1: Product Expert's Potential Reward in ECU

Consumer's Reward. Your reward in a round consists of three parts:

(a) You will receive the value of the product you choose:

  (i) 100 ECU if you have chosen Product A; and

  (ii) either 0 or 80 ECU if you have chosen Product B, depending on the value randomly selected by the computer according to the 40%–60% chance.

(b) You will receive an extra reward of 250 ECU if you have chosen the first ranked product (regardless of whether you purchase the ranking report or not).

(c) You will pay a cost of 70 ECU if you choose to purchase the ranking report.

Table A-2 on the next page summarizes your potential reward in a round as a Consumer, which depends on

(a) your decision on purchasing the ranking report (*Purchase* or *Pass*);

(b) your product choice (*A* or *B*);

(c) the Product Expert's ranking method (*Method 1* or *Method 2*);

(d) chance (the randomly selected value of Product B); and

(e) the interaction between (c) and (d) which determines the generated report (*"Product A is Ranked 1st"* or *"Product B is Ranked 1st"*).

	Product Expert's Choice of Ranking Method and the Generated Report		
	Method 1	*Method 2*	
	"Product A is Ranked 1st" (always generated and seen by you)	*"Product A is Ranked 1st" (generated with 40% chance and seen by you)*	*"Product B is Ranked 1st" (generated with 60% chance and seen by you)*
Purchase & A	280	280	30
Purchase & B	-70 with 40% chance 10 with 60% chance	-70	260
	Method 1	*Method 2*	
	"Product A is Ranked 1st" (always generated and NOT seen by you)	*"Product A is Ranked 1st" (generated with 40% chance and NOT seen by you)*	*"Product B is Ranked 1st" (generated with 60% chance and NOT seen by you)*
Pass & A	350	350	100
Pass & B	0 with 40% chance 80 with 60% chance	0	330

Your Decisions

Table A-2: Consumer's Potential Reward in ECU

Information Feedback

At the end of each round, you will be provided with a summary of what happened in the round, including the Product Expert's ranking method, the selected value of Product B, the generated ranking report, the Consumer's decision on purchasing the ranking report, the Consumer's product choice, and your reward in the current round. A history of all previous rounds will also be provided.

Your Cash Payment

The experimenter randomly selects 3 rounds out of the 40 rounds to calculate your cash payment. So it is in your best interest to take each round equally seriously. The average of the ECU you earn in the 3 selected rounds will be converted into U.S. dollar at an exchange rate of 20 ECU for 1 USD. You will also separately receive a "show-up fee" of 5 USD.

Quiz and Practice

To ensure your understanding of the instructions, we will provide you with a quiz below. After the quiz, you will participate in a practice round. The practice round is part of the instructions and is not relevant to your earnings. Its objective is to get you familiar with the computer interface and the flow of the decisions in each round.

Once the practice round is over, the computer will ask you to "Click 'Next' to start the official rounds."

C Analysis for the case $\tilde{\Delta}_1 = \tilde{\Delta}_2$

We consider $\tilde{\Delta}_1 = \tilde{\Delta}_2$ as a special case because under this condition (1) no firm has comparative advantage and (2) consumers' purchase decision will not be influenced by attribute preference w but only by prices. As a result, we call this as a Bertrand competition case.

When $\tilde{\Delta}_1 = \tilde{\Delta}_2$, $U_A - U_B = \tilde{\Delta}_1 - (P_A - P_B)$, so the demand for each firm is

$$\begin{cases} D_A = 1, D_B = 0 & \text{if } P_A - P_B \leqslant \tilde{\Delta}_1; \\ D_A = 0, D_B = 1 & \text{if } P_A - P_B > \tilde{\Delta}_1. \end{cases}$$

Note that this demand could be integrated into demand function (2.3). Therefore, given Assumption 1, demand function (2.3) characterizes all the market share possibilities.

The following proposition characterize the price equilibrium for this case.

