HARVARD | BUSINESS | SCHOOL

DOCTORAL PROGRAMS |

DISSERTATION ACCEPTANCE CERTIFICATE

(To be placed in First Library Copy)

The undersigned, approved by the Chair of the Doctoral Programs, have examined a dissertation entitled

Essays on the Design and Industrial Organization of Online Markets

presented by **Pavel Kireyev**

candidate for the degree of Doctor of Business Administration and hereby certify that it is worthy of acceptance.

Signature ....

Elie Ofek, Chair

Signature ....

Sunil Gupta

Signature ....

Ariel Pakes

Date 4/11/2016

SOLDIERS FIELD | BOSTON, MA 02163 | Ph 617.495.6101 | Fx 617.496.2203 | doctoralprograms@hbs.edu | GEORGE F. BAKER FOUNDATION

Essays on the Design and Industrial Organization of Online Markets

A dissertation presented

by

Pavel Kireyev

to

The Harvard Business School

in partial fulfillment of the requirements

for the degree of

Doctor of Business Administration

in the subject of

Marketing

Harvard University

Cambridge, Massachusetts

April 2016

Dissertation Advisor:
Professor Elie Ofek

Author:
Pavel Kireyev

Essays on the Design and Industrial Organization of Online Markets

Abstract

The internet has revolutionized marketing. Firms use the internet to procure advertising content, reach consumers, and offer a convenient channel of purchase. Given the growing importance of the internet, marketers must learn to take advantage of new marketplaces and channels. This research examines how the economic design of electronic marketplaces and online channels affects consumer and firm behavior. The first chapter examines the effects of prize structure and entry limits on participant behavior and idea quality in a freelance marketplace where popular advertisers such as P&G and Unilever organize contests to procure ideas for advertising content. It presents a structural model and uses counterfactual simulations to show that although the number of prizes does not appear to affect contest outcomes, prize amount and submission limits may have a significant impact that depends on participant heterogeneity and information. The second chapter uses a structural model and counterfactual simulations to explore how different expiration and pricing policies in a marketplace that offers deeply discounted but expiring deals for products affects the purchase and redemption behavior of consumers and the pricing decisions of merchants. This chapter sheds light on recent marketing regulation that befell the daily deals industry. The third and final chapter studies the decisions of retailers who operate both an online and a store channel to match their own prices across channels in a variety of competitive settings. It uses an analytical game theory model to show that different self-matching configurations may emerge in equilibrium, and that self-matching pricing policies may increase retailer profits.

Contents

List of Tables

List of Figures

Acknowledgments

I owe a debt of gratitude to my advisors, Elie Ofek, Sunil Gupta, and Ariel Pakes, for their enduring support, instruction, and insight. Each contributed greatly to my development as a scholar, helped direct my interests towards relevant topics and methodological innovations in marketing and economics, and made my experience as a doctoral student truly memorable.

The research presented herein benefited greatly from the efforts of co-authors, Vineet Kumar, Xueming Luo, and Koen Pauwels, and colleagues, Joseph Davin, Tilman Dette, Michael Egesdal, Zhenyu Lai, Clarence Lee, Daniel Pollmann, Thomas Wollmann, and Lingling Zhang. I thank Doug Chung, John Gourville, Ayelet Israeli, Uma Karmarkar, Karim Lakhani, Robin Lee, Donald Ngwe, and Thales Teixeira for their invaluable advice. I thank John Korn, LuAnn Langan, Jen Mucciarone, and Marais Young of the doctoral office for their incredible support. Parts of this dissertation benefited from excellent conversations with participants of the marketing unit seminar at the business school and the industrial organization seminar at the economics department.

I am grateful to my parents, grandparents, and friends for their enduring support throughout my studies.

Introduction

This research explores marketing procurement and practice in the online economy. The emphasis is on how the economic design of electronic marketplaces and online channels affects consumer and firm behavior.

The first chapter investigates the impact of contest design parameters on participation and quality outcomes in a popular freelance marketplace for idea generation. In marketing, design, and other creative industries, firms use freelance marketplaces to organize contests and obtain high-quality ideas for ads, new products, and even business strategies from participants. The central question faced by contest sponsors is how to appropriately structure prizes and entry regulations. I develop an empirical model of idea generation (ideation) contests and investigate the impact of the number of prizes, prize amount, and submission limit on participation and quality outcomes using data from a popular marketing ideation platform. The model explains participant submission decisions, jury ratings, and sponsor rankings of winning submissions. Counterfactuals reveal the impact of design parameters and provide guidance for the optimal design of ideation contests and platforms.

The second chapter, co-authored with Vineet Kumar and Xueming Luo, examines the impact of a daily deal marketplace's expiration and pricing policy on consumer and merchant behavior. This chapter sheds light on recent regulation that required the extension of expiration dates on daily deals. The daily deals marketplace has emerged as a new online platform for merchants to connect with local consumers by offering them coupons for deeply discounted products. Consumers who purchase a daily deal coupon have a limited time to exercise the option to consume a product. We develop an empirical model of the marketplace, where merchants set discounts and consumers react by making purchase and redemption decisions. We employ a single-agent dynamic model of consumer purchase and redemption decisions that allows for inattention - a consumer may fail to redeem a purchased deal not by choice but because of forgetfulness. Merchants set discounts taking consumer demand into account. We estimate the model using data from an Asian daily deal marketplace and conduct counterfactual simulations to compare consumer choice, merchant discounts, and platform revenue under alternative expiration and pricing policies.

The third and final chapter, co-authored with Vineet Kumar and Elie Ofek, investigates the implications of a novel multichannel pricing strategy. Multichannel retailing has created several new strategic choices for retailers. With respect to pricing, an important decision is whether to offer a "self-matching policy." Self-matching allows a multichannel retailer to offer the lowest of its online and store prices to consumers. In practice, we observe considerable heterogeneity in self-matching policies: there are retailers that offer to self-match and retailers that explicitly state they will not match prices across channels. Using a game-theoretic model, we investigate the strategic forces behind the adoption (or non-adoption) of self-matching across a range of competitive scenarios, including a monopolist, two competing multichannel retailers, as well as a mixed duopoly. Though self-matching can negatively impact a retailer when consumers pay the lower price, we uncover two novel mechanisms that can make self-matching profitable in a duopoly setting. Specifically, self-matching can dampen competition online and enable price discrimination in-store, and its effectiveness in these respects depends on the decision-making stage of consumers and the heterogeneity of their preference for the online vs. store channels. Surprisingly, self-matching strategies can also be profitable when consumers use "smart" devices to discover online prices in stores. Our findings provide insights for managers on how and when self-matching can be an effective pricing strategy.

Chapter 1

Markets for Ideas: Prize Structure and Entry Limits in Ideation Contests

1.1 Introduction

Contests have a rich history as a mechanism for the procurement of innovation in design and technology. With the growth of the internet, firms have begun using contests to procure ideas for advertising, new products, and marketing strategies. For example, when motorcycle manufacturer Harley-Davidson split with its ad agency of 31 years, it turned to the crowd to create its next generation of advertising (Klaassen, 2011). With the help of a crowdsourcing firm, Harley organized an *ideation contest* - fans of the brand could submit short ad ideas for a chance to win a cash prize. The winning submissions motivated a series of popular Harley marketing campaigns. Contests carry many advantages over the traditional ad agency model of advertising procurement: brands can expect a large number of ideas at a relatively low cost; participants tend to be actual end users of the product; and contests build awareness by engaging consumers in conversation with the brand (Kirby, 2013).

Harley is not alone in adopting the contest model of ideation. Government agencies and firms in the private sector across a variety of industries have implemented ideation contests. For example, Challenge.gov, a government operated ideation platform, solicits ideas from participants for projects organized by different federal agencies, such as DARPA and NASA. Innocentive, a popular platform for scientific innovation, hosts ideation contests for companies such as Ford, GlaxoSmithKline, and MasterCard. The crowdsourcing studio Tongal organizes advertising ideation contests for AT&T, General Electric, Google, Lego, P&G, and Unilever, among others.[1]

The success of an ideation contest hinges on its design - the choice of how to structure prizes and contest entry regulations. In this research, I empirically examine the impact of three broadly applicable design decisions - how many prizes to award, how much money to award per prize, and how many submissions to accept per participant - on contest participation and idea quality outcomes such as expected total and maximum submission quality.

Prior research has explored how incentives affect ideation from an agency theory perspective (Toubia, 2006; Girotra *et al.*, 2010), but few papers have empirically examined the use of contest mechanisms for the procurement of ideas. I develop and estimate a structural model of ideation contests to assess the impact of different design parameters on contest outcomes. The model captures participant, jury, and sponsor decision processes. Participants choose how many ideas to submit to a contest based on their expected returns and costs of effort. A Tongal jury assigns a quality rating to all submissions. The sponsor then ranks submissions and rewards the winners.

Participants may differ in their *abilities* and *costs*. Ability heterogeneity reflects the notion that idea quality may differ across participants. Ideation contests attract a wide array of entrants with different backgrounds and experiences. Certain participants may submit higher quality ideas than others. For example, we may expect a Harley veteran to generate higher quality ideas for a motorcycle ad than someone with limited riding experience. Cost heterogeneity allows for participants to differ in how easy or difficult it is for them to think of ideas for a particular contest. For example, individuals with more outside commitments may have less time to participate in online contests, increasing their costs of making submissions. I allow for abilities and costs to differ by participant and contest. Moreover, participants can select into contests based on an unobservable (to the researcher) component of costs.

I use data from crowdsourcing platform Tongal to estimate the model in three stages. First, data on sponsor rankings of winning submissions identify sponsor preferences as a function of ob-

[1] Some of the earliest ideation contests in marketing date back to the 1950s and 1960s (Kirby, 2013). Popular brands would organize contests through newspapers and specialized publications to obtain ideas for ads, commercial jingles, and new product names from consumers.

servable participant characteristics and a rating assigned to the submission by a jury. Second, jury ratings assigned to all submissions identify the distribution of ratings conditional on observable and unobservable participant characteristics. Third, participant submission decisions identify the costs of ideation. I estimate the final stage as an empirical discrete game where participants choose how many ideas to submit to a given contest to maximize their expected payoffs. I use moment inequalities to partially identify parameters of the cost function. This methodology allows for multiple equilibria, a non-parametric cost unobservable, and yields estimates that are robust to different specifications of participant information sets. I estimate the model separately by industry of the contest sponsor.

Counterfactual simulations reveal the impact of alternative prize allocation and submission limit decisions on contest outcomes under different assumptions about the information sets of participants. I experiment with two information structures, which I label complete and incomplete information. In the complete information scenario, participants know their own characteristics, as well as sponsor and jury preferences, and the characteristics of their competitors. In the incomplete information scenario, participants do not know sponsor or jury preferences, or competitor characteristics, but are aware of the joint density of these variables conditional on contest structure. I find that both information structures imply similar counterfactual outcomes on average across contests. However, the outcome of each individual contest may differ depending on the informational assumption.

First, I investigate the impact of offering a single prize instead of multiple prizes. I find that although multiple prizes motivate weaker (low ability, high cost) participants and demotivate stronger (high ability, low cost) participants, the number of prizes, holding fixed total award, has a negligible impact on participation and quality - the change in expected marginal returns to most participants is small compared to submission costs. Second, I explore the impact of increasing prize money. I find that a strong response from stronger participants leads to an increase in idea quality but may not lead to a substantial increase in the total number of entrants. Finally, I examine the effect of reducing the maximum number of submissions allowed per participant. This policy benefits weaker participants who would have otherwise been discouraged from entry by the presence of stronger participants who submit multiple times to the same contest. A more stringent submission limit restricts stronger participants, increasing the number of entrants but reducing expected quality outcomes.

The remainder of this paper is organized as follows. Section 1.2 summarizes the relevant theoretical and empirical literature on contest design. Section 1.3 presents the data and Section 1.4 presents descriptive evidence of the importance of prize allocation. Section 1.5 outlines the structural model, Section 1.6 details the estimation routine, and Section 1.7 presents the estimates. Section 1.8 examines the impact of counterfactual contest designs and presents practical implications. Section 1.9 concludes.

1.2 Contest Design

A contest is a game in which players invest costly effort in an attempt to win a prize. Throughout, I refer to players who consider entering a contest as *participants*. Of all participants, those who enter the contest are referred to as *entrants*, and the rest, as *non-entrants*. The *sponsor* organizes the contest and ultimately selects winners and awards prizes. Effort in the contest literature is typically viewed as a non-negative continuous decision variable. I view effort as the discrete number of idea *submissions* a participant makes to a given contest.[2]

Traditionally, contests have been modeled as either imperfectly discriminating (Tullock, 1980), all-pay auctions (Baye *et al.*, 1994), or rank-order tournaments (Lazear and Rosen, 1981). Imperfectly discriminating contests and rank-order tournaments typically allow for uncertain outcomes - the participant exerting the highest effort is not guaranteed to win. However, a higher effort increases the participant's chances of winning. In all-pay auctions, highest effort typically guarantees victory. Ideation contests share similarities with imperfectly discriminating contests and rank-order tournaments - participants who submit the most ideas are not guaranteed to win, and in contests with multiple prizes, submissions are ranked in order of the sponsor's preferences.

A key aspect of ideation contests is participant heterogeneity. Participants, with different levels of skill and experience, can freely join the platform and enter contests. Although a greater number of entrants improves the sponsor's chances of obtaining an extreme-value, high quality idea, especially in contests with significant participant uncertainty about sponsor preferences

[2] The ideation contests I study require participants to submit 140 character ideas for ads. Each participant can submit at most 5 ideas to a single contest. Section 1.4 show evidence that the number of submissions is a good measure of participant effort - submissions react in expected ways to changes in prize allocation.

(Boudreau *et al.*, 2011), increased participant asymmetries typically result in reduced effort (Baye *et al.*, 1993; Stein, 2002). Intuitively, participants with a low chance of winning are discouraged and "give up," which in turn reduces the level of competition for participants with a high chance of winning, resulting in a lower level of effort from all types. However, an appropriate choice of prize allocation can mitigate this concern.[3]

Theory literature has examined the impact of prize allocation on the effort of heterogeneous participants. Moldovanu and Sela (2001) explore the impact of multiple prizes on effort in all-pay auctions. The authors show that, holding fixed total award, a greater number of prizes encourages weaker participants, as they have a chance of winning one of the lower ranking prizes. On the other hand, stronger participants exert less effort, as with multiple prizes, the payoff from "losing" increases. The optimality of offering multiple prizes depends on participant heterogeneity and the convexity of their costs of effort. If costs are sufficiently convex, a smaller number of prizes will not encourage stronger participants to increase effort by enough to compensate for the reduced effort of weaker participants, and the sponsor may find it optimal to offer multiple prizes. Szymanski and Valletti (2005) argue that stronger participants may increase effort in response to multiple prizes in imperfectly discriminating contests. The added uncertainty of winning may motivate stronger participants to react to increasing competition from weaker participants. Few papers have examined the impact of prize allocation on outcomes other than effort. Terwiesch and Xu (2008) consider expected maximum and average quality outcomes in imperfectly discriminating innovation contests and all-pay auctions. The authors similarly show that multiple prizes may be optimal in contests with heterogeneous participants, but a single prize works best for contests with ex-ante identical participants.[4] Overall, the effect of multiple prizes on effort is ambiguous and depends on participant heterogeneity and cost function shape.[5] I contribute to the literature

[3]Fullerton and McAfee (1999) argue that restricting entry can also benefit sponsors. By imposing an appropriate entry auction mechanism that encourages higher ability participants to enter, the sponsor can expect greater effort while minimizing the costs of procurement.

[4]See Sisak (2009) for a survey of the theoretical literature on multiple prizes in contests.

[5]Apart from heterogeneity, participant risk-aversion may also motivate sponsors to adopt multiple prizes. Kalra and Shi (2001) show that in rank-order sales contests with sufficiently risk-averse homogeneous participants, multiple prizes may increase effort. However, experimental research suggests that risk-averse participants are less likely to enter contests altogether (Eriksson *et al.*, 2009; Dohmen and Falk, 2011). Throughout, I focus on settings with risk-neutral heterogeneous participants and show further evidence in support of this model tenet in Section 1.4.

by presenting estimates of different prize allocation policies on participation and quality outcomes, and suggesting practical implications for ideation contest design.

Although the question of how many submissions to accept per participant is unique to contests where participants can make multiple submissions, researchers have investigated the related aspect of restricted bidding in all-pay auctions. Che and Gale (1998) consider the impact of caps on investments in political lobbying in an all-pay auction with one high-valuation (strong) player and one low-valuation (weak) player. The authors find that bid caps can increase total spending by limiting the strong participant and encouraging the weak participant. Che and Gale (2003) similarly show that handicapping a stronger participant in research contests can improve the contest outcome. I investigate the impact of restricting the number of submissions per participant - a relevant and easy to implement policy in the context of ideation contests. My results show that submission limits constrain higher ability participants and increase overall entry. However, I find that a more stringent submission limit may reduce expected total and maximum idea quality.

Substantial progress in the empirical literature on contests has been achieved with the increasing availability of online data. Boudreau *et al.* (2016) examine the impact of competition on the effort of heterogeneous participants in the context of the popular TopCoder platform for programming contests. The authors examine a number of contest design policies but do not focus on the question of how many prizes to award or how many submissions to accept per participant. Yoganarasimhan (2016) presents a model of beauty contest auctions, or procurement auctions with uncertain outcomes. The author applies the model to a freelance marketplace where each auction can have at most one winner and selection on unobserved components of cost is less of a concern as participants do not invest costly effort to enter. In the setting of online design contests, research has explored the impact of feedback and entry visibility on participation and submission quality (Wooten and Ulrich, 2015a,b; Gross, 2016) as well as the effects of competition on experimentation (Gross, 2014). Gross (2016) presents a structural model to study the impact of performance feedback on submission quality in logo design contests but does not study prize allocation and submission limits or allow for non-entry in estimation. I contribute to the empirical literature by presenting estimates of the impact of a number of key design parameters, such as the number of prizes, prize amount, and submission limit, on participation and quality outcomes in ideation contests organized by popular brands. I suggest an empirical model of contests with multiple prizes, participant heterogeneity, and the possibility of selection into contests based on costs of effort. Furthermore, I address a recent call in literature to allow for more flexibility in the information structures of empirical games (Borkovsky *et al.*, 2015) and derive contest outcome predictions that are robust to different informational assumptions.

1.3 Data and Setting

I use data from Tongal, a popular crowdsourcing platform. Major brands such as AT&T, General Electric, Google, Lego, P&G, and Unilever use the platform to organize ideation contests. Brands typically use the obtained ideas to develop advertising content, either independently or through Tongal.

Ideation contests on Tongal operate as follows. Tongal and the contest sponsor jointly decide on how many prizes to offer and how much money to offer per prize. The sponsor presents participants with the contest prize allocation, rules and regulations, and a description of the ideation topic. Participants can then enter the contest by submitting at least one 140 character idea for an ad based on the topic suggested by the sponsor. Each entrant can submit at most 5 ideas to a single contest. After the contest ends, a Tongal jury reviews and rates each submission without knowledge of the identity of its creator. Winning submissions are selected and ranked by the sponsor and their creators receive prize money. The platform does not display the identities or actions of participants during the contest period. Only after the sponsor selects winners does the platform make public the list of winning submissions.

I focus on a sample of 181 ideation contests that ran from 2011 to 2015 (the platform was founded in 2009) and a set of 8,875 participants who entered at least one of these contests. A total of 127 sponsors organized at least 1 and at most 11 of the contests, with 24 sponsors hosting more than one contest. For each contest, I observe the number of submissions made by each entrant, the rating assigned to each submission, the ranking of the winning submissions, the number of prizes awarded, and prize amount. All contests divide prizes evenly among winners. For example, each winning submission receives $250 if a contest offers 4 prizes with a total award of $1,000. I classify each contest into a category based on the industry of the sponsor. Table 1.1 further describes the classification criteria and shows the distribution of contests by category.