Proposition 9. *When consumers have expectations $\tilde{\Delta}_1$ and $\tilde{\Delta}_2$, if $\tilde{\Delta}_1 = \tilde{\Delta}_2$, it is a Bertrand competition and there exists a unique price equilibrium as:*

$$P_A^* = \tilde{\Delta}_1, P_B^* = 0.$$

Proof. We start from firm B. In order to obtain a positive market share, firm B needs to set a price so that $P_A - P_B > \tilde{\Delta}_1$, that is, $P_B < P_A - \tilde{\Delta}_1$, which is negative when $P_A = \tilde{\Delta}_1$. As firm B does not have incentive to obtain a positive market share with a negative profit, it does not have incentive to lower price. Then we analyze firm A. First of all, it does not have incentive to reduce price when it occupies the entire market. Secondly, we can prove that it does not want to increase the price either. Because if it increases the price, firm B would take over the entire market, so that firm A would lose all his profits. Therefore, firm A is better off to charge $\tilde{\Delta}_1$. $\square$

Notice that this equilibrium is just a special case of the general equilibrium that characterized by Proposition 2. Therefore, we do not lose generality by excluding this Bertrand competition.

D Outcomes for the case $\tilde{\Delta}_1 < \tilde{\Delta}_2$

If $\tilde{\Delta}_1 < \tilde{\Delta}_2$, the demand for each firm is given by:

$$\begin{cases} D_A = 1, D_B = 0 & \text{if } P_A - P_B \leqslant \tilde{\Delta}_1; \\ D_A = \hat{w}, D_B = 1 - \hat{w} & \text{if } \tilde{\Delta}_1 < P_A - P_B \leqslant \tilde{\Delta}_2; \\ D_A = 0, D_B = 1 & \text{if } P_A - P_B > \tilde{\Delta}_2. \end{cases}$$

To maximize profits in equation (2.4), two firms' best-response functions are:

$$P_A^{BR} = \frac{1}{2}(P_B + \tilde{\Delta}_2)$$
$$P_B^{BR} = \frac{1}{2}(P_A - \tilde{\Delta}_1)$$

Solving for the price equilibrium, we obtain:

$$P_A^* = \frac{2\tilde{\Delta}_2 - \tilde{\Delta}_1}{3},\ P_B^* = \frac{\tilde{\Delta}_2 - 2\tilde{\Delta}_1}{3},$$
$$\pi_A^* = \frac{(2\tilde{\Delta}_2 - \tilde{\Delta}_1)^2}{9(\tilde{\Delta}_2 - \tilde{\Delta}_1)},\ \pi_B^* = \frac{(\tilde{\Delta}_2 - 2\tilde{\Delta}_1)^2}{9(\tilde{\Delta}_2 - \tilde{\Delta}_1)}.$$

And the market share is divided by the indifferent consumer which is given by:

$$\hat{w}^* = \frac{\tilde{\Delta}_1 - 2\tilde{\Delta}_2}{3(\tilde{\Delta}_1 - \tilde{\Delta}_2)}.$$

The condition for this equilibrium to exist is:

$$\tilde{\Delta}_2 > 2\tilde{\Delta}_1$$
$$(\text{given by: } \tilde{\Delta}_1 < P_A - P_B < \tilde{\Delta}_2)$$

The derivatives are:

$$sign(\partial\pi_A^*/\partial\tilde{\Delta}_1) = sign(\tilde{\Delta}_1),\ \partial\pi_A^*/\partial\tilde{\Delta}_2 > 0;$$
$$\partial\pi_B^*/\partial\tilde{\Delta}_1 < 0,\ sign(\partial\pi_B^*/\partial\tilde{\Delta}_2) = sign(\tilde{\Delta}_2);$$
$$\partial\hat{w}^*/\partial\tilde{\Delta}_1 > 0,\ \partial\hat{w}^*/\partial\tilde{\Delta}_2 < 0.$$

This equilibrium result is symmetric to the main analysis with a switch between two firms. Therefore, we do not lose the generality from Assumption 1.