Table 1.1: *Contest Categories*

Category	Description	Number of Contests
Consumer	Consumer packaged goods	22
Food	Food and beverages: snacks, ingredients, soft and alcoholic drinks	45
Utility	Hardware: tires, tools, paint, etc.	12
Health	General and male personal care and medical products	21
Health(F)	Female personal care products	18
Tech	Electronics and internet services	19
Toy	Toys and games	20
Other	Sporting goods, clothing, social cause, professional services	24

An important aspect of many contests is that not all participants who consider entering choose to do so. I use browsing data to define the set of likely non-entrants, or participants who considered entering a contest but chose not to. Specifically, participants who did not enter the contest but viewed the contest page more than once and were active in the past 3 months are considered likely non-entrants. I restrict non-entrants to this subset to avoid including participants who were simply “surfing" the site without seriously considering entry into the contest. This procedure yields a total of 9,732 instances of non-entry by likely non-entrants. On 35,011 occasions, participants make at least one submission. Figure 1.1 presents a plot of the distribution of submissions per participant within a contest for all 181 contests in the data. A significant proportion of participants does not enter, submits once, or makes the maximum number of submissions allowed. For each one of the participants, I observe a set of characteristics collected by the platform, which is further summarized in Section 1.6.

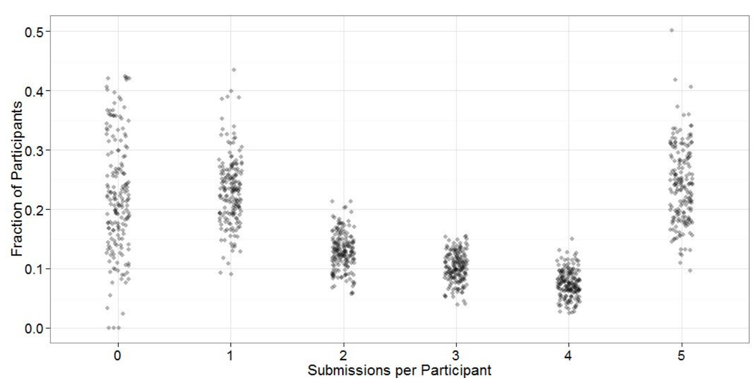

Figure 1.1: *Distribution of Submissions Within Contests*

Note: An observation is a contest. Plot shows the fraction of participants who made d submissions within each contest, where $d \in \{0, ..., 5\}$.

Table 1.2 presents summary statistics for the contests considered. The contests tend to attract a high number of entrants and submissions, with the average contest securing 193 entrants and 572 submissions. There is also substantial variation in prize allocation across contests, with the number of prizes ranging from 1 to 50 and prize amount per winning spot ranging from $100 to $1,250. Figure 1.2 shows that contests predominantly offer a $250 or $500 prize per spot and that contests with more prizes tend to have a higher total award.

Table 1.2: *Summary of Contest Characteristics*

Per-Contest Characteristics	Min	Median	Mean	Max
Non-Entrants	0	48	54	124
Entrants	58	187	193	499
Submissions	178	551	572	1,875
Number of Prizes	1	4	5	50
Prize Amount per Spot	$100	$250	$323	$1,250
Total Award	$500	$1,000	$1,450	$10,000

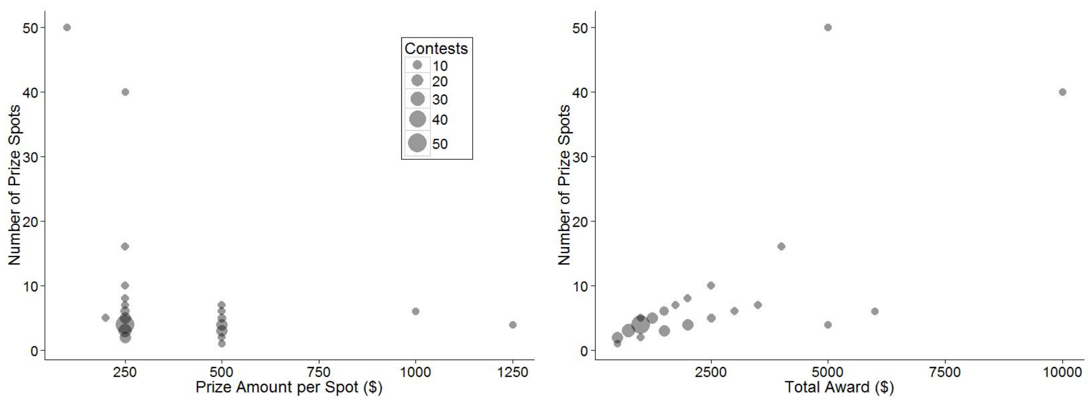

Figure 1.2: *Scatter of Contest Characteristics*

Approximately 0.9% or 950 out of a total of 103,554 submissions win a prize. I observe the rating assigned to each submission by a Tongal jury. Ratings are assigned on a 5-point scale and based on the jury's perceived quality of the submission. Submissions receiving below a 3 are considered inadequate by the jury but may still win if sponsor preferences differ significantly from the jury's. I refer to a rating below 3 as a low rating. Otherwise, the submission is said to have received a high rating. Of all submissions, 68% receive a low rating. A submission with a low rating has a 0.1% chance of winning and less than 10% of all winning submissions have a low rating. A high-rating submission has a 2.4% chance of winning and roughly 90% of all winning submissions have a high rating. Jury ratings are a strong predictor of a submission's success.

1.4 Descriptive Evidence

Is there evidence in the data that participants respond to prizes? Such evidence would suggest that prize allocation is an important design parameter that can alter behavior.

First, consider the impact of prize amount on submissions. Figure 1.3 shows the raw correlation between total award and three outcomes: the total number of submissions a contest receives, the number of entrants, and the number of submissions made by each entrant. Contests that award a higher prize attract more entrants and receive more submissions in total and more submissions per entrant. I regress the outcome metrics on the logarithm of total award and include fixed effects to control for differences in contest category, sponsor, and the number of prizes. Identifying variation comes from differences in the outcome across contests that share the same set of fixed effects but offer different prizes. Table 1.3 shows the estimated coefficients on the logarithm of total award. All columns show positive coefficients, consistent with the notion that a larger total award increases entry and effort.

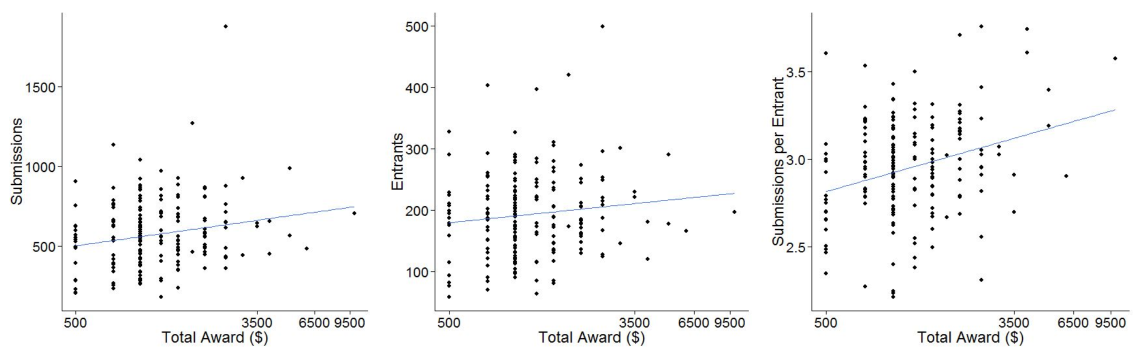

Figure 1.3: *Scatter of Submission Outcomes and Total Award*

Note: An observation is a contest. Blue line shows best-fitting linear model.

Next, consider the impact of the number of prizes on submission behavior. Table 1.4 shows the coefficient estimates for a series of regressions of submission outcomes on the logarithm of the number of prizes offered in a contest, controlling for total award as well as category and sponsor fixed effects. I find a negative relationship between the outcome metrics and the number of prizes, suggesting that participants are possibly not sufficiently heterogeneous or risk-averse for multiple prizes to be optimal. Alternatively, contests that award more prizes may be more difficult, even after controlling for category or sponsor.

To further investigate the impact of the number of prizes on submission behavior, I draw on

Table 1.3: *Contest-Level Regressions of Outcomes on* $\log(Total\_Award_t)$

DV: $\log(\text{Submissions}_t)$	0.154	0.187	0.350	0.260	0.238	0.555
	(0.053)	(0.025)	(0.058)	(0.061)	(0.069)	(0.112)
R^2	0.045	0.071	0.319	0.050	0.055	0.377
DV: $\text{Submissions}_t/\text{Entrants}_t$	0.156	0.166	0.129	0.147	0.097	0.060
	(0.039)	(0.051)	(0.074)	(0.095)	(0.065)	(0.075)
R^2	0.082	0.090	0.071	0.032	0.015	0.012
Category Fixed Effects	N	Y	N	N	Y	N
Sponsor Fixed Effects	N	N	Y	N	N	Y
Number of Prizes Fixed Effects	N	N	N	Y	Y	Y
Observations	181	181	78	178	167	40

Note: Robust standard errors in parentheses.

Table 1.4: *Contest-Level Regressions of Outcomes on* $\log(Number\_of\_Prizes_t)$

DV: $\log(\text{Submissions}_t)$	-0.296	-0.306	-0.175
	(0.096)	(0.116)	(0.118)
R^2	0.034	0.032	0.018
DV: $\text{Submissions}_t/\text{Entrants}_t$	-0.085	-0.068	-0.078
	(0.089)	(0.099)	(0.061)
R^2	0.006	0.002	0.004
Category Fixed Effects	N	Y	N
Sponsor Fixed Effects	N	N	Y
Total Award Fixed Effects	Y	Y	Y
Observations	179	158	33

Note: Robust standard errors in parentheses.

individual participant-level submission patterns. Theory (Moldovanu and Sela, 2001; Terwiesch and Xu, 2008) predicts that stronger participants prefer a smaller number of prizes, holding fixed total award, whereas the reverse is true for weaker participants. I classify participants into segments based on their participation frequency, defined as the number of contests they viewed. I expect that participants who view a large number of contests either have a low cost of participation or a high expected probability of winning. Each segment contains a similar number of participant decision instances. Contests are grouped based on their observable characteristics. I compare the number of submissions made by the same participant across contests offering a different number of prizes within the same contest group. The regression equation is given by

$$\text{Submissions}_{it} = \alpha \log(\text{Number\_of\_Prizes}_t) + \xi_{iG(t)} + \epsilon_{it}, \tag{1.1}$$

where Submissions_{it} is the number of submissions made by participant i in contest t. The fixed effects $\xi_{iG(t)}$ control for unobserved participant and observed contest heterogeneity, where $G(t)$ denotes the group of contest t. Finally, α is the parameter of interest and ϵ_{it} is an error term.

Figure 1.4 illustrates the estimate of the coefficient α when Regression 1.1 is applied separately to each segment of participants. Participants who view a small number of contests appear to prefer multiple prizes, but participants who view a moderate number of contests show a distaste for multiple prizes. No effect is found for the most frequent participants. Further inspection reveals that participants who consider over 33 contests tend to submit near the maximum number of times to each contest, perhaps because of low costs or high abilities. As a result, there is limited variation in their submission behavior, resulting in a near-zero coefficient for the most frequent participants. Figure 1.4 shows evidence consistent with the theoretical prediction that stronger participants may prefer fewer prize, holding fixed total award, but is not consistent with explanations that rely solely on risk-aversion or unobserved contest difficulty level.

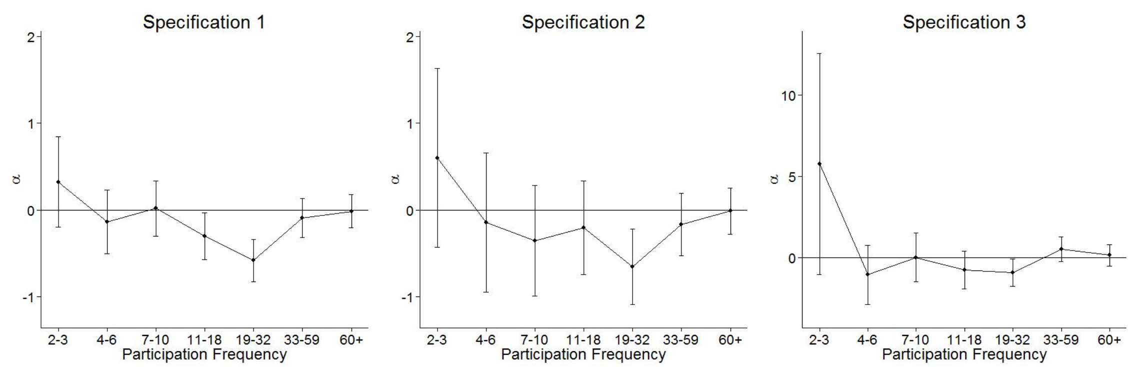

Figure 1.4: *Participant-Level Regression Estimates by Participation Frequency*

Note: Plots show estimates and robust 95% confidence intervals for α in Regression 1.1. Participants segmented based on the number of contests they viewed. Specification 1 groups contests by total award. Specification 2 groups contests by category and total award and category. Specification 3 groups contests by sponsor and total award.

The descriptive evidence presented in this section suggests that submission decisions respond to changes in prize allocation. Furthermore, the evidence supports models with risk-neutral heterogeneous participants such as Moldovanu and Sela (2001), Stein (2002), and Terwiesch and Xu (2008). I proceed to derive a structural model motivated by these theoretical contributions. A structural model of sponsor preferences is required to assess the impact of contest design on quality, a variable not observed in the data. Furthermore, a model of participant submission decisions would enable an analysis of the impact of submission limits, a variable that remains unchanged across contests in the data. Finally, structure would allow for the investigation of the impact of participant information sets on contest outcomes.

1.5 Model

I model each ideation contest as an independent game consisting of three stages. First, participants decide on how many submissions to make given their costs and expected payoffs. I consider a trial-and-error model of ideation, whereby participants sample from a quality distribution with each submission in an attempt to generate an idea of high quality for the sponsor. This approach is common in models based on the statistical view of innovation in new product design (Dahan and Mendelson, 2001; Loch *et al.*, 2001) and in models of research contests with uncertain outcomes (Taylor, 1995; Fullerton and McAfee, 1999; Terwiesch and Xu, 2008). Second, a Tongal jury assigns a binary quality rating to all submissions. Finally, the sponsor reviews all submissions and ranks the top submissions by its perception of submission quality. Most sponsors in the data offer only one contest. Sponsors who offer multiple contests tend to focus on different products and ideation topics across contests. I use sponsor and contest interchangeably and denote both by t.

The model will capture three key features of ideation contests. First, participant abilities may differ across contests. If two participants make the same number of submissions, the participant with the higher ability has a larger expected payoff. Ability heterogeneity captures the notion that the “fit" between a participant and a contest may depend on the participant’s background and the contest ideation topic. Second, participants exhibit cost heterogeneity and may select into contests they find most convenient based on an unobservable component of costs. If two participants have the same ability but different costs, the participant with the lower cost may increase her expected payoff by making a larger number of submissions. Cost heterogeneity accounts for the possibility that certain participants may be busier than others at different times, or find it more difficult to think of ideas for certain contests. Third, participants may view their own abilities, the number of competitors, competitor abilities, and competitor actions with uncertainty and form expectations of their own expected payoffs given their information sets.

I work backwards and first present the model for the final stage sponsor decision (Section 1.5.1), followed by the model for jury ratings (Section 1.5.2). I present the model for participant entry decisions in Section 1.5.3. The empirical implementation of the two stages is presented in Section 1.6. A discussion of key model assumptions is presented in Section 1.6.5.

1.5.1 Sponsor Choice Model

Consider the sponsor's decision process after it receives a set of submissions. From the perspective of the sponsor, submission s by participant i in contest t has quality

$$q_{st} = \beta X_i + \gamma W_{st} + \epsilon_{st},$$

where X_i is a vector of participant characteristics, W_{st} is the rating assigned to submission s, β and γ are sponsor preference parameters, which I assume are common to all sponsors within a category, and $\epsilon_{st} \sim T1EV$ is an iid submission-specific quality shock.

The interaction of sponsor preferences β and participant characteristics X_i reflects the differences in participant submission quality that can be explained by observed participant characteristics. The parameter γ captures the effect of the rating assigned to the submission. The rating W_{st} may explain unobserved components of submission quality that are not captured by X_i. The shock ϵ_{st} captures all heterogeneity in submission quality that cannot be explained by participant characteristics or submission rating.

The sponsor observes q_{st} for each s and ranks submissions by quality. Only the best N_t submissions receive a ranking, where N_t is at least as large as the number of prizes. In other words, the sponsor chooses a ranking $s_{(1)}, ..., s_{(N_t)}$ such that $q_{s_{(1)}t} \geq q_{s_{(2)}t} \geq ... \geq q_{s_{(N_t)}t} \geq q_{kt}$, where q_{kt} is the quality of any other submission k not in $s_{(1)}, ..., s_{(N_t)}$.

1.5.2 Jury Rating Model

Before the sponsor reviews the submissions and selects winners, a Tongal jury gives a rating to each submission in contest t. From the perspective of the jury, submission s has quality

$$u_{st} = \alpha X_i + \xi_{it} + \eta_{st},$$

where X_i is the same vector of participant characteristics used in the sponsor choice model, α is a parameter that reflects jury preferences and is assumed constant within a category, $\xi_{it} \sim N(\phi, \sigma)$ is a participant-contest specific quality unobservable, distributed iid across participants and contests, and η_{st} is an iid submission-specific quality shock that follows a standard logistic distribution. If $u_{st} > 0$, the jury assigns a high rating and $W_{st} = 1$. Otherwise, the submission receives a low rating and $W_{st} = 0$.

The unobservable ξ_{it} allows for correlation in the unobserved components of quality of submissions made by the same participant in contest t. For example, a participant may submit ideas with a similar level of humor that cannot be explained by her X_i. This source of variation in ratings will be explained by her quality unobservable ξ_{it}.

1.5.3 Participant Entry Model

Risk-neutral participants form expectations of their contest payoffs with respect of the distribution of jury ratings conditional on X_i, the distribution of quality shocks, and participant perceptions of competitor actions and characteristics. Participants know their own X_i and the contest prize structure but may view sponsor and jury preferences, the number of competitors, competitor characteristics, and competitor actions as random variables because of incomplete information. I further make the following assumption:

Assumption 1 *Participants do not know the realizations but do know the distributions of ϵ_{st}, η_{st}, and ξ_{it} before making submission decisions.*

I require that participants cannot select into contests based on an unobservable (to the researcher) component of sponsor preferences or jury rating. In other words, participants have the same information as the researcher regarding unobserved components of submission quality.[6]

Suppose that a total of I_t participants consider entering contest t. Participant i chooses to make $d_{it} \in \{0, 1, ..., D\}$ submissions in contest t, where D is the submission limit. Expected payoffs are given by

$$\pi_{it} = E\left[R_t(d_{it}, d_{-it}; X_i, X_{-it}) | \mathcal{J}_{it}\right] - c_{it}(d_{it}).$$

[6] A similar assumption is made in empirical models of contests by Yoganarasimhan (2016) and Gross (2016), and in two-stage entry and demand models in industrial organization such as Ishii (2008), Eizenberg (2014), and Wollman (2014). In Section 1.6.5 and Appendix A.1, I show evidence that supports this assumption.

The expected returns function $R_t(d_{it}, d_{-it}; X_i, X_{-it})$ captures the expected winnings of a participant with characteristics X_i who makes d_{it} submissions given competitor characteristics X_{-it} and actions d_{-it}. For example, in a contest with one prize, the expected winnings of a participant making $d_{it} > 0$ submissions are

$$R_t(d_{it}, d_{-it}; X_i, X_{-it}) = \int \frac{\sum_{k=1}^{d_{it}} \exp\{\beta X_i + \gamma W_{kt}^i\}}{\sum_{j=1}^{I_t'} \left(\sum_{k=1}^{d_{jt}} \exp\{\beta X_j + \gamma W_{kt}^j\} \right)} dF_{W_t}(W_t^1, ..., W_t^{I_t'})$$

where I_t' is the total number of participants who made at least one submission, W_{kt}^i is the rating assigned to submission number k belonging to participant i in contest t, F_{W_t} is the distribution of ratings and $W_t^i = (W_{1t}^i, ..., W_{d_{it}t}^i)$. To determine her optimal action, the participant must form an expectation of her expected returns $R_t(.)$ with respect to her information set $\mathcal{J}_{it}$, which will vary depending on what the participant knows about her competitors.

I consider cost functions of the form

$$c_{it}(d_{it}) = (\theta_1 + \theta_2 d_{it} + \nu_{it}) d_{it},$$

where ν_{it} is a mean-zero participant-contest specific cost unobservable, and θ_1 and θ_2 are cost parameters with $\theta = (\theta_1, \theta_2)$. Prior to entry, each participant observes her cost shock ν_{it} and chooses how many submissions to make to maximize expected payoffs π_{it}. She may also choose to make no submissions and obtain zero payoffs.

1.6 Estimation

Estimation proceeds in two stages. In the first stage, I estimate the sponsor choice model and the jury rating model. Given the first stage results, I estimate the participant entry model using moment inequalities. The underlying game is likely to have multiple equilibria because of the discrete action space. The moment inequalities methodology allows for multiple equilibria, does not require explicit specification of participant information sets, and permits a flexible distribution of cost unobservables. In the second stage, I follow the estimation procedure for discrete games with ordered choices suggested by Ishii (2008) and Pakes *et al.* (2015).

1.6.1 Sponsor Choice Model

I use data on sponsor ranking decisions, participant characteristics, and submission ratings to estimate the sponsor choice model. Identification relies on rankings data and heterogeneity in participant characteristics and submission ratings. Variation in sponsor decisions given different sets of submission characteristics identifies the sponsor preference parameters. The likelihood of observing a ranking $s_{(1)}, ..., s_{(N_t)}$ is

$$L_t\left(s_{(1)}, ..., s_{(N_t)}\right) = \prod_{r=1}^{N_t} \left(\frac{\exp\{\beta X_{s_{(r)}} + \gamma W_{s_{(r)}t}\}}{\sum_{j=r}^{N_t} \exp\{\beta X_{s_{(j)}} + \gamma W_{s_{(j)}t}\} + \sum_{k \in \emptyset} \exp\{\beta X_k + \gamma W_{kt}\}} \right),$$

where N_t is the number of ranked submissions, $\emptyset$ is the set of all unranked submissions, and $X_s = X_i$ if submission s belongs to participant i. The likelihood of the data corresponds to the likelihood of a rank-ordered logit model. Prior research has used rank-ordered logit models (also known as exploded logits) to recover preferences from rankings in consumer survey data (Beggs *et al.*, 1981; Chapman and Staelin, 1982). In my setting, a structural model of sponsor choice generates a statistical rank-ordered logit model that can be estimated using data on sponsor rankings of contest winners. I estimate the model separately for each category using maximum likelihood methods.

1.6.2 Administrator Rating Model

Data on jury ratings and variation in participant characteristics within a contest identify jury preference parameters α and ϕ. The standard deviation of participant-contest specific quality unobservables σ is identified from instances where multiple submissions made by the same participant receive a similar rating that cannot be explained by the participant's observed characteristics. The likelihood of observing a sequence of ratings $W^i_{1t}, ..., W^i_{d_{it}t}$ for participant i conditional on ξ_{it} is

$$M_t(W^i_{1t}, ..., W^i_{d_{it}t}|\xi_{it}) = \prod_{k=1}^{d_{it}} \left(\frac{\exp\{\alpha X_i + \xi_{it}\}}{1 + \exp\{\alpha X_i + \xi_{it}\}} \right)^{W^i_{kt}} \left(\frac{1}{1 + \exp\{\alpha X_i + \xi_{it}\}} \right)^{1-W^i_{kt}}.$$

The likelihood of observing all of the ratings in a contest is

$$\int \prod_{i=1}^{I'_t} M_t(W^i_{1t}, ..., W^i_{d_{it}t}|\xi_{it}) dF_{\xi}$$

where F_ξ is the distribution of ξ_{it}, parameterized by ϕ and σ. I use simulated maximum likelihood to estimate model parameters separately for each category.

1.6.3 Participant Entry Model

I use moment inequalities to partially identify cost parameters for each contest category. Pakes *et al.* (2015) show how moment inequalities can be used to obtain upper and lower bounds on cost parameters for discrete choice games where agents make ordered choices. With moment inequalities, I need not explicitly specify an equilibrium selection mechanism. Furthermore, the methodology allows for a flexible distribution of cost unobservables and yields estimates that are robust to different specifications of participant information sets. However, parameters will typically be set identified and not point identified. In other words, moment inequalities yield a set of parameters as opposed to a point, and confidence bounds must be obtained taking this into account.

The participant entry model can be rewritten as follows. I define the expectational error $\omega_{itd_{it}}$ as the difference between a participant's expected and actual returns:

$$\omega_{itd_{it}} = E\left[R_t(d_{it}, d_{-it}; X_i, X_{-it}) | \mathcal{J}_{it}\right] - R_t(d_{it}, d_{-it}; X_i, X_{-it}).$$

Sources of expectational error may include participant uncertainty about competitor actions (as a function of costs) and characteristics, and may also incorporate optimization mistakes made by the participant in evaluating her expected returns. Then, the payoff equation can be written as

$$\pi_{it} = R_t(d_{it}, d_{-it}; X_i, X_{-it}) - c_{it}(d_{it}) + \omega_{itd_{it}}.$$

I require that participants are correct on average and, at this stage, place no additional restrictions on the distribution of expectational errors.

Assumption 2 $E[\omega_{itd_{it}}] = 0.$

Note that Assumption 2 holds trivially if participants have correct expectations as

$$\begin{aligned} E[\omega_{itd_{it}}] &= E\left[E\left[R_t(d_{it}, d_{-it}; X_i, X_{-it}) | \mathcal{J}_{it}\right]\right] - E\left[R_t(d_{it}, d_{-it}; X_i, X_{-it})\right] \\ &= E\left[R_t(d_{it}, d_{-it}; X_i, X_{-it})\right] - E\left[R_t(d_{it}, d_{-it}; X_i, X_{-it})\right] = 0. \end{aligned}$$

However, participants may have incorrect expectations (perhaps because of incorrect perceptions about equilibrium action distributions) as long as they are correct on average.

I proceed by first deriving a lower bound for marginal costs, where I take into account the possibility that participants who made the maximum number of submissions may have had particularly low costs. Then, I derive an upper bound for marginal costs, where I use a selection correction technique to account for the possibility that non-entrants may have had particularly large costs. Additional assumptions about the distributions of $\omega_{itd_{it}}$ and ν_{it} are introduced as they become relevant.

Lower Bound

First, consider the derivation of the lower bound for marginal costs. Define a function of the difference in observable returns from making one additional submission as

$$\Delta R^*_{it}(d_{it}+1, d_{it}) = \begin{cases} R_t(d_{it}+1, d_{-it}, X_i, X_{-it}) - R_t(d_{it}, d_{-it}, X_i, X_{-it}), \text{ if } d_{it} < 5, \\ 0, \text{ if } d_{it} = 5, \end{cases}$$

and let $\omega_{itd_{it}+1,d_{it}} = \omega_{itd_{it}+1} - \omega_{itd_{it}}$. By revealed preference, for a participant who made less than 5 submissions,

$$\underbrace{\Delta R^*_{it}(d_{it}+1, d_{it}) + \omega_{itd_{it}+1,d_{it}}}_{\text{expected marginal return}} \leq \underbrace{\theta_1 + \theta_2(2d_{it}+1) + \nu_{it}}_{\text{marginal cost}},$$

as the expected marginal return from making one additional submission must be no greater than the marginal cost of making one additional submission. Otherwise, the participant would have made $d_{it}+1$ instead of d_{it} submissions. For a participant who made 5 submissions, the expected marginal return from making one additional submission is likely an overestimate of the marginal cost of doing so, as the participant may have chosen to make more submissions under a more lenient submission limit. I make the assumption that the marginal cost of making one additional submission is at least zero for entrants who made the maximum permitted number of submissions.

Assumption 3 *The condition $\theta_1 + \theta_2(2d_{it}+1) + \nu_{it} \geq 0$ holds for entrants with $d_{it} = 5$.*

Taking the expectation over participants, it must be the case that

$$E\left[\underbrace{\theta_1+\theta_2(2d_{it}+1)}_{\text{marginal cost}}-\underbrace{\Delta R^*_{it}(d_{it}+1,d_{it})}_{\text{marginal return}}\right]\geq 0.$$

The expectational errors $\omega_{itd_{it}+1,d_{it}}$ average out to zero because participants are correct on average. The cost unobservables ν_{it} average out to zero because the expectation does not condition on the participant's action. The ability to take an expectation over cost unobservables for all participants, regardless of their action, is crucial for the estimation of bounds on cost parameters.

An empirical analogue for the lower bound for marginal costs can be written as

$$m^L(\theta)=-\frac{1}{T}\sum_{t=1}^{T}\frac{1}{I_t}\sum_{i=1}^{I_t}\Delta r^*_{it}(d_{it}+1,d_{it};\theta),$$

where T is the total number of contests used in estimation and

$$\Delta r^*_{it}(d_{it}+1,d_{it};\theta)=\Delta R^*_{it}(d_{it}+1,d_{it})-\theta_1-\theta_2(2d_{it}+1).$$

Any θ that satisfies $m^L(\theta)\geq 0$ must lie in the identified set of cost parameters.

In practice, $R_t(d_{it},d_{-it};X_i,X_{-it})$ is not analytically tractable but is required as an input to $\Delta r^*_{it}(d_{it}+1,d_{it})$ in the definition of $m^L(\theta)$. To obtain expected returns, it is necessary to consider the probability of observing all possible combinations of winning submissions from the set of all submissions. For contests with multiple prizes and hundreds of submissions, this expression can be analytically intractable. I use simulation to obtain an approximation of the expected returns function for each participant in every contest.[7]

Upper Bound

Next, consider the upper bound for marginal costs. For entrants i in $L_t=\{i: d_{it}>0\}$, define the difference in observable returns from making one less submission as

$$\Delta R_{it}(d_{it},d_{it}-1)=R_t(d_{it},d_{-it},X_i,X_{-it})-R_t(d_{it}-1,d_{-it},X_i,X_{-it}).$$

[7] It can be shown that simulation error averages out in the moment inequalities framework.

Then, by revealed preference, for $i \in L_t$,

$$\underbrace{\Delta R_{it}(d_{it}, d_{it}-1) + \omega_{itd_{it},d_{it}-1}}_{\text{expected marginal return}} \geq \underbrace{\theta_1 + \theta_2(2d_{it}-1) + \nu_{it}}_{\text{marginal cost}}.$$

In other words, the expected marginal return of increasing submissions from $d_{it}-1$ to d_{it} must have been greater than the associated marginal cost. Otherwise, entrants would have made one less submission than they actually did.

The above condition holds only for participants who submitted at least once. I must take into account the possibility that non-entrants, or participants with $d_{it} = 0$, may have had particularly large cost unobservables. If an empirical analogue, only for entrants, is developed based on the above inequality, the estimated upper bound on costs may be too low. Pakes *et al.* (2015) suggest using symmetry of the ν_{it} distribution to obtain an upper bound on the ν_{it} for non-entrants. Intuitively, the negative of the lowest lower bound for ν_{it} can be used as the highest upper bound for the negative of the ν_{it} of non-entrants. This result holds as long as the ν_{it} density is not skewed left.

Assumption 4 *For each contest, the cost unobservables ν_{it} follow a mean-zero distribution that is not skewed left.*

For exposition, I derive all subsequent inequalities assuming that the ν_{it} follow a symmetric distribution, which will yield conservative bounds if the actual distribution is skewed right. Assumption 4 allows for the cost unobservables to correlate with participant characteristics but requires that contests do not differ in difficulty level, conditional on contest category. The symmetry property of the ν_{it} distribution can be used to implement the selection correction technique suggested by Pakes *et al.* (2015) and obtain upper bounds for the unobserved costs of non-entrants. As long as the number of entrants exceeds the number of non-entrants for a given contest, the negatives of the lowest lower bounds on cost unobservables over all participants can be used as upper bounds for the negatives of the cost unobservables of non-entrants.

For a given contest, the moment conditions can be developed as follows. First, rank all entrants by $r_{it} = -\Delta r^*_{it}(d_{it}+1, d_{it}; \theta)$ so that $r_{(1)t} \leq r_{(2)t} \leq ... \leq r_{(I_t)t}$. Next, construct a set of size equal to the number of non-entrants, such that $U_t = \{i : r_{it} \geq r_{(n_t+1)t}\}$, where n_t is the number of entrants in contest t. The negative lowest lower bounds for ν_{it} become the upper bounds for the $-\nu_{it}$ of non-entrants. Define the moment

$$m^U(\theta) = \frac{1}{T}\sum_{t=1}^{T}\frac{1}{I_t}\left(\sum_{i\in L_t}\Delta r_{it}(d_{it}, d_{it}-1;\theta) - \sum_{i\in U_t}\Delta r^*_{it}(d_{it}+1, d_{it};\theta)\right),$$

where

$$\Delta r_{it}(d_{it}, d_{it}-1;\theta) = \Delta R_{it}(d_{it}, d_{it}-1) - \theta_1 - \theta_2(2d_{it}-1)$$

is the difference in observable profits from making one less submission.

Consider the expectational error $\omega_{itd_{it}}$. The lowest lower bounds on cost unobservables used as part of the selection correction technique originate from a selected subset of participants. I require an assumption on the joint density of expectational errors and cost unobservables to ensure that participants with the lowest costs do not consistently underestimate their expected marginal returns. Otherwise, the upper bounds I obtain for non-entrants may be too low. This assumption would only affect the observations used in constructing $m^U(\theta)$ for participants in U_t with $d_{it} < 5$ because the inequality condition for participants with $d_{it} = 5$ does not contain an expectational error term (Assumption 3). I find that this applies to less than 5% of all participant entry occasions and, as a result, does not have a consequential impact on estimated identified set of cost parameters. I provide the exact condition for the joint density of expectational errors and cost unobservables in Appendix A.2. The proof that if $m^U(\theta) \geq 0$, then θ lies in the identified set of cost parameters follows naturally from the proof presented in Pakes *et al.* (2015) and is reproduced in Appendix A.2 for completeness.

Identified Set

The identified set for parameters $\theta = (\theta_1, \theta_2)$ is defined as

$$\{\theta : m^L(\theta) \geq 0 \text{ and } m^U(\theta) \geq 0\}.$$

Identification of the cost parameters follows naturally from the restrictions imposed by the moment inequalities. However, it is not possible to obtain lower and upper bounds for both θ_1 and θ_2 (a total of 4 bounds) using only 2 moment inequalities. Additional restrictions on the covariance of ν_{it} and participant characteristics X_i can generate additional inequalities. However, there is no reason to expect that characteristics that affect the quality of a participant's submissions do not also affect her costs. Instead, I choose to restrict the shape of the cost function.

Assumption 5 $\theta_1 = 0$ *so that* $c_{it}(d_{it}) = (\theta_2 d_{it} + \nu_{it}) d_{it}$.

I find that a cost function with $\theta_1 > 0$ and $\theta_2 \leq 0$ is unlikely. Given the large number of participants in each contest, the marginal expected returns of each participant are almost linear in the number of submissions. If the cost function were also linear or concave, a small change in prize amount would lead all participants to submit either 0 or 5 times, which does not appear reasonable as over one-third of all participants in the data make an intermediate number of submissions.

1.6.4 Confidence Bounds

I obtain confidence bounds using a block-bootstrap procedure. The procedure is applied separately to each category. I sample a dataset of size equal to the number of contests within a category (with replacement) and estimate the sponsor choice and jury rating models on the re-sampled set of contests. I repeat this procedure multiple times and recover the standard deviation of the parameter estimates across bootstrapped datasets.

The confidence set for the cost parameter includes the true parameter 95% of the time and is obtained using a procedure suggested by Andrews and Soares (2010). Intuitively, the procedure consists of simulating via bootstrap the distribution of a criterion function that penalizes violations of the moment inequalities. The simulated distribution is used to obtain a critical value, which is compared to the actual value of the criterion function in the observed sample. Points where the value of the criterion function falls below the critical value are included in the confidence set. The above procedure, first described by Chernozhukov *et al.* (2007), may produce very conservative confidence sets, primarily because of the influence of very positive moments that satisfy the inequality restrictions by a wide margin. Andrews and Soares (2010) suggest a moment selection procedure that yields more precise coverage by excluding very positive moments before simulating the criterion function. I use the bootstrapped datasets obtained in the inference procedure for the sponsor choice and jury rating models to incorporate first-stage estimation error in the Andrews and Soares (2010) criterion function.

1.6.5 Discussion of Model Assumptions

Selection and Participant Heterogeneity

The model allows for rich sources of observed and unobserved participant heterogeneity. Before entering a contest, participants differ in their observed characteristics X_i and cost unobservables ν_{it}, and can choose how many submissions to make based on these variables. Hence, the model allows for selection on observable components of ability and unobservable components of costs. Furthermore, participants can exhibit persistent differences in ability through X_i and persistent differences in costs as the ν_{it} may be correlated across contests for the same participant.

The model does not allow for persistent unobserved heterogeneity in the quality of a participant's submissions. Participants cannot choose how many submissions to make based on an unobserved component of expected submission quality. I include an indicator in X_i for whether or not a participant won money from Tongal prior to her first ideation contest to allow for persistent differences in skill and submission quality across participants. Participant submissions may depend on a quality unobservable ξ_{it} which is iid across participants and contests and not known to participants before entry. The quality unobservable may explain correlation in the quality of submissions made by the same participant. A similar assumption on the role of unobservable components of demand is made in recent empirical work on contests (Yoganarasimhan, 2016; Gross, 2016) and two-stage entry models (Ishii, 2008; Eizenberg, 2014; Wollman, 2014) to allow for two-step estimation. Incorporating unobserved components of demand known to participants in entry games with multiple equilibria is an active area of research.

Submission Order

I test for the importance of unobserved components of submission quality by exploiting data on submission order. Participants may choose to submit their best idea first. Evidence of a relationship between submission quality and order may suggest that participants have private information about the quality of their submissions that cannot be explained by participant characteristics alone. In addition, it would provide evidence of decreasing returns to submissions resulting from quality deterioration. I run a series of regressions to test for a relationship between jury rating and submission order. In Appendix A.1, I show that there does not appear to be a strong connection between the two variables.

Cost Function Shape

Incorporating non-linearities in the cost function requires instruments for participant actions, restrictions on the distribution of cost unobservables, or covariance restrictions between observable characteristics and cost unobservables. Note that d_{it} cannot be used as an instrument to construct additional inequalities unless ν_{it} is assumed to be zero as $E[\nu_{it}|d_{it}] \neq 0$. In Section 1.7.3, I present estimates of θ_1 under the assumption that $c_{it}(d_{it}) = (\theta_1 + \nu_{it})d_{it}$ and estimates of θ_2 under the assumption that $c_{it}(d_{it}) = (\theta_2 d_{it} + \nu_{it})d_{it}$. I find that the assumption of constant marginal costs is unlikely to hold as small perturbations in a contest's prize would lead all participants to submit either 0 or 5 times given the approximate linearity of participant marginal expected returns. It is possible that the expected returns function is not sufficiently concave as it does not incorporate risk-aversion or deteriorating submission quality. However, the descriptive evidence in Section 1.4 does not support risk-aversion and the analysis presented in Appendix A.1 does not show significant evidence of submission quality deterioration for participants who make multiple submissions.

1.7 Structural Model Estimates

1.7.1 Sponsor Choice Model

I use a set of characteristics collected by the platform to account for possible sources of heterogeneity in the quality of a participant's submissions. Variables in the set of characteristics X_i include the participant's age, country, gender, an indicator for whether or not the participant won a contest on Tongal prior to the first contest she considered entering in my sample, an indicator for whether or not the participant has production skills, and an indicator for whether or not the participant was referred to the platform. Table 1.5 presents the definitions for all variables used in estimation. The set of observable characteristics is deliberately discretized to ensure that there exists only a finite number of participant types, which facilitates simulation of expected payoffs and counterfactuals.