Additional Figures

Figure A-1: Screenshot of Subaru's Official Website

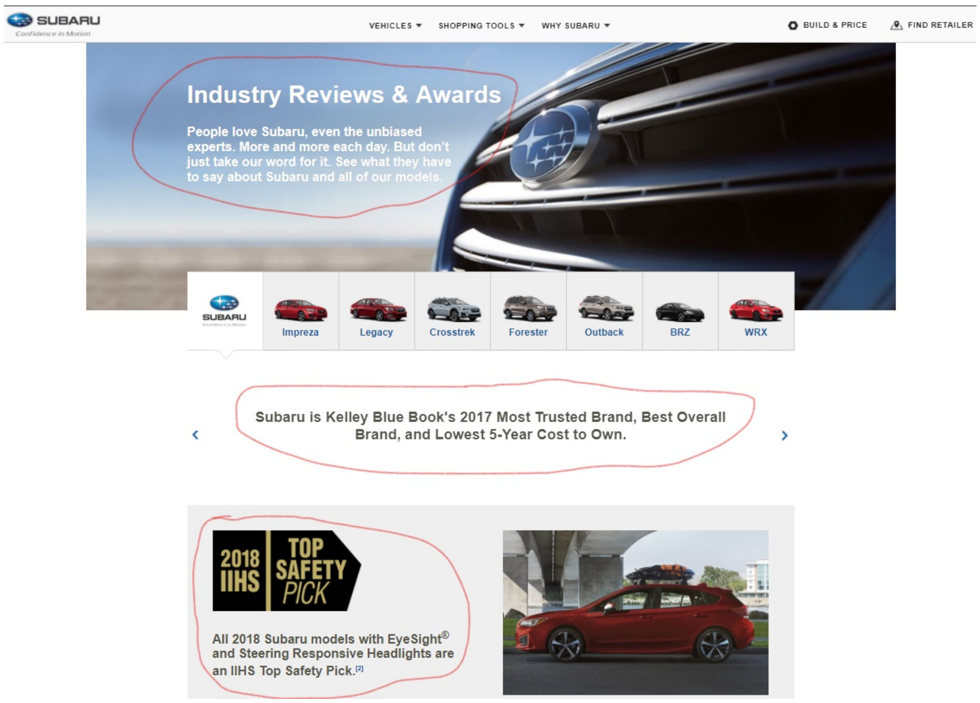

E Competition: Online and Offline Retailers

In this section, we consider the competition between online and offline retailers. When choosing either channel for shopping, consumers are heterogeneous in two dimensions: (1) they have different travel cost, $t \sim U[0, 1]$, when shopping in offline stores, while there is no travel cost when shopping online; (2) they have different *ex post* valuations of the product, $v \in \{\underline{v}, \bar{v}\}$, and only consumers with valuation $\bar{v}$ would finally either buy the product from the offline store or keep it when purchasing online. Note that consumers have *ex ante* identical uncertainty of the valuation v, and they can realize it only after they see the product either ordering online or shopping in an offline store. We denote the prior belief that the probability of higher valuation $\bar{v}$ as $\delta \in (0, 1)$.

A consumer needs to make two sequential decisions: (1) shopping channel decision, that is, to choose online or offline; and (2) purchase decision, that is, keep or return it if shopping online, and pay for it or leave without purchase if shopping offline. If returning a product, the consumer gets a refund of $r \in [0,\ p_o]$, where $r = p_o$ for "free return" and $r = 0$ for "no return".

A consumer chooses a shopping channel based on her expected utilities from online and offline, which are given by the following two equations:

$$EU_o = \bar{v} \cdot \delta + r \cdot (1 - \delta) - p_o \tag{A-1}$$

$$EU_s = (\bar{v} - p_s) \cdot \delta - t \tag{A-2}$$

The consumer goes online if and only if $EU_o \geqslant EU_s$, which gives us $t \geqslant \hat{t} = p_o - p_s\delta - r(1 - \delta)$. Figure A-2 summarizes consumers' decision.