Table 1.5: *Description of Participant Characteristics*

Variable	Definition	Share of Participants
Demographics		
Age_i	1 participant i was born after 1984 and 0 otherwise.	0.373
Country_i	1 if participant i is from the US and 0 otherwise.	0.814
Gender_i	1 if participant i is female and 0 otherwise.	0.252
Participant-Platform Characteristics		
Paid_i	1 participant i was paid prior to her first contest and 0 otherwise.	0.047
Producer_i	1 if participant i has video production skills and 0 otherwise.	0.232
Referred_i	1 if participant i was referred to the platform and 0 otherwise.	0.210

Table 1.6 presents parameter estimates for the sponsor choice model by category. As expected, I find a significant effect of submission rating on chance of winning. Conditional on rating, estimates of the remaining parameters should be interpreted either as discrepancies between sponsor preferences and jury ratings, or as additional determinants of submission quality that cannot be explained by the rating alone, perhaps because of the coarseness of the rating measure. For example, participants with past success and with experience producing video have a higher chance of winning in many categories, conditional on rating, suggesting that rating alone does not fully capture differences in submission quality for these participants. I find heterogeneity in the effects of age and country but no significant effect of gender. Heterogeneity of coefficients across contests points to differences in the variance of unobserved components of quality and idiosyncrasies in the sponsor's choice process, as well as discrepancies between sponsor choice and jury rating patterns.

I examine the model's explanatory power by testing its ability to predict characteristics of the set of winning submissions. The model successfully predicts 4.4% of winning submissions and can predict the identity of the winning participant correctly 6.1% of the time. In 11.6% of the contests, the model is able to predict at least one winner correctly. Note that these prediction tasks are very difficult as only 0.9% of submissions win a prize.

1.7.2 Jury Rating Model

Table 1.7 presents parameter estimates for the jury rating model by category. I find that participants from the US with past success and video production skills tend to receive higher ratings in all categories. Female participants appear to receive higher ratings in the toy and other categories,

Table 1.6: *Sponsor Choice Model Parameter Estimates by Category*

	Consumer	Food	Hardware	Health	Health(F)	Tech	Toy	Other
Age	-0.703	-0.066	0.170	-0.009	0.001	0.144	0.073	0.294
	(0.211)	(0.158)	(0.290)	(0.207)	(0.205)	(0.229)	(0.178)	(0.173)
Country	-0.442	0.283	0.494	-0.540	0.279	0.391	-0.420	0.062
	(0.233)	(0.274)	(0.525)	(0.252)	(0.338)	(0.378)	(0.240)	(0.255)
Gender	-0.008	-0.204	0.131	0.043	0.169	-0.460	0.115	-0.179
	(0.213)	(0.194)	(0.347)	(0.241)	(0.212)	(0.307)	(0.211)	(0.204)
Paid	0.525	0.274	0.093	-0.181	0.474	0.342	0.655	0.324
	(0.218)	(0.188)	(0.370)	(0.284)	(0.249)	(0.279)	(0.201)	(0.205)
Producer	0.390	0.800	0.977	0.482	-0.028	-0.117	0.369	0.600
	(0.214)	(0.181)	(0.338)	(0.222)	(0.225)	(0.242)	(0.230)	(0.187)
Referred	-0.489	-0.683	-0.037	0.043	-0.538	0.221	-0.067	-0.484
	(0.288)	(0.245)	(0.373)	(0.243)	(0.352)	(0.279)	(0.244)	(0.248)
Rating (γ)	3.214	2.212	3.904	2.767	4.042	4.845	5.155	3.046
	(0.357)	(0.195)	(0.617)	(0.292)	(0.463)	(0.707)	(0.694)	(0.263)
Contests	22	45	12	21	18	19	20	24
Choice Instances	120	172	49	100	101	82	133	148
Log-Likelihood	-657	-998	-267	-561	-508	-412	-636	-810

Note: Bootstrapped standard errors in parentheses.

and a slightly higher rating in the category dedicated to female health and personal care products. Older participants perform better in the female health, technology, and other categories. I also find evidence of unobserved heterogeneity in submission quality across participants as indicated by the estimates of σ, ranging from 1.116 to 1.553. Estimates of the variance of participant-level quality unobservables are higher in the consumer, food, hardware, and toy categories but lower in the health categories.

1.7.3 Cost Estimates

Table 1.8 shows estimates of the cost parameters. First, I assume that $c_{it}(d_{it}) = (\theta_1 + \nu_{it})d_{it}$ and obtain a confidence set for θ_1. Subsequently, I assume that $c_{it}(d_{it}) = (\theta_2 d_{it} + \nu_{it})d_{it}$ and obtain a confidence set for θ_2. The cost parameters are estimated separately for each category. As discussed previously in Section 1.6.3 and Section 1.6.5, I focus on the quadratic cost function in the remainder of the analysis.

The estimated cost parameters suggests that, on average, participants incur a cost of \$0.33-1.30 for producing a single submission.[8] This cost estimate captures the cognitive and mental effort required to think of a 140 character idea as well as the opportunity cost of time that could have

[8]The average cost across categories is obtained by taking the weighted average of category-specific costs.

Table 1.7: *Jury Rating Model Parameter Estimates by Category*

	Consumer	Food	Hardware	Health	Health(F)	Tech	Toy	Other
Age	-0.035	0.007	0.040	-0.021	-0.114	-0.085	0.008	-0.101
	(0.039)	(0.025)	(0.047)	(0.039)	(0.046)	(0.045)	(0.047)	(0.039)
Country	0.292	0.289	0.207	0.224	0.202	0.299	0.213	0.293
	(0.056)	(0.036)	(0.063)	(0.055)	(0.064)	(0.062)	(0.062)	(0.055)
Gender	0.053	0.010	-0.070	0.046	0.088	0.061	0.257	0.125
	(0.044)	(0.029)	(0.055)	(0.046)	(0.049)	(0.053)	(0.055)	(0.044)
Paid	0.367	0.101	0.390	0.169	0.177	0.417	0.225	0.177
	(0.055)	(0.037)	(0.069)	(0.056)	(0.064)	(0.060)	(0.063)	(0.057)
Producer	0.079	0.175	0.219	0.152	0.195	0.149	0.170	0.101
	(0.041)	(0.026)	(0.050)	(0.041)	(0.049)	(0.047)	(0.049)	(0.040)
Referred	-0.082	0.031	0.082	-0.017	-0.145	-0.207	-0.018	-0.044
	(0.051)	(0.031)	(0.058)	(0.047)	(0.061)	(0.055)	(0.057)	(0.049)
Mean (ϕ)	-1.006	-1.098	-1.119	-1.039	-1.078	-1.113	-0.873	-1.191
	(0.059)	(0.039)	(0.066)	(0.059)	(0.069)	(0.066)	(0.067)	(0.058)
Std. Dev. ($\log \sigma$)	0.408	0.350	0.440	0.153	0.110	0.284	0.399	0.220
	(0.034)	(0.022)	(0.040)	(0.038)	(0.049)	(0.041)	(0.041)	(0.035)
Contests	22	45	12	21	18	19	20	24
Choice Instances	4006	10161	2966	4170	2971	3221	2893	4623
Log-Likelihood	-7096	-17819	-5165	-7540	-5370	-5744	-5161	-7654

Note: Bootstrapped standard errors in parentheses.

Table 1.8: *Ideation Cost Estimates by Category*

Cost Function		Consumer	Food	Hardware	Health	Health(F)	Tech	Toy	Other
Linear	LB	1.575	1.492	1.407	2.246	1.767	1.979	3.231	1.951
($\theta_1 \neq 0, \theta_2 = 0$)	UB	3.136	2.436	2.025	3.917	4.597	3.502	6.886	3.197
Quadratic	LB	0.279	0.257	0.242	0.391	0.322	0.352	0.623	0.358
($\theta_1 = 0, \theta_2 \neq 0$)	UB	1.197	0.793	0.664	1.379	2.559	1.443	6.694	1.472

Note: Bootstrapped 95% confidence bounds.

been spent elsewhere. Costs increase in a convex manner, with the average cost of making five submissions in the range of $8.30-32.43. For comparison, the median hourly salary of a writer, copywriter or editor in the US is $28.71 (Bureau of Labor Statistics, 2016), which falls in the range of costs required to think of five original ideas.[9]

I find heterogeneity in costs across categories. For example, contests in the hardware category appear less costly for participants than contests in the toy category. Differences in costs may arise for several reasons: participants may find it easier to think of ideas for certain topics; contests within a category may be scheduled at times that are inconvenient relative to contests in a different category; the set of participants who typically consider entering into a contest within a category

[9] Assume that writers incur a cost of effort less than their wage and that the wage estimate for writers applies to Tongal creatives. Then, an upper bound for the time spent per 140 character idea falls in the range of 3-14 minutes, or 16-76 seconds per word (assuming 11 words per idea).

may differ in their availability from other participants.

1.8 Counterfactuals

Although moment inequalities allow for flexible information sets in estimation, I require an explicit specification of the information sets of participants to simulate counterfactuals. I experiment with two specifications, which I refer to as *complete information* and *incomplete information*.

In the complete information setup, I assume that participants play a Nash equilibrium in submission strategies and know the prize structure of the contest, sponsor and jury preferences, the number of competitors they face, their own characteristics, as well as competitor characteristics and actions. Formally, participant i's information set in contest t is given by $\mathcal{J}_{it}^{CI} = \{d_{it}, d_{-it}, X_i, X_{-it}, I_t, M_t, \delta_t\}$, where M_t represents the prize structure of contest t and includes the prize amount and number of prize spots, $\delta_t = (\alpha_t, \beta_t, \gamma_t, \phi_t, \sigma_t)$, and $\delta_t = \delta_s$ for contests t and s within the same category. I introduce the subscript t on β and γ to reflect the notion that sponsor and jury preferences may differ across categories. For a uniformly sampled point in the identified set, I recover bounds on cost unobservables for each participant. These bounds ensure that at the sampled parameter, the observed decisions constitute an equilibrium. I uniformly sample cost unobservables that satisfy the bounds for each participant and compute equilibrium actions under alternative contest designs. I repeat the procedure for different sample parameters and cost draws, and recover bounds on the outcome of interest across simulations. Details of the counterfactual simulation procedure are provided in Appendix A.3.1.

The complete information assumption may require a high level of participant sophistication. Tongal reveals neither the identities nor the submissions of competitors. Furthermore, participants may not have a good sense of sponsor or jury preferences. In the incomplete information scenario, I focus on a subset of contests with the same prize structure and allow for participant uncertainty with regards to sponsor and jury preferences, and the quantity, characteristics, and actions of competitors. Participant i's information set in contest t is given by $\mathcal{J}_{it}^{II} = \{d_{it}, X_i, M_t\}$, and the participant knows the conditional joint density of the number of participants and sponsor/jury preferences $H(I_t, \delta_t|M_t)$, and the conditional joint density of competitor actions and characteristics $G(d_{-it}, X_{-it}|I_t, M_t, \delta_t)$. I assume that participants use an iterative updating procedure, described further in Section 1.8.1 and Appendix A.3.2, to converge to a new equilibrium from their current state. The procedure can be interpreted as a learning algorithm that participants use to find a new equilibrium under a different contest structure.

Equipped with parameter estimates and an assumption about the information structure of the game, I run simulations of alternative contest designs. I focus on a number of outcome metrics. The total number of entrants $\sum_{i=1}^{I_t} 1\{d_{it} > 0\}$ and total submissions $\sum_{i=1}^{I_t} d_{it}$ are important metrics for the data provider. Increasing entry cultivates participant engagement with the platform and allows for sponsors to communicate with a large number of potential consumers and build brand awareness. Quality outcomes are also important if the goal of the sponsor is to implement the best idea or to incorporate information from all submitted ideas into its marketing strategy. I consider expected total quality, defined as $\int \left(\sum_{i=1}^{I'_t} \sum_{k=1}^{d_{it}} e^{\beta_t X_i + \gamma_t W^i_{kt}}\right) dF_{W_t}$, and expected maximum quality, defined as $\int \log \left(\sum_{i=1}^{I'_t} \sum_{k=1}^{d_{it}} e^{\beta_t X_i + \gamma_t W^i_{kt}}\right) dF_{W_t}$. A sponsor may be interested in total quality if it wishes to combine data from all submissions to create an ad or improve its product offerings. Maximum quality becomes more important for a sponsor interested in implementing only the best idea.

1.8.1 The Impact of Incomplete Information

To simulate counterfactuals under incomplete information, it is necessary to recover $H(I_t, \delta_t | M_t)$ and $G(d_{-it}, X_{-it} | I_t, M_t, \delta_t)$, which can theoretically be achieved by flexible density estimation. However, I find this to be infeasible given the large number of contest-specific variables. Instead, I focus on a subset of 49 contests that offered four \$250 prizes and treat each contest as an independent draw from the joint density of sponsor/jury parameters, the number of competitors, and competitor actions and characteristics conditional on contest structure. All incomplete information counterfactual analyses are conducted only for this subset of contests, labeled $\mathcal{W}$.

To understand the impact of incomplete information on behavior, I recover participant expectational errors, which capture the difference between a participant's expected returns under incomplete information and her expected returns under complete information. To do so, I draw a large sample of contests of size B from $\mathcal{W}$ (with replacement) and label these contests $b = 1, ..., B$. Then, assuming that the cost unobservables ν_{it} are independent conditional on X_i for all participants i within a contest t and letting j_b denote a random participant in contest b, the expected returns $E\left[R_t(d_{it}, d_{-it}; X_i, X_{-it}) | \mathcal{J}_{it}\right]$ can be approximated by

$$ER_{it}(d_{it}) = \frac{1}{B}\sum_{b=1}^{B} R_b(d_{it}, d_{-j_b b}; X_i, X_{-j_b b})$$

for $t \in \mathcal{W}$. This is akin to assuming that the participant knows the variables associated with each contest in $\mathcal{W}$ but does not know which one of these contests she is playing. An estimate of participant i's expectational error is given by $\hat{\omega}_{itd_{it}} = ER_{it}(d_{it}) - R_t(d_{it}, d_{-it}; X_i, X_{-it})$.

I find that participants who make higher quality submissions tend to underestimate their expected returns ($\hat{\omega}_{itd_{it}} < 0$) as they do not know with certainty that they are the most skilled participants in their contests. Similarly, participants who make lower quality submissions tend to overestimate their expected returns ($\hat{\omega}_{itd_{it}} > 0$) as they do not know with certainty that they fall in the lower range within the contests they participate in. Given their observed actions, this implies that participants with higher quality submissions will have lower cost unobservables ν_{it} than had they had complete information. Similarly, participants with lower quality submissions will have higher cost unobservables than in a complete information scenario. A counterfactual simulation analysis is required to compare equilibrium outcomes under complete and incomplete information.

1.8.2 Counterfactual Outcomes Across Contests

I obtain bounds on the outcomes of all contests for each one of the three design counterfactuals under the assumption of complete information. For the 49 contests in $\mathcal{W}$, I also obtain counterfactual outcomes under the assumption of incomplete information using an iterative procedure described in Appendix A.3.2. Figures 1.5, 1.6 and 1.7 show the impact of counterfactual design policies on different contests, and Table 1.9 shows the average impact across contests.

Single Prize

For most of the contests, reducing the number of prizes while holding fixed total award does not have a substantial impact on outcomes. The change in expected marginal returns to participants is low as the number of prizes is usually small compared to the number of submissions. As a

Table 1.9: *Average Counterfactual Design Outcomes Across Contests*

	Entrants		Submissions		Total Quality		Max Quality	
	LB	UB	LB	UB	LB	UB	LB	UB
Complete Information (all contests)								
Single Prize	0.2	3.9	-0.2	1.1	-0.8	2.3	-0.1	0.3
20% Prize Increase	1.1	8.7	2.1	8.9	3.3	12.8	0.4	1.3
3 Submission Limit	0.7	6.9	-9.3	-5.0	-10.1	-3.9	-1.3	-0.5
Incomplete Information (49 contests offering four $250 prizes)								
Single Prize	-0.9	1.8	-0.6	0.7	-0.2	1.8	0	0.2
20% Prize Increase	0.8	4.9	1.7	5.5	6.5	13.5	0.7	1.5
3 Submission Limit	0.5	4.0	-9.5	-6.9	-8.4	-3.7	-1.0	-0.4

Note: Average lower bound (LB) and upper bound (UB) of percentage change in counterfactual outcomes reported.

result, few participants alter their actions.

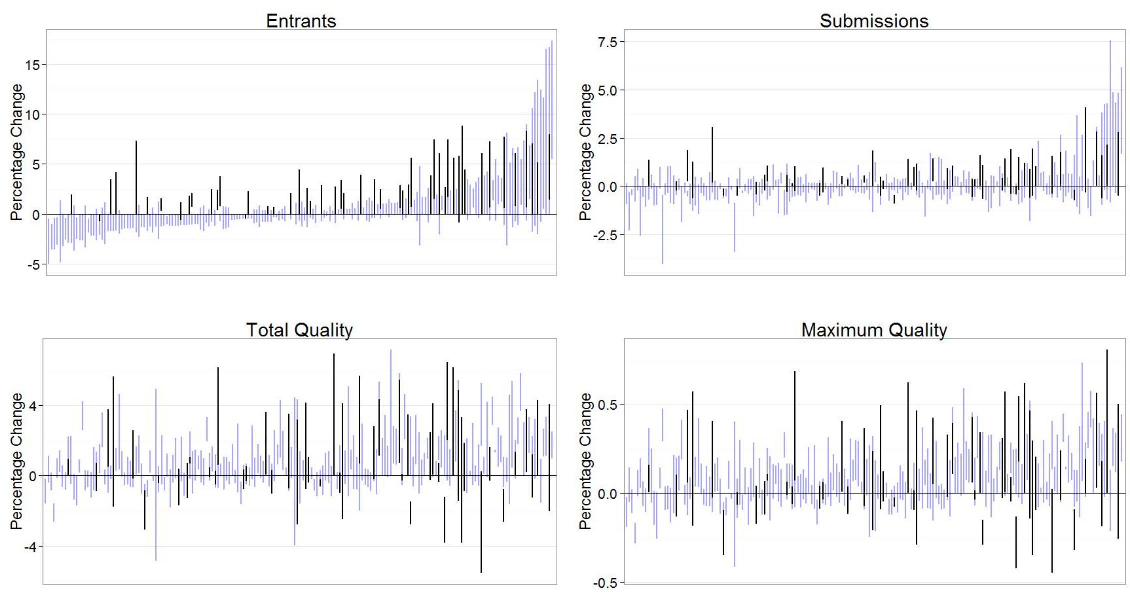

Figure 1.5: *Impact of Reducing the Number of Prizes Across Contests*

Note: Each segment represents the range of counterfactual outcomes for a single contest under complete information (light) and incomplete information (dark). Contests ordered by increasing impact on the number of entrants in the complete information scenario.

For certain contests with very low heterogeneity in the expected marginal returns and costs of participants, a single prize may significantly increase entry but only under the assumption of complete information. In a setting with limited participant heterogeneity, there is no longer a reason to motivate participants with a lower chance of winning, and a single prize reduces the incentive for all participants to achieve a worse rank. However, in a setting with incomplete information, participants do not know the extent of heterogeneity within their contest. As a result, they average over possible states and do not react as strongly to a reduction in the number of prizes in contests with limited participant heterogeneity.

I find that the counterfactual simulation results are in line with the implications of the regression estimates in Table 1.4. In particular, I fail to find a significant effect of prize allocation on submission behavior in both models. This is encouraging validation of the structural model and not a mechanical result as the regressions and the structural model rely on different assumptions and different sources of variation for identification.[10]

Prize Increase

A 20% prize increase improves the outcome metrics, especially expected total quality, but may not lead to as significant an increase in entry in a complete information scenario if there is substantial participant heterogeneity. The added prize incentive predominantly encourages participants with a higher chance of winning to submit more, limiting the benefits of making more submissions for participants with a lower chance of winning. If the contest is highly asymmetric, the prize increase will only significantly affect the behavior of a small number of participants with high expected marginal returns and low costs, which will increase quality outcome metrics but may not lead to a significant increase in entry.

I compare the implied elasticity of submissions with respect to prize to the elasticities obtained from the regression results in Table 1.3. The counterfactual simulations imply a prize elasticity of submissions in the range of 0.105-0.445 for the complete information scenario and 0.085-0.275 for the incomplete information scenario, whereas the estimates in Table 1.3 suggest elasticities of 0.154-0.555. It is encouraging that both sets of elasticities fall within the same range.

Submission Limit

The platform requires that all participants submit at most five times to each contest. What if participants could submit at most three times? Participants with a higher chance of winning would be restricted by a lower submission limit as they tend to make more submissions, creating

[10]Regressions leverage variation in the number of submissions across contests offering different prize structures, whereas the structural model recovers the costs that rationalize participant submission decisions and then simulates a counterfactual outcome for each contest individually. The structural model does not require variation in prize structure for identification.

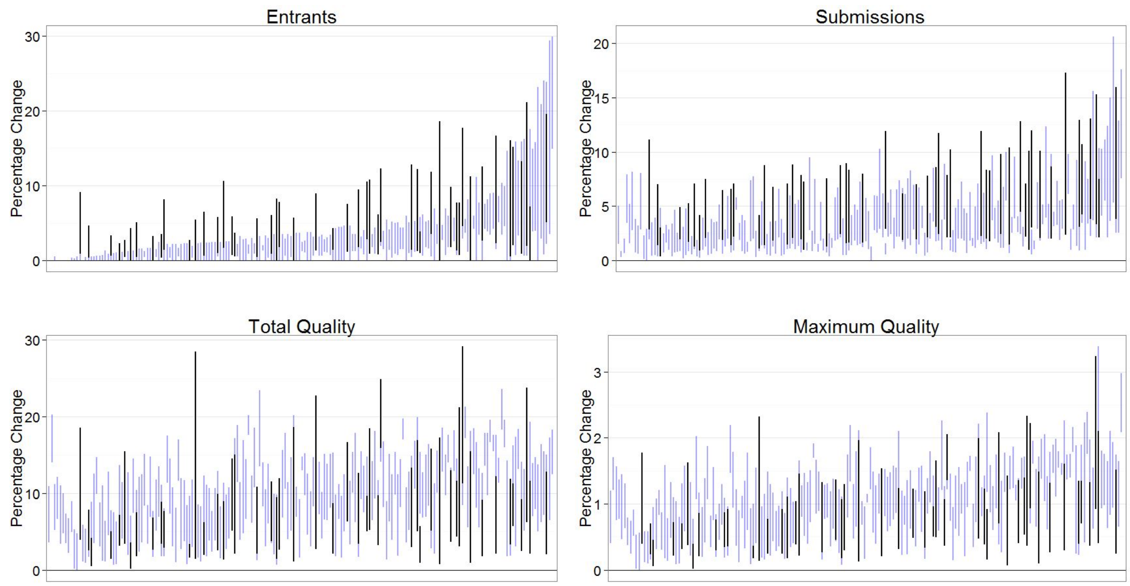

Figure 1.6: *Impact of a 20% Prize Increase Across Contests*

Note: Each segment represents the range of counterfactual outcomes for a single contest under complete information (light) and incomplete information (dark). Contests ordered by increasing impact on the number of entrants in the complete information scenario.

an opportunity for participants with a lower chance of winning to enter the contest and make more submissions.

Higher ability participants with low costs no longer crowd out other potential entrants. A more stringent submission limit encourages entry but restricts participants with higher expected marginal returns and lower costs, reducing expected total and maximum quality across both complete and incomplete information scenarios.

Comparison of Complete and Incomplete Information

In the complete information scenario, the impact of a design parameter depends crucially on the extent of participant heterogeneity in expected marginal returns and costs within a contest. For example, a prize increase will be more effective if participants know that they are competing in a contest with limited participant heterogeneity. Under incomplete information, participants must form an expectation of their expected returns by averaging over states (contests). Building on the previous example, participants who are engaged in a contest with limited participant heterogeneity do not know the extent of this heterogeneity, and hence, will not react as strongly to a prize increase under incomplete information.

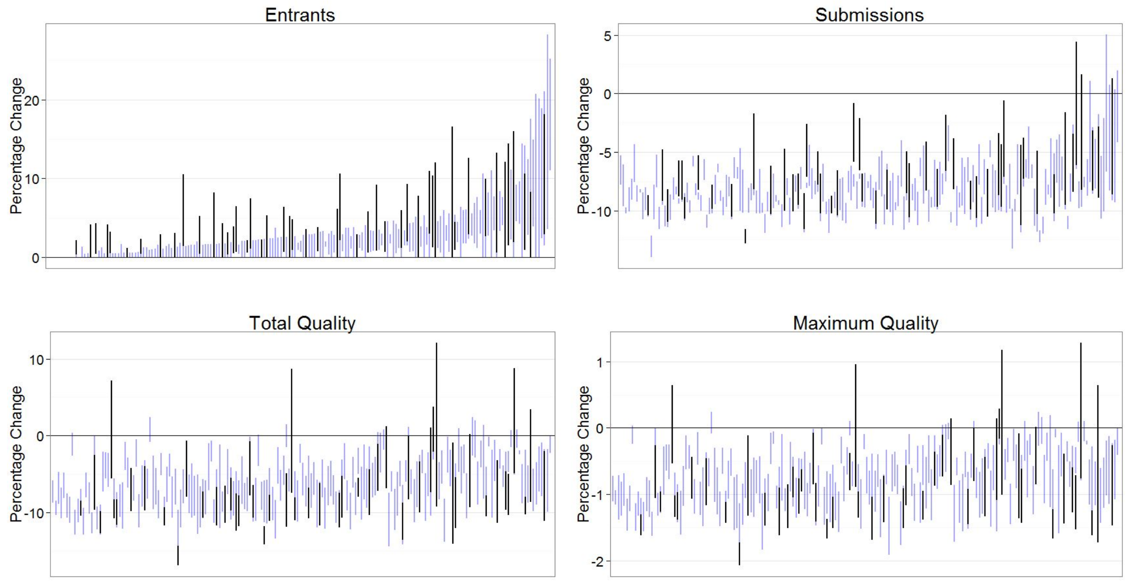

Figure 1.7: *Impact of a 3 Submission Limit Across Contests*

Note: Each segment represents the range of counterfactual outcomes for a single contest under complete information (light) and incomplete information (dark). Contests ordered by increasing impact on the number of entrants in the complete information scenario.

Interestingly, many other contest-based freelance marketplaces (such as 99designs for graphic design) offer sponsors the option to organize a contest where participants can see the submissions of their competitors. Research has considered the impact of visibility and free-riding on entry incentives and submission quality (Wooten and Ulrich, 2015b). My results suggest that offering contests with complete information can have an impact beyond free-riding that directionally depends on the extent of participant heterogeneity.

1.8.3 Practical Implications

These findings have a number of practical implications for ideation contest and crowdsourcing platform design. First, the choice of how many prizes to offer does not substantially affect the outcome of the contest, as long as the contest attracts a large number of submissions. The choice of how many prizes to offer should be driven by institutional considerations. For example, in many ideation contests, the sponsor retains intellectual property of the winning submissions. In these settings, the sponsor would benefit from offering more prizes, without significantly altering the outcome of a contest. Only in complete information settings where the sponsor expects to receive submissions from a very homogeneous set of participants does a single prize appear preferable.

Second, if a sponsor's intention is to increase the number of entrants, increasing prize award may not have as strong an impact if there is no limit on the maximum number of submissions a participant can make. Third, a submission limit can be used as an effective strategy to encourage entry but may come at the cost of expected total and maximum idea quality. If a sponsor seeks to attract a large number of entrants and use the contest as a mechanism for engaging potential consumers, it should implement a more stringent submission limit.

1.9 Conclusion

Firms across a range of industries use ideation contests to procure ideas for ads, new products, and marketing strategies. An appropriate design can improve the outcome of a contest. Moreover, different firms may care about different outcome metrics. Brands interested in engaging consumers may focus on increasing entry, whereas a manufacturer interested in designing a new product may value the maximum quality of submitted ideas.

I empirically investigate the impact of three design parameters - number of prizes, prize amount, and submission limit - on contest participation and quality outcomes, using data from a popular crowdsourcing studio that runs ad ideation contests for major brands. I present a structural model of ideation contests that allows for multiple equilibria, incomplete information, and heterogeneity in participant submission quality and costs. Counterfactual simulations reveal the impact of different contest designs. The results show that, on average, the number of prizes does not significantly affect contest outcomes, prize amount increases submissions and all expected quality metrics but may not necessarily increase entry, and submission limits encourage entry but significantly reduce expected total and maximum quality.

I make several simplifying assumptions to ensure the model remains feasible. First, I assume that each contest is an independent game. Participants face no dynamic incentives and do not have constraints that prevent them from entering multiple contests at the same time. Future research may examine the implications of dynamics and competing contests on the optimal design of contest platforms. Second, I assume that participants choose the quantity but not the quality of submissions to make. Future research may allow for participants to choose not only how many submissions to make but also how much effort to invest into each individual submission. Finally, I do not observe the actual applications of ideas obtained through the contests. Absent these data, quality is inferred from jury ratings and sponsor rankings, and does not necessarily represent the market's perception of idea quality. Future work may incorporate post-contest outcomes to assess the impact of contest design on the true value of winning ideas.

An appropriately designed ideation contest can yield interesting concepts and spur innovation. Crowdsourcing platforms, sponsors, and contest designers must carefully consider the effects of different design parameters on outcomes.

Chapter 2

Expiration and Pricing Policy with Inattentive Consumers: The Case of Daily Deals[1]

2.1 Introduction

Daily deal sites are online marketplaces for discount vouchers such as Groupon and LivingSocial that offer local merchants a platform to sell consumers deeply discounted deals for products. Daily deals typically offer deep discounts of 50% or more for a period of time. Merchants benefit from participation by gaining access to a large customer base, and consumers benefit from access to a large number of discounted products from merchants they would have otherwise not known about. The daily deal marketplace has access to a number of levers that can affect consumer decisions and platform profits. We investigate the impact of expiration dates, prepayment, and revenue-sharing structures on consumer choice, merchant decisions, and platform profits.

How are consumer purchase and redemption decisions driven by deal characteristics, inattention, and uncertainty? What would change if firms were required by law to offer deals with extended expiration date? Does a prepayment policy increase profits? Do alternative revenue-sharing agreements favor merchants, the platform, or both? To address these questions, we develop

[1]Co-authored with Vineet Kumar and Xueming Luo.

an empirical model of the marketplace. Merchants set discounts to maximize their expected future stream of payoffs, subject to the expiration and pricing policies enforced by the platform. Consumers purchase deals anticipating future redemption opportunities and the possibility of forgetting about the deal. After purchase, a consumer who forgets cannot redeem until she remembers. Consumers who remember decide whether or not to redeem in the current period or in the future. Uncertainty in future redemption availability and future attention implies that consumers face a dynamic trade-off. Immediate redemption would yield immediate gratification, whereas delayed redemption may yield a higher level of gratification in the future if redemption conditions are more favorable and the consumer remembers about the deal. At the same time, if the consumer delays redemption for too long, she may fail to redeem the deal altogether as a result of inattention or a lack of favorable redemption opportunities.

We use a novel consumer-level dataset from a large Asian daily deal operator to estimate the model. Individual-level purchase and redemption patterns help us identify behavioral parameters, such as consumer preferences for deals and the probabilities of forgetting and remembering. Our estimates show that inattention significantly influences consumer purchase and redemption decisions. In a series of counterfactuals we explore the impact of different expiration and pricing policies on consumers, merchants, and the platform.

2.1.1 Industry Background

In 2011, the expiration policy of daily deals stated that a consumer who failed to redeem a purchased deal before it expires could no longer leverage the discount. However, state and federal regulations in the US prohibited the sale and marketing of prepaid vouchers with expiration dates.[2] This inconsistency resulted in a series of class-action lawsuits against major daily deal operators. Consumer advocates argued that according to the Credit Card Accountability Responsibility and Disclosure (CARD) Act of 2009, daily deals could not expire. In 2012, Groupon settled its lawsuits for $8.5 million. In 2013, competitor LivingSocial settled for $4.1 million after roughly 26,830 consumers filed claims against the company.

Daily deal platforms typically require consumers to undertake a risk by prepaying for products

[2] See Edelman and Kominers (2011) for further details on the regulatory and consumer protection issues facing daily deals.

consumed in the future. Consumers often forget to redeem purchased deals before expiration. Non-redemption rates can be as high as 20%. In contrast, online coupon exchanges, such as RetailMeNot, only require consumers to pay when they retrieve the product. Prepayment exacerbates the legal risks faced by daily deal operators but may potentially lead to greater merchant and platform profits.

Revenue-sharing agreements between the daily deal platform and its merchants often condition on the redemption status of a deal. If a deal goes unredeemed, merchants obtain a smaller share of the revenues. This behavior may drive merchants to offer smaller discounts and draw on revenues from unredeemed deals. We use counterfactual simulation to investigate the impact of extended expiration dates, a no-prepayment policy, and an alternative revenue sharing agreement on consumers, merchants, and the platform in the context of a popular Asian daily deal marketplace.

2.1.2 Literature

We build on the literature on the impact of expiration dates and pricing on inter-temporal consumer decision-making and firm behavior in marketing contexts. Traditionally, these questions have been studied in the context of rebates, gift cards, and coupons. Silk and Janiszewski (2008) consider the design of mail-in rebate programs and develop a number of hypotheses of how length of rebate period and reward amount affect product sales and rebate redemptions. Rebates are similar to daily deals in that consumers must first purchase a product to later secure the benefit of the rebate but differ from daily deals and coupons as consumers must incur a cost today (mail-in the rebate) to secure the reward in the future. Although the author does not test the hypotheses in the paper, he points out that expiration and pricing policy is important in the context of rebate programs.

Expiration policy is also very important in gift card markets. Gift cards are similar to daily deals in that they must be purchased prior to consumption and may expire. Different states in the US differ in their gift card expiration policies. Many states requiring by law that gift cards never expire, whereas others allow for gift cards to expire in as little as two years (Offenberg, 2007).

Inman and McAlister (1994) study coupon redemption. Daily deals differ from coupons in that consumers need to decide on whether or not to purchase a daily deal for the option to consumer a product at a later date, whereas a coupon is typically awarded to the consumer. The authors identify a U-shaped coupon redemption pattern driven by an increase in redemptions near expiration for spaghetti sauce and use anticipated regret theory to build a predictive model and explain this pattern. We offer an alternative explanation in the context of daily deals. Consumers may forget about deals and redeem near expiration because of reminders. We observe daily deal purchase and redemption patterns at the consumer-level, which allows for us to show that a model of consumer inattention and preference heterogeneity can explain the observed purchase behavior and U-shaped redemption patterns.

A small but growing empirical literature considers the implications of inattention in consumer behavior. Mehta *et al.* (2004) develop a structural model of learning and forgetting about brand quality in a repeated purchase decision context. The authors cast forgetting as the noise in a consumer's recall of brand quality from past purchase experiences. Sahni (2015) develops a memory-based model to explain why a wider spacing of ads may lead to higher purchases. In contrast to the above studies where consumers forget a brand characteristic or have a weaker ability to recall past information, we study inattention in consumer choice. In other words, consumers may forget that they have an action to make altogether.

In the context of daily deals, a number of authors consider the effectiveness of daily deals as a promotional device for merchants, but few papers explicitly study consumer purchase and redemption behavior for daily deals. An exception is Luo *et al.* (2014), who study the impact of deal popularity and social influence on deal purchases and redemptions using the same data as in this paper. Edelman *et al.* (2014) argue that daily deals may be profitable if they attract new customers or allow the firm to reach more price-sensitive customers. Byers *et al.* (2012a,b) find that merchant ratings on popular online review site Yelp tend to decrease after the merchant offers a deal on Groupon. The authors argue that the decrease occurs as Groupon users tend to be "real customers,"in the sense that they reduce the bias resulting from fake or inflated reviews on the Yelp platform. We do not focus on the long-term implications of daily deal marketing for merchants. Instead, we explore how merchants would react differently to alternative platform policies, and the implications this has for consumers, merchants, and the platform.

We contribute to the literature on two fronts. First, we develop a model of a daily deal marketplace that takes into consideration both consumer purchase and redemption behavior, and

Table 2.1: *Summary Statistics*

Weeks	46
Consumers	1559
Purchases	3341
Redemption Failures	312

merchant incentives to offer discounts. Second, we apply our model to explore how alternative designs would affect the daily deal marketplace. Our analysis may shed light on the impact of different expiration dates and pricing strategies for products other than daily deals, such as rebates, gift cards and coupons. Moreover, few studies consider the impact of alternative policies on both consumer and merchant decisions. The remainder of the paper is organized as follows. Section 2.2 presents the data and the conceptual framework. In Sections 2.3 and 2.4 we present an empirical framework for modeling behavior in daily deal marketplaces. Section 2.5 presents the results of our estimation routine. In Section 2.6, we summarize the insights of several counterfactuals. Finally, Section 2.7 concludes.

2.2 Data

We obtain data from a popular daily deal platform operating in a large Asian metropolis. We observe the purchase and redemption times of registered consumers, as well as the deals available for purchase over the course of six months from the inception of the platform in July 2011. We focus on a sample of 39 deals and 1,559 consumers who purchased more than once. Table 2.1 presents summary statistics of the sample of consumers. No individual consumer purchases a substantially large number of deals. The median consumer purchases 2 deals over the course of 46 weeks, and the heaviest consumer purchases 5 deals. For each deal, we observe its availability and expiration dates, category, price, savings, and an indicator for location. Table 2.2 provides a summary of deal characteristics. Across the 39 deals, we observe that 26 are offered in the city center. Deal redemption periods range from 7 to 29 weeks. Purchase prices range from $2.56 to $47.68 and savings range from 54 to 97 percent. For almost all deals, the purchase period lasts one week and coincides with the first week of the redemption period.

For each consumers, we observe when they purchased and redeemed a deal, or if the deal

Table 2.2: *Summary of Deal Characteristics by Category*

Category	# Deals	In City Center	Price ($)			Savings (%)			Duration (Weeks)		
			Min	Avg	Max	Min	Avg	Max	Min	Avg	Max
Restaurant	15	10	3.04	12.94	22.08	55	67	83	7	11	14
Beauty	4	4	4.64	8.52	12.48	70	84	97	10	12	15
Photography	5	5	2.56	21.99	47.68	81	87	94	12	19	29
Hair Styling	7	6	3.04	11.31	15.36	90	92	96	7	11	14
Entertainment	4	1	3.52	5.76	9.28	63	71	78	9	15	22
Health	4	4	10.88	22.91	36.48	54	67	83	10	15	20
Total	39	26									

expired. Figure 2.1 presents redemption frequency as a function of time since purchase before expiration, aggregated across all deals used in estimation, controlling for the fraction of time a deal was redeemable in the first and last week. The left panel shows the fraction of consumers who redeemed at a given time since purchase out of the total number of purchasers. The curve is noticeably U-shaped, with peaks at purchase and just before expiration. The right panel shows the fraction of redemptions conditional on a consumer having waited a given time since purchase. In other words, it is the conditional probability that a consumer will redeem a deal at a given time since purchase as estimated from the aggregated data. Interestingly, this curve also decreases at first and then steeply increases. If consumers were identical and forward-looking, we would expect the conditional probability of redemption to increase as the deadline approaches. As the expiration date approaches, the option value of holding a deal falls. As the utility of redeeming does not change with time (product characteristics are time invariant), the decrease in a deal's option value reduces the expected utility a forward-looking consumer receives from waiting, increasing the probability of redemption.[3] Figure 2.1 suggests that forward-looking behavior is not the only force in play.

2.2.1 Conceptual Framework

Our goal is to develop a model that explains the U-shaped redemption pattern observed in the left panel of Figure 2.1. This aggregate pattern may arise from a number of behavioral phenomena. We review the most popular explanations below and motivate our modeling framework.

[3] If consumers were myopic and identical, we would observe a constant conditional redemption probability.