Figure A-2: Consumers' Decision Map

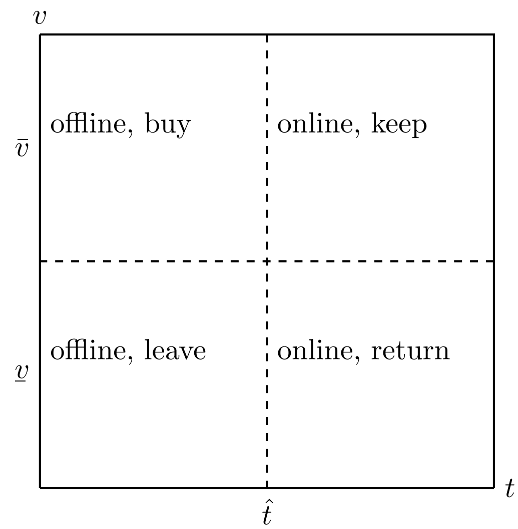

With consumers' shopping channel decision, two retailers' expected profits are given by:

$$\pi_o = (p_o - c_0)(1 - \hat{t})\delta + (p_o - r - c_r)(1 - \hat{t})(1 - \delta) \tag{A-3}$$
$$= (1 - \hat{t})[p_o - c_0\delta - (r + c_r)(1 - \delta)],$$
$$\pi_s = (p_s - c_0)\hat{t}\delta. \tag{A-4}$$

Suppose there are two consumer types, $\theta = H, L$, for high and low types. Two types differ in the probability of high valuation for the product, denoted as δ_θ and the fraction of high type is λ.

We assume firms cannot tell the consumer type, so that they have to set one uniform price and return policy. Each type of consumers choose shopping channel based on the travel cost $\hat{t}_\theta = p_o - p_s\delta_\theta - r(1 - \delta_\theta)$.

Solving for the equilibrium of profit-maximization, we obtain:

$$
\begin{aligned}
p_o^* &= c_0 + (\frac{W}{6Z} - \frac{1}{6})c_r + \frac{W}{6Z} + \frac{1}{2}, \\
r^* &= c_0 + (\frac{W}{6Z} - \frac{2}{3})c_r + \frac{W}{6Z}, \\
p_s^* &= c_0 + (\frac{W}{3Z} - \frac{1}{3})c_r + \frac{W}{3Z}.
\end{aligned}
$$

where $W \equiv \lambda(\delta_H - \delta_L) + \delta_L$, $Z \equiv \lambda(\delta_H^2 - \delta_L^2) + \delta_L^2$, and $W^2 - Z \neq 0$ for the above equilibrium to exist.[1]

We can observe that $p_o^* > r^*$ given any values of parameters, which suggests that "free return" is not optimal for the online seller when it cannot tell the consumer type. This result is consistent with our main analysis.

F Preference for "Free Returns"

In the basic model, we maintain an assumption that the only impact of return policy is on expected utility. However, this may not the case in reality. For instance, Drake (2014) pointed out that "according to Internet analytic service comScore, consumers check out an e-commerce business' policies with an eye on not having to pay to return an item. Eighty-two percent of online shoppers polled in 2014 said they would make a purchase on the Internet if the company offered free return shipping."[2]

Therefore, it is plausible to assume that "free return" as a salient signal may have another impact. One possibility is that there are some consumers would purchase online only when "free return" is offered. We add this assumption in this section to examine the optimal return policies.

Again we models a unit mass of consumers with *ex post* "valuation" v that is drawn randomly from $[0,\ V]$, by a distribution function $G(v)$. However, we add one assumption that some consumers would purchase only when "free return" is offered, and others normally

[1] Otherwise, there is only one type consumers, and the equilibrium still exist but not unique.

[2] available at https://www.entrepreneur.com/article/238701.

make purchase decision based on expected utility. The fraction of those normal consumers are λ.

A monopoly retailer would consider two return policies: "partial return" and "free return". In the following analysis, we assume a uniform distribution $G(v) = v/V$.

Partial return

When partial return is employed, only normal consumers would purchase, thus the optimal price and refund would be the same as what we have in section 3.3 with a lower profit:

$$
\begin{aligned}
p^* &= \frac{V}{2} + \frac{(c_0 - c_r)^2}{2V}, \\
r^* &= c_0 - c_r, \\
\pi^{PR} &= \lambda[\frac{V}{2} - c_0 + \frac{(c_0 - c_r)^2}{2V}].
\end{aligned}
$$

Free return

When free return is employed, all consumers would purchase and keep the product if $v \geqslant p$, thus the retailer faces the following profit function:

$$
\pi = (p - c_0)[1 - G(p)] - c_r G(p)
$$

The optimal price to maximize profit is:

$$
\begin{aligned}
p^* &= \frac{c_0 - c_r + V}{2}, \\
\pi^{FR} &= \lambda[\frac{(c_0 - c_r + V)^2}{4V} - c_0].
\end{aligned}
$$

Comparison

It is easy to show that $\pi^{FR} \geqslant \pi^{PR}$ when c_r and λ are small enough. Figure A-3 reports a numerical analysis when assuming $V = 10$ and $c_0 = 1$.

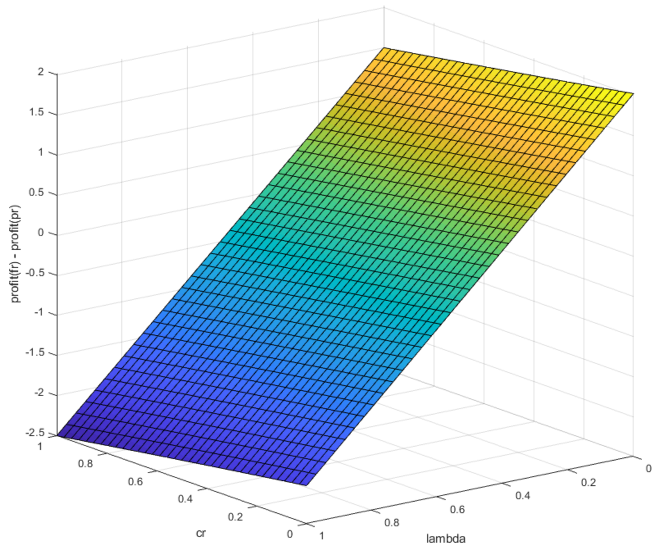

Figure A-3: Profits with Preference for Free Return

G Proof of Proposition 9

Proof. The retailer's optimization program is

$$\max_{p,r} \pi = p - c_0 + (c_0 - c_r - r)G(r-h),$$
$$\text{subject to } E[U(p,r;v)] = V - \int_{r-h}^{V} G(v)dv - p \geqslant 0,$$

which is equivalent to

$$\max_{r} \pi = V - \int_{r-h}^{V} G(v)dv - c_0 + (c_0 - c_r - r)G(r-h),$$
$$\text{with } p = V - \int_{r-h}^{V} G(v)dv.$$

The derivative of profit with respect to r is

$$\frac{\partial \pi}{\partial r} = (c_0 - c_r - r)g(r-h)$$

Since $g(r-h) > 0$ for any $r \in [h, V+h]$, then if $h \leqslant c_0 - c_r \leqslant h + V$, we can obtain $\tilde{r} = c_0 - c_r$; if $c_0 - c_r \leqslant h$, we have $\frac{\partial \pi}{\partial r} < 0$, which gives us a corner solution $\tilde{r} = h$; and if $c_0 - c_r \geqslant V + h$, we have $\frac{\partial \pi}{\partial r} > 0$ which gives us another corner solution $\tilde{r} = V + h$. In addition, Lemma 6 shows that $\tilde{r} = h$ is equivalent to $\tilde{r} = 0$. Thus, we have proved all the cases. $\square$

Qichao Shi

Department of Economics
Lehigh University
621 Taylor Street
Bethlehem, PA 18015, U.S.