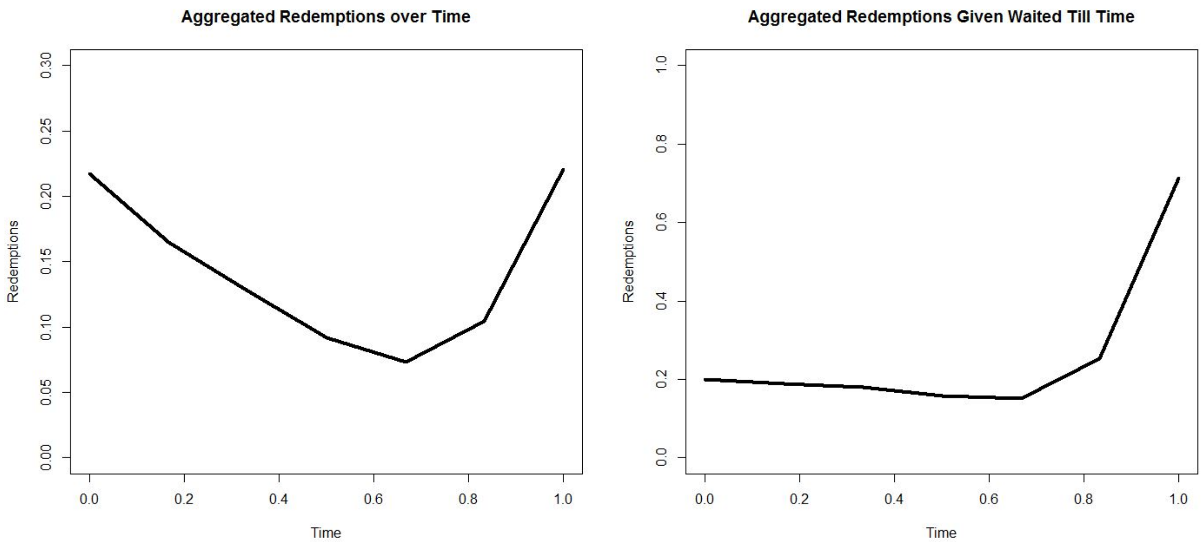

Note: Left panel shows fraction of redemptions at each time period out of all total purchasers. Right panel shows fraction of redemptions at each time period out of purchasers who waited until that time period. Duration of all deals is normalized to 1. Availability of deals in each week is normalized to 1.

Figure 2.1: *Aggregate Redemption Curves*

The U-shaped redemption pattern may arise as a result of consumer inattention. We observe a high level of redemptions at the beginning of the redemption period as most consumers had just purchased and are attentive. As time progresses, consumers begin to forget and fall into a state of inattention. As a result, we observe low levels of redemption for intermediate time periods. Finally, as the deadline approaches, consumers are more likely to remember about the deal as a result of reminders, which leads to a high number of redemptions near the deadline. Inattention is a popular explanation for why redemptions occur near expiration in the context of rebate programs (Silk and Janiszewski, 2008). Furthermore, the right panel in Figure 2.1 supports the inattention mechanism. As time progresses, consumers are more likely to fall inattentive, resulting in a lower conditional probability of redemption near the purchase date. Alternatively, this pattern can be explained by unobserved preference heterogeneity - a possibility we consider below.

An alternative explanation for the U-shaped curve is regret aversion. In the context of coupon redemptions, Inman and McAlister (1994) argue that as coupon expiration approaches, consumers anticipate the loss they may feel if the coupon expires and are more likely to redeem as a result. We lack the data to separate regret aversion from inattention in our framework. More specifically, regret-aversion is confounded with reminder effects. We can incorporate regret aversion in the following sense. Near expiration, consumers may be more likely to respond to reminders if they are regret-averse. That is, regret aversion will affect the probability that a consumer will be attentive near expiration but not the utility of redemption.

Unobserved heterogeneity in redemption preferences can also explain the U-shaped redemption curve and the conditional redemption frequencies in Figure 2.1. Namely, consumers with a low redemption utility would tend to redeem later than consumers with a high redemption utility. For example, certain consumers may be busier than others, and as a result, may be more likely to delay redemption into the future. If there is a sufficient amount of both types of consumers in the sample, we would observe a bi-modal redemption curve with peaks near purchase and expiration. Moreover, as high redemption utility consumers redeem and exit the sample, the consumers who remain have a lower redemption utility, resulting in a decreasing conditional probability of redemption near purchase. We are able to incorporate unobserved preference heterogeneity in our framework by drawing on repeated purchase and redemption observations per consumer and by studying the discrepancies in purchase and redemption behavior across deals. Using repeated purchase instances we can show that it is not the case that two a priori different types of consumers are driving the redemption pattern. The correlation of redemption times between the first and second purchase for consumers is 0.12. A consumer's prior redemption behavior has a limited ability to predict her future redemption, suggesting that it is unlikely that unobserved heterogeneity alone drives the U-shaped redemption pattern.

We proceed by developing a model that incorporates unobserved preference heterogeneity and inattention with reminder effects near expiration. We do not model learning about deals or about inattention. This modeling decision is motivated by the fact that we do not find evidence of learning in the data.[4]

[4]Namely, the redemption curves for first-time and second-time purchasers look very similar. If there was learning about inattention, or about the quality of deals, we would expect the distribution of redemptions to differ as consumers with experience would select deals more intelligently and be more likely to redeem deals. We also observe that there is no significant difference in first-time redemptions for consumers who end up purchasing twice and consumers who end up purchasing three or more times, suggesting that past experience with a deal does not influence future purchase behavior. We may fail to find evidence of learning simply because we do not observe the market for a sufficiently long time. Nevertheless, given the lack of evidence in the data, we proceed to develop a model without learning.

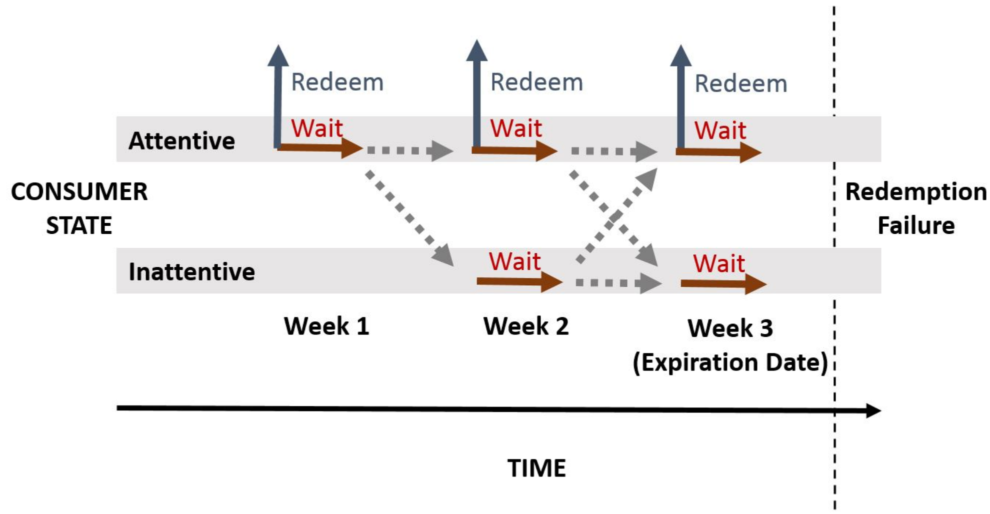

Figure 2.2: *Consumer Decision Process for a Deal Redemption Decision*

2.3 Consumer Choice Model

We develop a dynamic discrete choice model of consumer purchase and redemption decisions. We treat each deal independently and model the consumer decision process for purchasing and redeeming each deal. Consumers purchase a deal if the utility from holding or redeeming the deal exceeds the utility of the outside option. At any point in time after purchasing a deal, the consumer can be attentive or inattentive. An inattentive consumers cannot redeem a purchased deal. An attentive consumers can choose to redeem the deal or to delay the redemption decision into the future. Uncertainty surrounding weekly redemption conditions and future attention requires that consumers form expectations of the future utility they may derive from holding a deal, taking into account the possibility of redemption failure.

Figure 2.2 summarizes the consumer decision process for deal redemption. In the first week that the deal is redeemable, the consumer decides whether or not to redeem immediately. If the consumer decides to wait, she faces the same redemption decision in the following week given that she remembers about the deal, and so on, until the deal expires. At any time $t > 1$, the consumer may transition between states of attentiveness and inattentiveness given a process we refer to as memory activation.

2.3.1 State Space and Information Set

Let i denote the consumer, j denote the deal, and $t \in \{1, ..., T_j\}$ denote the time period, where $t = 1$ denotes the first week when the deal becomes available for purchase, and T_j denotes the week of expiration. The state of consumer i for deal j at time t is defined as $s_{ijt} = \{X_{jt}, \epsilon_{ijt}, \{a_{ijk}\}_{k=1}^{t}\}$ where X_{jt} is a vector of time-invariant and time-varying deal characteristics, which we assume the consumer knows for all deals that have been offered up to time t, ϵ_{ijt} is a vector of independent and identically distributed shocks observed by the consumer but not by the researcher, and $\{a_{ijk}\}_{k=1}^{t} = \mathcal{A}_{ijt}$ is the consumer's history of attention, where $a_{ijt} \in \{0, 1\}$ is the consumer's state of attentiveness at time t, with $a_{ijt} = 0$ denoting inattention. The consumer is always attentive at $t = 1$. For $t > 1$, the consumer's attentiveness evolves according to $a_{ijt} = I(m_{ijt} > 0)$, where $I(.)$ is the indicator function and $m_{ijt} = m_{ijt}(\mathcal{A}_{ijt-1})$ is a measure of the consumer's memory activation that is allowed to depend on the history of attention. The consumer has rational expectations about the evolution of a_{ijt} and ϵ_{ijt}.

2.3.2 Choice Set

We denote by $d_{ijt} \in \{R, W, NP\}$ the consumer's decision, where R stands for 'Redeem', W stands for 'Wait', and NP stands for 'Not Purchase'. At time period t, the consumer's choice set depends on the consumer's state and whether or not deal j is purchasable and/or redeemable. At $t = 1$, the consumer can always not purchase the deal. If the deal is available for redemption at $t = 1$, the consumer can also choose to purchase and redeem immediately (R) or to purchase and wait (W). For $t > 1$, if the consumer has purchased the deal and waited, and is attentive in that period, she can choose to redeem the deal or continue waiting. This decision structure is repeated until the deal expires or the consumer redeems the deal. To summarize, $d_{ij1} \in \{R, W, NP\}$ if the deal is redeemable at time 1, and $d_{ij1} \in \{W, NP\}$ otherwise. For $1 < t \leq T_j$, $d_{ijt} \in \{R, W\}$ if the consumer is attentive ($a_{ijt} = 1$) and has not redeemed the deal prior to period t, and $d_{ijt} = W$ otherwise.

2.3.3 Consumer Decision Process

Suppose that consumer i must decide when to redeem a deal j. The per-period utility of redeeming deal j at time t is $u_{ijt}^{R} + \epsilon_{ijt}^{R}$, and the utility for waiting and not redeeming is $u_{ijt}^{W} + \epsilon_{ijt}^{W}$, where

u^W_{ijt} and u^R_{ijt} are the deterministic components of utility always known to consumers and the ϵ^R_{ijt} and ϵ^W_{ijt} are the stochastic components such that $\epsilon_{ijt} = \{\epsilon^R_{ijt}, \epsilon^W_{ijt}\}$. At the beginning of time period t, the consumer is revealed values of ϵ^R_{ijt} and ϵ^W_{ijt}, but she is uninformed of their future values. The consumer inherently faces an optimal stopping problem. She can either redeem immediately or wait to redeem later when the realizations of the stochastic components may be more favorable. Recall, that the consumer may be inattentive in the future period. As a result, she must consider the possibility of forgetting when forming future expectations. As a result, both uncertainty in the future realizations of the structural errors and future inattention play a role in the consumer's decision process. We develop an expression for the deal-specific continuation value $v_{ijt}(\epsilon^R_{ijt}, \epsilon^W_{ijt}, \mathcal{A}_{ijt})$, which is a function of the realizations of the stochastic components $\epsilon^R_{ijt}, \epsilon^W_{ijt}$ for redeeming or holding deal j at time t and attention history $\mathcal{A}_{ijt}$. If the consumer looses the deal, either by redeeming it or by allowing it to expire, the continuation value is zero. Let T_j denote the expiration time of deal j. The value function for the dynamic program solved by the consumer for $t < T_j$ can be written as

$$v_{ijt}(\epsilon^R_{ijt}, \epsilon^W_{ijt}, \mathcal{A}_{ijt}) = \max\{u^R_{ijt} + \epsilon^R_{ijt}, u^W_{ijt} + \epsilon^W_{ijt} + \delta E_{\epsilon_{ijt+1}, \mathcal{A}_{ijt+1}|\mathcal{A}_{ijt}} \left[v_{ijt+1}(\epsilon^R_{ijt+1}, \epsilon^W_{ijt+1}, \mathcal{A}_{ijt+1}) \right] \}$$

if $a_{ijt} = 1$ and

$$v_{ijt}(\epsilon^R_{ijt}, \epsilon^W_{ijt}, \mathcal{A}_{ijt}) = u^W_{ijt} + \epsilon^W_{ijt} + \delta E_{\epsilon_{ijt+1}, \mathcal{A}_{ijt+1}|\mathcal{A}_{ijt}} \left[v_{ijt+1}(\epsilon^R_{ijt+1}, \epsilon^W_{ijt+1}, \mathcal{A}_{ijt+1}) \right]$$

if $a_{ijt} = 0$, where δ is the discount factor. At terminal period $t = T_j$, the value function is given by

$$v_{ijT_j}(\epsilon^R_{ijT_j}, \epsilon^W_{ijT_j}, \mathcal{A}_{ijT_j}) = \max\{u^R_{ijT_j} + \epsilon^R_{ijT_j}, u^W_{ijT_j} + \epsilon^W_{ijT_j}\}.$$

if $a_{ijt} = 1$ and $v_{ijT_j}(\epsilon^R_{ijT_j}, \epsilon^W_{ijT_j}, \mathcal{A}_{ijT_j}) = u^W_{ijT_j} + \epsilon^W_{ijT_j}$ otherwise.

Now suppose that $t = 1$, and consumer i is considering whether or not to buy deal j. We model the utility of not purchasing the deal and consuming the outside option as $u^{NP}_{ij} + \eta_{ij}$, where u^{NP}_{ij} is the deterministic component of utility, and η_{ij} is a structural error revealed to the consumer but not the researcher when deal j becomes available at $t = 1$. All deals in our analysis are available for purchase during the first week only, and hence adding a subscript t to the no-purchase decision utility is unnecessary. At this stage, if the deal is redeemable in the first week as well, the consumer has three options. She can not purchase the deal and derive utility $u_{ij}^{NP} + \eta_{ij}$, she can purchase the deal and redeem it immediately in the same time period to derive utility $u_{ij1}^{R} + \epsilon_{ij1}^{R}$, or she can purchase the deal and choose to wait until next period, gaining $u_{ij1}^{W} + \epsilon_{ij1}^{W} + \delta E_{\epsilon_{ij2},\mathcal{A}_{ij2}|\mathcal{A}_{ij1}} \left[v_{ij2}(\epsilon_{ij2}^{R}, \epsilon_{ij2}^{W}, \mathcal{A}_{ij2}) \right]$ in the current period. The utility of not purchasing captures the payoff the consumer derives from consuming an outside option. The utility of redeeming captures the payoff the consumer receives from consuming the product offered by the daily deal. The utility of waiting captures the consumer's payoff from spending the time needed to redeem the deal on other activities, as well as the option value of holding deal j and possibly redeeming it in the future.

2.3.4 Parametrization

We now discuss each component of the model in turn and its chosen parametrization. Namely, we require parametrizations of the components of the consumer decision process - the deterministic utilities u_{ijt}^{R}, u_{ijt}^{W} and u_{ijt}^{NP}, the stochastic components $\epsilon_{ijt}, \eta_{ij}$, and the attention process, governed by the signal function m_{ijt}. We fix the consumer discount factor at 0.99.

Components of the Consumer Decision Process

The deterministic utility of redeeming deal j at time t is given by

$$u_{ijt}^{R} = \alpha_i^0 + \alpha_i^h \text{HEALTH}_j + \alpha_i^e \text{ENTERTAINMENT}_j + \alpha_i^l \text{LOCATION}_j + \alpha_i^s \text{SAVINGS}_j + \alpha_i^a \text{AVAILABLE}_{jt} - \alpha_i^p I(t=1)\text{PRICE}_j$$

where α_i^0 is an intercept term, α_i^h and α_i^e are category specific coefficients that capture differences in individual preferences for products and differences in how products from different categories are consumed. We pool certain categories together because of their similarity and to increase the number of observations in each category. The variable HEALTH_j is an indicator for beauty, haircare, and health deals, and the variable ENTERTAINMENT_j is an indicator for photography and entertainment deals. The remaining deals are in the restaurant category, and the associated effect is captured by the intercept term α_i^0. The remaining terms capture consumer preferences for deal characteristics. These terms may have a structural interpretation but may also proxy for deal quality. The LOCATION_j variable is an indicator for whether or not the deal merchant is located in the city center. The parameter α_i^l captures consumer preferences for deal location. The parameter α_i^s captures the impact of deal savings, SAVINGS_j, on the consumer's redemption utility. High savings may increase the consumer's willingness to redeem and leverage the savings but may also correspond to a low quality deal. The AVAILABLE_{jt} variable measures the fraction of week t that deal j is redeemable. If a deal is redeemable for only a fraction of a time period, the consumer may find it more difficult to find a redemption opportunity and the extent of this is captured by the parameter α_i^a. The term PRICE_j is the price of deal j, and the coefficient α_i^p captures consumer price sensitivity. The consumer must pay the price of the deal in the first period when she chooses to purchase (and redeem immediately or wait). We normalize the mean of the deterministic utility of waiting and not redeeming a deal to zero. We can then write $u_{ijt}^W = -\alpha_i^p I(t=1)\text{PRICE}_j$. The deterministic component of utility associated with the no-purchase decision can be written as

$$u_{ij}^{NP} = \beta^r + \beta^h \text{HEALTH}_j + \beta^e \text{ENTERTAINMENT}_j$$

where we allow for category-specific fixed effects to capture the time and effort costs of purchasing a deal (such as completing payment forms, accessing the category page on the website, etc.).

The Attention Process and Unobserved Heterogeneity

Note that the model features two main sources of unobserved heterogeneity. First, consumers may differ in their preferences for redeeming deals. This is a source of persistent unobserved heterogeneity, captured by the preference parameters indexed by i. These parameter enters not only through the redemption utility but also through the value function of the utility of waiting and control for the possibility of dynamic selection. Namely, a consumer may wait long before redeeming a deal because of a low redemption utility and not inattention. Second, consumers may differ in their attention state at any time period $t > 1$. This is a source of time-varying unobserved heterogeneity that allows for consumers to differ endogenously as a result of the different paths their attention states may follow.

Recall that all consumers are attentive at $t = 1$, or $a_{ijt} = 1$, but for $t > 1$, $a_{ijt} = I(m_{ijt}(\mathcal{A}_{ijt-1}) > 0)$. That is, only if the memory activation function m_{ijt} exceeds a threshold will consumer i be attentive in period $t > 1$ for deal j. We parametrize this function as

$$m_{ijt}(a_{ijt-1}) = \gamma^0 + \gamma^a a_{ijt-1} + \sum_{k=T_j-l}^{T_j} \xi_k I(t=k) + \nu_{ijt}$$

which depends on an intercept, the consumer's prior attentiveness state a_{ijt-1}, a set of fixed effects for time periods prior to expiration, and a stochastic component ν_{ijt}. We would expect that a consumer who was attentive in period t is more likely to be attentive in period $t+1$ and vice versa. Hence, the lagged attention term in m_{ijt} will capture any serial correlation in attention states.[5]

The set of fixed effects, $\{\xi\}_{T_j-l}^{T_j}$ for $l \geq 0$, capture the additional impact of any reminders or other contextual cues in the weeks just prior to expiration on the consumer's memory. Given the many different ways and time periods in which a consumer may choose to record a reminder, we flexibly model the effect of reminders on attention through fixed effects. In practice, we must fix an l and estimate the associated set of fixed effects.[6] We assume $\xi_{T_j-l} = \xi_{T-l}$ are identical across deals and estimate specifications with $l=1$ and $l=2$. We choose to focus the case when $l=2$ based on model fit. Furthermore, we find that restricting $\xi_{T-l} = \xi$ for $l = 1, 2$ does not yield substantially different results. The resulting memory activation function then becomes

$$m_{ijt}(a_{ijt-1}) = \gamma^0 + \gamma^a a_{ijt-1} + \xi I(t \geq T_{j-1}) + \nu_{ijt}$$

where the parameter ξ captures deadline effects on memory. Finally, the model includes a stochastic component ν_{ijt}. This error term will capture weekly consumer-deal specific variation in memory activation not explained by the observables and past attention. We assume ν_{ijt} follows a standard logistic distribution. This implies that the probability a consumer is attentive in period $t > 1$ given her state of attention in period $t-1$ is $\Pr(a_{ijt} = 1 | a_{ijt-1}) = e^{m_{ijt}(a_{ijt-1})} / (1 + e^{m_{ijt}(a_{ijt-1})})$.

Unobserved attention states endogenously induce differences across consumers over time, whereas persistent unobserved heterogeneity captures time-invariant differences in the preferences

---

[5]We assume that attention states prior to period $t-1$ do not affect memory activation at time t to reduce the complexity and computational burden of the model. As a result, $m_{ijt}(\mathcal{A}_{ijt-1}) = m_{ijt}(a_{ijt-1})$.

[6]In theory, it is possible to set $l = T_j - 1, \gamma^0 = \gamma^a = 0$ and estimate the model $m_{ijt} = \xi_{jt} + \nu_{ijt}$ for $t > 1$. However, this specification would not get at the mechanism driving inattention and substantially increase the number of parameters. As a result, we restrict l to be a small number to capture deadline effects and develop a parametric form based on theory to explain behavior in all remaining time periods.

of consumers. We allow for a latent class structure for the parameters governing persistent unobserved heterogeneity.[7] Let $\theta_i = \{\alpha_i^0, \alpha_i^h, \alpha_i^e, \alpha_i^l, \alpha_i^s, \alpha_i^a, \alpha_i^p\}$ denote the vector of random parameters for consumer i. Each consumer i belongs to a type $k \in \{1, ..., K\}$. A consumer of type k has preferences governed by parameters θ_k and the set of homogenous parameters $\eta = \{\beta^r, \beta^h, \beta^e, \gamma^0, \gamma^a, \xi\}$. We assume 2 types of consumers labelled $k = 1$ and $k = 2$, such that the fraction of type 1 consumers in the data is given by the parameter ϕ.

Integrated Value Function

To complete the decision model, the stochastic components $\epsilon_{ijt}^R, \epsilon_{ijt}^W$ and η_{ij} are independent and identically distributed across consumers, deals and time periods, and follow a demeaned type-1 extreme value distribution. We label Ev_{ijt} the deal-specific integrated value function, or the expectation at time t of the value function just before the $\epsilon_{ijt}^R, \epsilon_{ijt}^W$ are revealed. Letting

$$w_{ijt}(a_{ijt}) = u_{ijt}^W + \delta \Pr(a_{ijt+1} = 1|a_{ijt}) Ev_{ijt+1}(1) + \Pr(a_{ijt+1} = 0|a_{ijt}) Ev_{ijt+1}(0)$$

the deal-specific integrated value function can be written as

$$Ev_{ijt}(a_{ijt}) = \log\left(e^{u_{ijt}^R} + e^{w_{ijt}(1)}\right) I(a_{ijt} = 1) + w_{ijt}(0) I(a_{ijt} = 0)$$

The deal-specific integrated value function represents to option value of holding deal j at time t for consumer i.

Discussion

We discuss some of our modeling assumptions and how they relate to some of the findings in the behavioral literature on memory and redemption. First, we assumed that consumers exhibit rational inattention. In other words, the consumers know the parameters of the function m_{ijt} and behave accordingly, taking correctly into account the possibility of falling in and out of a state of attention in the future when making current-period decisions. This assumption does not allow for consumers to have incorrect expectations about the probability that they will forget about

[7] Longer consumer panels yield more accurate estimates of consumer-specific parameters. Given the relatively small number of redemption decisions we observe per consumer, we restrict parameters of the signal function to be the same across consumers.

a purchased deal. The behavioral literature has identified the possibility of overconfidence in redemption behavior in experimental studies (Shu and Gneezy, 2010; Taubinsky, 2013; Tasoff and Letzler, 2014). Overconfident consumers overestimate the probability of remembering about a deal. This would imply that consumers attach a greater option value to a deal, as they do not discount it by the correct probability of inattention. Overconfidence about memory can lead to lower redemption rates for longer deals, whereas a model of rational attention predicts that increasing a deals duration increases redemptions (Taubinsky, 2013). As we lack the data to identify consumer expectations, we make the assumption that consumers have rational expectations.[8]

Second, we assumed that the consumer utility function does not incorporate a cost term for when a deal expires. Regret-aversion theory states that consumers experience an emotional cost when a deal expires. To avoid this cost, a forward-looking consumer who had not redeemed earlier would be inclined to redeem near expiration. In the context of our model, increased redemptions near expiration are driven by reminders (or what we refer to as deadline effects). We cannot separate the effect of reminders on the memory activation function from the impact of regret on the utility function as both imply very similar data patterns.[9] In our framework, the effects of any loss or regret will be absorbed by the ξ term in the memory activation function that captures deadline effects.

Third, we assume that category-specific fixed effects and the set of variables included in the redemption utility suffice to control for unobserved differences in quality and advertising across deals, such that the price is not correlated with the unobservables. Redemption data contains information about deal quality, as we would expect higher quality deals to obtain higher redemptions and yield a higher option value to consumers. This mitigates concerns about price endogeneity, as higher quality and possibly more heavily advertised deals would yield a higher option value to consumers. In other words, option value is a measure of the unobserved characteristics of a deal that may be correlated with price.

[8]Namely, overconfidence is confounded with unobserved preference heterogeneity in our context.

[9]Inman and McAlister (1994) apply a model of regret-aversion to study similar U-shaped redemption patterns in the spaghetti sauce category. The authors do not attempt to separate regret-aversion from inattention.

2.3.5 Identification

First, assume there is no unobserved preference heterogeneity by dropping the i subscript on all utility parameters. The only source of heterogeneity in the model (other than the structural errors) is then consumer attention. Given our assumptions of a known discount factor, extreme value errors, and rational inattention, we can identify (homogenous) utility parameters in u^R_{ijt}, u^W_{ijt}, and u^{NP}_{ij}, as well as the parameters of the memory activation function m_{ijt}.

We observe purchase decisions for a set of deals, as well as weekly redemption decisions for the set of consumers who purchased each deal. The utility and memory activation parameters are jointly identified from these data. The utility parameters are identified by the shares of purchases and redemptions in the first period, and the memory activation parameters are identified by redemption patterns in subsequence periods. Recall that all consumers are attentive in the first period. If we observe many consumers purchase a deal but few of the purchasers redeem in the first period, than it must be the case that the utility of purchasing the deal was high but the utility of redeeming it was low. Then if we observe that the number of redemptions in period 2 is significantly less than in period 1, then it must be the case that the probability of attention in time period 2 is low. This is because consumers anticipate the possibility of forgetting at time 1 and are more likely to redeem than wait, and also because consumers who choose to wait tend to forget at time 2 and fail to redeem. In general, redemption behavior at time periods $t = 2, ..., T_j$ identifies the associated memory activation function m_{ijt}. This is because the utility function associated with redemption u^R_{ijt} depends on time only through the variable AVAILABLE$_{jt}$, which also varies across deals in the first time period. Hence, u^R_{ijt} is identified from behavior in the first period when all consumers are attentive, and m_{ijt} is identified from deviations in all subsequent time periods from the behavior implied by redemption decisions in the first period. The utility of not purchasing u^{NP}_{ij} is simply identified from the share of purchasers for deal j.

Now, reintroduce unobserved preference heterogeneity. Purchase and first period redemption data for multiple deals with varying characteristics and multiple observations per consumer help identify unobserved persistent preferences for deal redemption. Namely, we observe each consumer purchase at least two deals, which helps identify consumer-specific preferences for deal characteristics. Given that our panels are fairly short, we restrict unobserved preference heterogeneity to the utility of redemption and the pricing coefficient, and estimate a model with two types of consumers. Another source of identification of unobserved heterogeneity is the assumption that certain deal characteristics enter only in the utility of redemption. This exclusion restriction implies that in a model with no unobserved preference heterogeneity, deals that are purchased the most will also be redeemed the most. The extent to which this prediction fails in the data will help identify the extent of unobserved preference heterogeneity.

2.3.6 Estimation

The likelihood of observing individual i of type m for deal j can be written as

$$l_{ij}(\theta_m, \eta) = \Pr\left(\{d_{ijt}\}_{t=1,\ldots,T} | \theta_m, \eta\right).$$

Dropping the ij subscripts for exposition, we can write the likelihood as a function of conditional probabilities,

$$\begin{aligned} l(\theta_m, \eta) = \Pr\left(\{d_t\}_{t=1,\ldots,T} | \theta_m, \eta\right) &= \int_{a_1,\ldots,a_T} \cdots \int \Pr\left(\{d_t, a_t\}_{t=1,\ldots,T} | \theta_m, \eta\right) da_1 \ldots da_T \\ &= \int_{a_1,\ldots,a_T} \cdots \int \Pr\left(d_1 | a_1\right) \prod_{t=2}^{T} \Pr\left(d_t | a_t\right) \Pr\left(a_t | a_{t-1}\right) da_1 \ldots da_T, \end{aligned}$$

where $\Pr(d_t|a_t)$ is the probability of making choice d_t conditional on attention state a_t, and $\Pr(a_t|a_{t-1})$ is the probability of being attentive in period t given the attention state in period $t-1$. To evaluate this likelihood expression we would need to execute a high-dimensional integral. We use the algorithm proposed by Reich (2014) to transform the problem of evaluating one T-dimensional integral into that of iteratively evaluating T one-dimensional integrals simultaneously with the expected value function and evaluate the likelihood $l(\theta_m, \eta)$ exactly for each consumer and deal. Integrating over consumer types and reintroducing the ij subscripts, the likelihood for the data can then be written as

$$L(\theta_1, \theta_2, \eta, \phi) = \prod_{i=1}^{N} \left(\phi \prod_{j=1}^{J} l_{ij}(\theta_1, \eta) + (1-\phi) \prod_{j=1}^{J} l_{ij}(\theta_2, \eta) \right).$$

We maximize this likelihood to obtain parameter estimates.

2.4 Merchant Pricing Model

We model each merchant as offering a single deal and use the subscript j for deal and merchant interchangeably.[10] A merchants sets a discount to maximize the expected discounted sum of future profits obtained from its deal. Platforms and merchants share revenue from deals sold. We assume that for each deal, the platform commits to offering 50% of its revenues to the merchant if the consumer redeems the deal. Otherwise, the platform retains 75% and pays the merchant 25% of revenue. The merchant gets paid whenever a consumer redeems the deal. After the deal expires, the merchant receives payment for all consumers who failed to redeem.

Merchants are likely to have dynamic incentives. Indeed, the main purpose of a daily deal is often to attract new consumers who may become regulars in the future (Edelman *et al.*, 2014). Merchants attribute a value ω_j to each redeeming consumer. The value ω_j may represent the merchant's estimate of the customer lifetime value of a new consumer acquired from the daily deal site (Gupta *et al.*, 2004). We assume that the customer lifetime value from a consumer who purchased but failed to redeem a deal is zero. Otherwise, all redeeming consumers have the same value to the merchant. Any marginal costs of serving a consumer are absorbed by the ω_j term.[11] As a result, ω_j may be negative if the merchant expects to obtain limited future business from consumers with daily deals.

Expected profits of merchant j are given by

$$\pi_j^{Merchant}(p_j) = \sum_{t=1}^{T_j} \tau^{t-1} R_{jt}(0.5p_j + \omega_j) + 0.25\tau^{T_j+1} F_j p_j,$$

where T_j is the duration of deal j, τ is the discount factor, R_{jt} is the total number of consumers who redeem the deal at time t, $p_j = \text{PRICE}_j$ is the price of deal j, ω_j is the merchant's perception of a new customer's lifetime value, and F_j is the total number of consumers who purchased but did not redeem the deal.

[10] The 39 deals in our data offer discounts for different products.

[11] The ω_j term may also capture add-on fees charged by the merchant at the time of redemption.

2.4.1 Estimation

We fix τ at 0.99 and obtain estimates of ω_j that rationalize merchants' pricing decisions. Merchant first-order-conditions imply that

$$\frac{d\pi_j^{Merchant}(p_j)}{dp_j} = \sum_{t=1}^{T_j} \tau^{t-1} \left[\frac{dR_{jt}}{dp_j}(0.5p_j + \omega_j) + 0.5R_{jt} \right] + 0.25\tau^{T_j+1} \left[\frac{dF_j}{dp_j} p_j + F_j \right] = 0,$$

at the price-levels chosen by the merchants. Solving for ω_j we obtain

$$\omega_j = \frac{0.5 \sum_{t=1}^{T_j} \tau^{t-1} \left[\frac{dR_{jt}}{dp_j} p_j + R_{jt} \right] + 0.25\tau^{T_j+1} \left[\frac{dF_j}{dp_j} p_j + F_j \right]}{\sum_{t=1}^{T_j} \tau^{t-1} \frac{dR_{jt}}{dp_j}}.$$

Substituting in estimates for R_{jt}, $\frac{dR_{jt}}{dp_j}$, F_j and $\frac{dF_j}{dp_j}$ implied by the consumer choice model, we find the ω_j values that satisfy the firs-order-conditions for each merchant.

2.5 Results

2.5.1 Demand Estimates

Table 2.3 shows the results of the estimation routine. We detect substantial heterogeneity in consumer *Redemption Utility Parameters*. Type 2 consumers appear more price sensitive (as suggested by the high estimate of the price-sensitivity parameter α_2^p). Estimates of the availability parameter α_i^a and the intercept α_i^0 suggest that type 2 consumers are also more time-constrained. Recall that the availability measure AVAILABLE_{jt} varies by week and measures the fraction of time that a deal is redeemable during that week. The fact that type 2 consumers have a higher availability parameter suggests that if a deal were redeemable for a smaller fraction of a given week, type 2 consumers would find it substantially more difficult to redeem than type 1 consumers. This evidence is consistent with our expectations of how the consumer population would behave. We would expect consumers to differ in their time commitments and hence in their costs of redeeming deals. The category-specific coefficients α_i^h and α_i^e show that consumers also differ in their preferences for different categories, with type 1 consumers exhibiting a substantial preference for restaurants and type 2 consumers largely indifferent between the three categories. We observe that type 1 consumers prefer deals located in the city center, whereas type 2 consumers are largely indifferent in terms of location. Also, type 2 consumers appear more heavily influenced by savings. Overall, the estimates may suggest that type 2 consumers are bargain hunters who purchase deals that are attractive because of their high savings, regardless of their location and category, even though these deals may be more difficult for them to redeem, whereas type 1 consumers prefer consuming deals for the experience, predominantly purchase restaurant deals, and enjoy a central location.

Parameters in the *Category-Specific Purchase Costs* section of Table 2.3 capture the costs of time and effort associated with completing the purchase forms and accessing the category-specific pages on the daily deal website. These estimates show that there are significant non-price costs associated with purchasing a deal across all categories.

The *Memory Activation Parameters* show that a consumer has a baseline probability of $e^{\gamma^0}/(1+e^{\gamma^0}) = 0.11$ of remembering about a purchased deal. This probability increases to $e^{\gamma^0+\gamma^a}/(1+e^{\gamma^0+\gamma^a}) = 0.91$ if the consumer was also attentive in the prior time period. Deadline effects increase the probability of remembering from 0.11 to 0.59 for an inattentive consumer, whereas a consumer who was attentive in the period before deadline effects is virtually guaranteed to remember about the deal. These estimates suggest that although a consumer who forgets about a purchased deal is very unlikely to remember about it again until it is near expiration, the extent of serial correlation in attention is very strong. Figure 2.3 plots the probability that a consumer will be attentive as a function of time for a duration of 10 weeks. This curve maps out the attention probabilities implied by the memory activation function. We can also examine the probability of being attentive for consumers who failed to redeem a deal. The model estimates imply that of all redemption failures, 28% can be attributed to inattention. In other words, of all consumers hadn't redeemed up to time T_j, 28% were inattentive at time T_j and failed to redeem as a result.

Demand Model Fit

Table 2.4 shows a comparison of a number of key sample moments with model predictions. We are able to closely reproduce key moments such as the total number of purchases and redemption failures, as well as the distribution of redemptions over time. Note that the predicted number of

Table 2.3: *Consumer Choice Model Parameter Estimates*

Redemption Utility Parameters			
		Type 1	Type 2
α_i^0	Intercept	-0.877	-2.870
		(0.328)	(0.836)
α_i^h	Health Category	-1.711	1.095^{NS}
		(0.356)	(0.589)
α_i^e	Entertainment Category	-2.599	-0.221^{NS}
		(0.328)	(0.499)
α_i^l	Location	1.160	0.174^{NS}
		(0.099)	(0.231)
α_i^s	Savings	0.413^{NS}	1.363
		(0.349)	(0.564)
α_i^a	Availability	1.598	2.451
		(0.174)	(0.452)
α_i^p	Price	0.021	0.037
		(0.004)	(0.005)
Category-Specific Purchase Costs			
β^r	Intercept	5.509	
		(0.169)	
β^h	Health Category	0.297^{NS}	
		(0.279)	
β^e	Entertainment Category	-0.830	
		(0.196)	
Memory Activation Parameters			
γ^0	Intercept	-2.125	
		(0.181)	
γ^a	Past Attention	4.428	
		(0.323)	
ξ	Deadline Effects	2.492	
		(0.214)	
ϕ	Share of Type 1 Consumers	0.69	
Log-Likelihood		-20850	

Note: Coefficients that are not significant at the 0.05*-level are marked with an* NS.

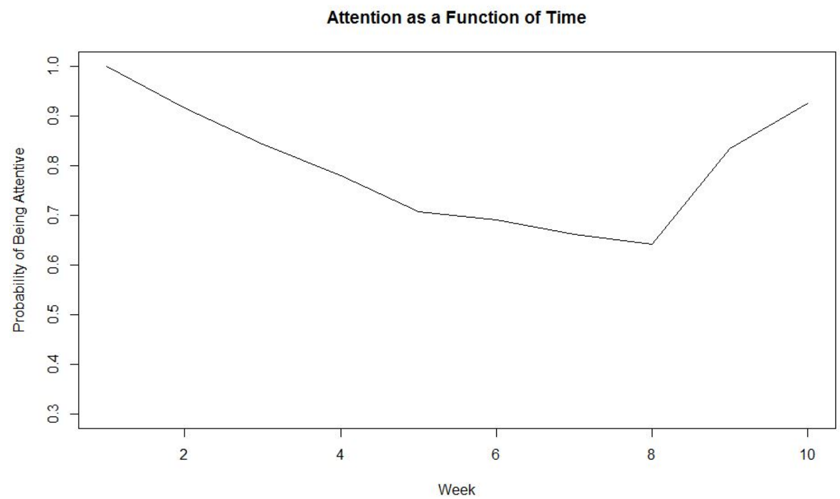

Figure 2.3: *Probability of Being Attentive*

Table 2.4: *Comparison of Sample Moments and Demand Model Predictions*

Moments	Purchases	Redemption Failures	Percentage of Redemption			
			1^{st} Quartile	2^{nd} Quartile	3^{rd} Quartile	4^{th} Quartile
Actual	3341	9%	33%	25%	16%	26%
Predicted	3338	11%	34%	25%	17%	24%

Note: We divide deal durations into quartiles and record the percentage of total redemptions that occur within each quartile.

redemption failures depends on the predicted number of purchases. As a result, any error in the purchase prediction will make it more difficult to predict redemption failure. We are satisfied to find that our model predicts both purchases and redemption failures well, yielding a correlation 0.70 between actual and predicted purchases and 0.65 between actual and predicted redemption failures across deals. Figure 2.4 further illustrates the ability of the model to rationalize redemption data by presenting the aggregate actual and fitted redemption curves.