Email: qis214@lehigh.edu
Phone: (+1) 857-919-4837
Web: qichaoshi.github.io

EDUCATION

Ph.D. Economics, Lehigh University, U.S. (expected) August 2019

Dissertation: Three Essays on Market Strategy with Incomplete Information
Committee: James Dearden (chair), Ernest Lai, Daisy Dai, David Goldbaum

M.S. Economics, Suffolk University, U.S. September 2013

Bachelor of Economics, Wuhan University, China June 2010

Bachelor of Medicine, Huazhong University of Science and Technology, China June 2010

RESEARCH FIELDS

Primary: Industrial Organization, Experimental Economics

Secondary: Applied Econometrics, Labor Economics

PUBLICATIONS

- "Impact of Economic Globalization on Income Inequality: Developed Economies vs Emerging Economies," with Inmee Baek. *Global Economy Journal*, 2016, 16.1: 49-61

WORKING PAPERS

- "An Experimental Study of Selling Expert Advice," (Job Market Paper), with James Dearden and Ernest Lai
- "Third-party Product Rankings and Advertising," with James Dearden, David Goldbaum and Ernest Lai
- "Sources of Income Inequality from a Global Perspective: Evidence of Country Heterogeneity"

RESEARCH IN PROGRESS

- "Behavioral Return Policies with Two-sided Asymmetric Information," with James Dearden, David Goldbaum and Ernest Lai

RESEARCH EXPERIENCE

Research Assistant, Department of Economics, Lehigh University, Fall 2018

Research Assistant, Prof. James Dearden, Lehigh University, Summer 2017, 2018

Research Assistant, Prof. Thomas Hyclak, Lehigh University, Summer 2016

TEACHING EXPERIENCE

Instructor, Lehigh University

- Principles of Economics, Summer 2018

Teaching Assistant, Lehigh University

- Money, Banking and Financial Markets, Spring 2017 – Spring 2018
- Principles of Economics, Fall 2014 – Fall 2016

Course Assistant, Lehigh University

- Principles of Economics, Summer 2015, 2016, 2017

HONORS & AWARDS

Doctoral Travel Grants, Lehigh University, June 2018 ($1000, competitive)

Citation for the Academic Achievement, Suffolk University, May 2013

Award for Academic Achievement, Huazhong University of Science and Technology, 2006

SKILLS

Software: Stata, Python, Matlab, Mathematica, EViews, oTree, LaTex

Languages: English (Fluent), Chinese (Native)

CONFERENCE PRESENTATIONS

"An Experimental Study of Selling Expert Advice"

- China Meeting of the Econometric Society (June 2018)
- Chinese Economists Society North America Conference (April 2018)

"Competition between Warehouse Clubs and Supermarkets: A Basic Model"

- Eastern Economic Association Conference (February 2017)

"Impact of Economic Globalization on Domestic Income Inequality"

- Eastern Economic Association Conference (February 2014)

Service Activities

Professional Societies and Affiliations

- American Economic Association, The Econometric Society, Eastern Economic Association, The Chinese Economist Society

University Service

- Graduate Student Senate Representative, Lehigh University, Fall 2017

References

James Dearden

Professor

Department of Economics
Lehigh University
621 Taylor Street RBC 320
Bethlehem, PA 18015, U.S.
Email: jad8@lehigh.edu

Ernest Lai

Associate Professor

Department of Economics
Lehigh University
621 Taylor Street RBC 426
Bethlehem, PA 18015, U.S.
Email: kwl409@lehigh.edu

Daisy Dai

Assistant Professor

Department of Economics
Lehigh University
621 Taylor Street RBC 402
Bethlehem, PA 18015, U.S.
Email: wed215@lehigh.edu

David Goldbaum

Associate Professor

Economics Discipline Group
University of Technology, Sydney
PO Box 123
Broadway, NSW 2007, Australia
Email: David.Goldbaum@uts.edu.au

Irina Panovska (Teaching)

Assistant Professor

Department of Economics
Lehigh University
621 Taylor Street RBC 457
Bethlehem, PA 18015, U.S.
Email: irp213@lehigh.edu