2.5.2 Supply-Side Estimates

Figure 2.5 shows a histogram of ω_j parameters implied by merchant pricing decisions. Overall, merchants assign a positive value to consumers obtained through the daily deal site. Only one merchant, the one that charges the highest price, perceives consumers as a cost. In general, merchants who charge a higher price have a lower perception of customer lifetime value. These

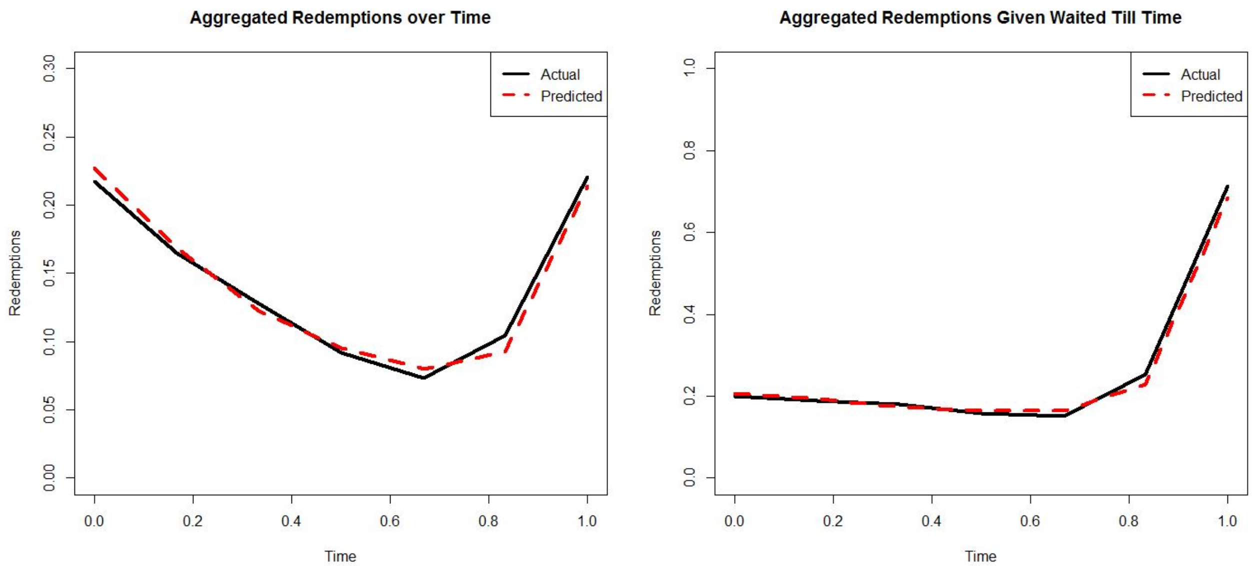

Figure 2.4: *Actual and Fitted Redemption Curves*

Note: Duration of all deals is normalized to 1. Availability of deals in each week is normalized to 1.

Table 2.5: *Merchant Perceptions of Customer Lifetime Value Across Categories*

Category	Mean	Standard Error
Restaurant	$111.15	$5.23
Health	$70.10	$7.42
Entertainment	$77.86	$17.88

merchants appear to be charging a higher price upfront to extract profit from the deal, rather than from the future visits from newly acquired consumers.

Table 2.5 shows the estimated mean customer lifetime value across categories. Restaurant merchants view customers as much more valuable than merchants in the health or entertainment categories, consistent with the popularity of daily deal sites for restaurants. It is important to note that ω_j captures the perceived customer value from the perspective of the merchant. The actual realization may be quite different. Indeed, merchants may be overly optimistic with regards to how much future value they can extract from customers acquired through daily deal sites.[12] The data we observe originate from the time period when daily deal popularity was at its peak, consistent with the high customer lifetime value estimates we recover.

[12] Absent data on future consumer visits and purchases we cannot estimate the extent of merchant optimism.

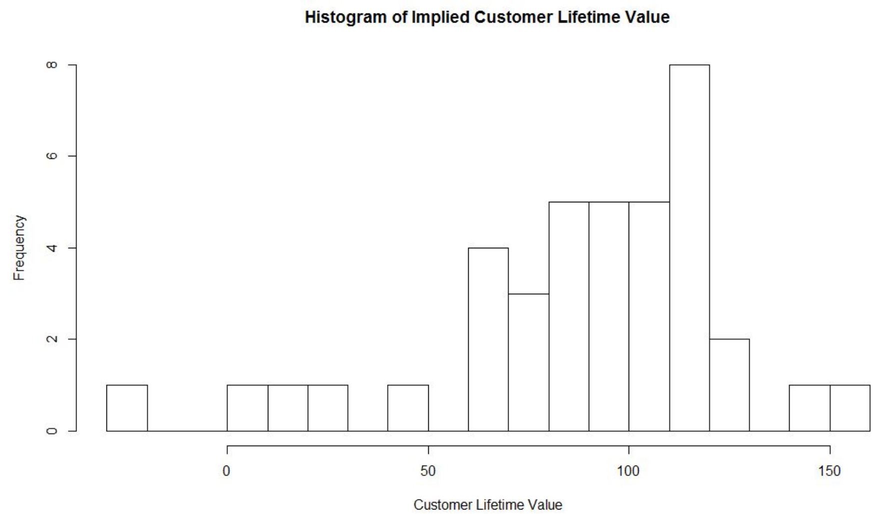

Figure 2.5: *Histogram of Merchant Perceptions of Customer Lifetime Value*

2.6 Counterfactuals

We conduct a series of counterfactual simulations to examine consumer behavior and how the design of a daily deal marketplace affects consumers, merchants, and the platform. Merchants react to counterfactual expiration and pricing policies by changing the discounts they offer according the the merchant first-order-conditions presented in Section 2.4. Consumers then make purchase and redemption decisions according to the model in Section 2.3. We assume that the platform has zero marginal costs. The platform obtains profits

$$\pi^{\text{Platform}} = \sum_{j=1}^{J} (0.5R_j + 0.75F_j) p_j$$

where $R_j = \sum_{t=1}^{T_j} R_{jt}$ is the sum of all redemptions obtained by merchant j.[13] Merchant profits reflect the expected discounted sum of profits obtained by each merchant throughout the duration of its deal, taking into account its perceptions of the value of daily deal-wielding consumers.

Table 2.6 summarizes key insights from the counterfactual policies. For each policy we compute the expected total number of deals sold, redemption failures, percentage of redemption failures attributable to inattention, profits for the merchants and the platform, and the minimum, mean,

[13]We assume that the platform does not discount future profits. Qualitatively similar results obtain if we allow for discount factors in the range of 0.9-1.

Table 2.6: *Counterfactual Policies*

Counterfactual Policy	Original	Attention	Duration	Commission	Prepayment
Deals Sold	3,338	4,865	4,127	3,477	3,455
Redemption Failures	358	206	383	375	448
Redemption Failures (% of sold)	11%	4%	9%	11%	13%
Forgotten Deals (% of redemption failures)	29%	0%	18%	30%	24%
Platform Profits	$21,336	$32,741	$27,276	$19,692	$19,140
Merchant Profits	$64,552	$100,113	$79,743	$62,875	$63,922
Minimum Savings (%)	54%	51%	54%	54%	50%
Mean Savings (%)	76%	75%	76%	82%	76%
Maximum Savings (%)	97%	96%	97%	99%	99%

Note: For the Duration counterfactual, we extend deal expiration dates by four weeks.

and maximum savings offered. The first column summarizes the predicted impact of the original policy of the platform as a benchmark.[14]

2.6.1 Attention

Had all consumers been perfectly attentive, how different would their purchase and redemption patterns be, and what are the implications for merchants and the platform? Perfect attention has a positive effect on the daily deal marketplace and the merchants. More consumers purchase, as the expected value of holding a deal increases, and fewer fail to redeem as inattention no longer causes redemption failure. Merchants can charge higher prices to take advantage of the increased value of holding a deal to consumers. Prices increase slightly, resulting in an increase in merchant profits. The reduction in redemption failure rate may cause a decline in platform profits, as it earns a greater commission on deals that go unredeemed. However, the large increase in sales outweighs the negative effects of an increase redemption rate. Platform profits increase.

2.6.2 Duration

We simulate a scenario where all deal durations are extended by 4 weeks. We observe an increase in purchases but no substantial change in redemption failures. The impact of extended deadlines is largely similar to the case of reduced inattention. Merchants do not change prices substantially, and the increase in sales drives up merchant and platform profits.

[14]We compare countefactual simulations to the simulated outcome under the original policy.

2.6.3 Commission

The original revenue sharing agreement requires that merchants receive 25% of the price for each consumer who fails to redeem a deal. We explore the impact of a more stringent policy, where the platform does not pay the merchant anything for unredeemed deals. In this case, platform profits become

$$\pi^{\text{Platform}} = \sum_{j=1}^{J} (0.5R_j + F_j) p_j,$$

such that the platform retains 100% of the price of an unredeemed deal. Merchants receive

$$\pi_j^{Merchant}(p_j) = \sum_{t=1}^{T_j} \tau^{t-1} R_{jt}(0.5p_j + \omega_j).$$

Interestingly, both merchant and platform profits fall under such a revenue sharing structure. Merchants no longer receive payment for unredeemed deals. They compensate by offering greater discounts to increase sales and redemptions. Sales do not increase sufficiently and merchant profits fall. The added gains to the platform from its increased retention of revenues for unredeemed deals do not suffice to cover the losses from lower prices charged by the merchants. Platform profits decrease. By choosing to retain a greater share of revenue, the platform actually induces merchants to lower prices and reduces its own profits.

2.6.4 Prepayment

A daily deal marketplace may consider transitioning to an online coupon marketplace business model. In such a scenario, consumers can download the deal for free, and the platform charges consumers only when they redeem the deal. Per-period utilities are adjusted as follows:

$$u_{ijt}^{R} = \alpha_i^r + \alpha_i^h \text{HEALTH}_j + \alpha_i^e \text{ENTERTAINMENT}_j + \alpha_i^l \text{LOCATION}_j + \\ \alpha_i^s \text{SAVINGS}_j + \alpha_i^a \text{AVAILABLE}_{jt} - \alpha_i^p \text{PRICE}_j,$$

$$u_{ijt}^{W} = 0,$$

so that the price sensitivity term only enters the utility of redeeming. This policy would require an alternative model for revenues, given that the daily deal marketplace will no longer have access to cash from expired deals to share with merchants. Platform profits can be written as

$$\pi^{\text{Platform}} = \sum_{j=1}^{J} 0.5 R_j p_j,$$

In other words, the platform shares 50% of revenues with merchants for deals that were redeemed, but obtains no payment from consumers with expired deals. Similarly, merchant profits can be written as

$$\pi_j^{Merchant}(p_j) = \sum_{t=1}^{T_j} \tau^{t-1} R_{jt}(0.5p_j + \omega_j).$$

We find that under this policy consumers purchase (download) slightly more (as purchase is now costless) and redeem less as the price of redemption increases. Merchants may increase or decrease prices. Prices may increase because as more consumers download deals, they have more opportunities to pay and redeem. Merchants can take advantage of this by increasing prices. However, if the merchant's perception of customer value is high and deal duration is short, it may lower price to get more consumers to redeem as it only derives a high profit from redeemers and convincing a consumer to redeem is more challenging for a short deal with a no prepayment policy. Overall, both platform and merchant profits fall, largely because redemptions do not increase substantially and neither the merchant nor the platform obtain any revenue when a redemption failure occurs.

2.6.5 Discussion and Implications

We find that increasing consumer attentiveness, perhaps through reminders, increases the value of the deal to consumers, which in turn allows merchants to reduce their savings offerings, increasing profits for both the merchant and the platform. Merchants and the platform no longer obtain as much revenue from consumers who don't redeem, but as non-redeemers form only a small of total sales, profits increase. We expect these findings to be sensitive to the assumption that consumers know their probability of falling inattentive. If consumers over-estimate their ability to remember, sales would not increase by as much in the presence of reminders, and merchants would not react as strongly by decreasing discounts. Nevertheless, as long as consumers recognize the possibility of forgetting, reminders will increase the value of holding a deal, leading to greater sales and greater merchant and platform profits.

Extended expiration dates achieve a similar objective. If a deal is available for longer, consumers value it more, which leads to greater sales and profits for both the platform and the merchant. It is possible that if consumers over-estimate their attentiveness, longer deal durations yield lower redemption rates than shorter deal durations. In such a scenario, merchants who assign a high value to daily deal consumers may lose profit, as they derive no value from non-redeemers. Merchants who assign a low value to daily deal consumers would gain profit, as they benefit from the increase in non-redemption. Overall, the impact on merchants would be ambiguous. As long as consumers do not over-estimate their attention probability so as to generate the implication that longer expiration periods yield fewer redemptions, our findings that both merchants and the platform benefit from extended expiration dates hold true.[15]

We find that by adjusting the revenue-sharing structure in its favor, the platform may actually decrease its profits, as well as merchant profits. Merchants respond to the new structure by offering lower discounts, which reduces profits for both merchants and the platform. A key assumption underlying the counterfactuals is that deal characteristics other than discount remain unchanged under different policies. For example, we may expect merchants to increase the quality of their offering if the platform restricts merchant earnings from non-redeemers. An increase in deal quality would offset the merchant's incentives to offer greater discounts and may potentially increase the profits of the platform. We expect our results to hold in markets where changing the quality of a product is non-trivial. Moreover, in most instances it is less costly to change price than it is to adjust the quality of the product.

We find that holding fixed deal characteristics, prepayment affects different merchants differently. Namely, merchants who attribute a high value to consumers increase their discounts to encourage redemption. Merchants who attribute a low value to consumers and benefit the most from non-redemption do not offer higher discounts and struggle to generate sales and experience a loss. Our findings suggest that, relative to daily deal platforms, coupon exchanges tend to attract merchants who assign a higher value to new consumers. Moreover, transitioning to a no-prepayment model would reduce the profits of a daily deal platform. When consumers no longer have to prepay, merchants may adjust product quality to encourage deal redemption and

[15]Typically, the extent of overconfidence must be substantial to rationalize an inverse relationship between deal duration and total redemptions (Taubinsky, 2013).

payment. Overall, this would increase merchant costs. Platform profits would not increase absent a reduction in discount offerings, suggesting that both merchant and platform profits would decrease even if merchants could adjust variables other than price.

2.7 Conclusions

We study consumer behavior and the effects of different expiration and pricing policies in an Asian daily deal marketplace. Our dynamic structural model helps explain the U-shaped redemption patterns observed in the data and provides estimates of the extent of preference heterogeneity and inattention in the sample. Counterfactuals highlight merchant and platform incentives, and shed light on alternative business models that may prove more attractive to the daily deal marketplace. Our results inform managers of the impact of consumer inattention and the potential advantages of different expiration and pricing policies for discount vouchers and other products with expiration dates and prepayment.

We make several simplifying assumptions to ensure the model remains estimable given our data. We treat deals as independent and do not model complementarity and substitution across deals. Complementarity is an important but challenging issue that arises in many empirical contexts (Berry *et al.*, 2014). We study purchase and redemption behavior for each deal in isolation for simplicity but posit that future research may develop a framework to incorporate purchase and redemption decisions for a portfolio of deals. Second, we do not model the negotiation process between the platform and the merchant. Although we believe that merchants have the final say in how great a discount to offer, the platform may influence merchant pricing decisions through sales and marketing efforts unobserved in our data. Third, we do not model biases in consumer perceptions of their ability to remember deals. Additional data on consumer expectations or calendars would help separately identify the actual inattention process and consumer perceptions of their inattention.

Similar consumer behavior may arise in relation to many other marketing phenomena, such as gift cards, loyalty points, coupons, and rebate programs. Our results suggest that decision-makers in the daily deals industry and other similar contexts must carefully consider the implications of different expiration and pricing policies.

Chapter 3

Match Your Own Price? Self-Matching as a Retailer's Multichannel Pricing Strategy[1]

3.1 Introduction

Many, if not most, major retailers today employ a multichannel business model - they offer products in physical stores and online. These channels tend to attract different consumer segments and allow retailers to cater to different buying behaviors and preferences. Consumers are also becoming more savvy in utilizing the various channels during the buying process: researching products, evaluating fit, comparing prices and purchasing (Neslin *et al.*, 2006; Grewal *et al.*, 2010; Verhoef *et al.*, 2015).

Retailers need to attend to all elements of the marketing mix as they strive to maximize profits. Not surprisingly, pricing has always been an important strategic variable for them to "get right." When retailers were predominantly bricks-and-mortar, this meant figuring out the most effective store price to set for their merchandise. However, having embraced a multichannel selling format, pricing decisions have become much more complex for these retailers to navigate. Not only do they need to price the products in their physical stores, they need to set prices for products in

[1]Co-authored with Vineet Kumar and Elie Ofek.

their online outlets *and* consider how the prices across the various channels should relate to one another. This complexity in devising a comprehensive multichannel pricing strategy is front and center for retailers today, as evidenced by the myriad of commentaries in the retail trade press. As Forrester Research reports (Mulpuru, 2012), "*...Given that the majority of retail sales are expected to be influenced by the Web to some degree, it becomes imperative for eBusiness professionals in retail to adopt cross-channel best practices, particularly related to pricing*."

Formulating an effective multichannel pricing strategy can be challenging. A recent survey of leading retailers (Retail Systems Research, 2013) revealed that their top two pricing challenges are: (1) increased price sensitivity of consumers, and (2) pricing aggressiveness from competitors. In a world where many consumers either buy online or conduct research online before entering a store, these findings suggest that the need to manage the heightened price sensitivity and combat intense competition are becoming even more important. We specifically investigate how retailers can leverage self-matching to diminish intense price competition while taking customer behavior into account. Interestingly, the above study did not find the item "*need to provide consistency in price across channels*" to be among the top challenges these retailers face, underscoring that they feel they have flexibility in customizing their price to the specific channel and customer mix that chooses to shop there. Indeed, according to Gartner's Kevin Sterneckert (Grau, 2012), "*Using a single-channel, consistent pricing strategy misses important opportunities in the marketplace...*"

With a self-matching policy, the retailer commits to charging consumers the lower of its online and store prices for the same product when consumers furnish appropriate evidence of a price difference. Note that even though self-matching can provide some degree of price consistency, it is fundamentally different from committing to consistent prices and setting exactly the same prices across all channels. Commonly, this policy features store to online self-matching, allowing consumers to pay typically lower online prices for store purchases.[2] This policy is a novel marketing instrument that is uniquely available to multichannel retailers and not relevant in the single-channel case. The primary objective of our paper is to understand the strategic and profitability consequences of self-matching policies. To our knowledge, the implications of

[2]A webpage printout or a mobile screenshot of the webpage usually suffices as appropriate evidence. Policies allowing self-matching in the other direction, i.e., allowing web customers to match store prices, are rarely observed in practice as prices online are typically lower than in-store (Grau, 2012; Mulpuru, 2013).

this important and growing issue of self-matching have not been systematically evaluated in the literature, in contrast to competitive price-matching policies that have been studied extensively (and reviewed in Section 3.2).

Examining a number of retail markets, there are two distinct self-matching patterns that come to our attention. First, we observe considerable heterogeneity in the adoption of self-matching policies *across retailers* in the US, including those competing for the same market. For example, Best Buy, Target, Walmart, Sears, Staples, Office Depot, Sports Authority, Toys"R"US, and Lowe's price-match their online channels in-store, whereas JCPenney, Macy's, Bloomingdale's, Urban Outfitters, and Petco explicitly state that they will not match their prices across channels.[3] Second, we also observe heterogeneity in self-matching *across industries*. In consumer electronics, discount retailing and office supplies, major players choose to offer self-matching, whereas in mid-upscale department stores and clothing, most or all retailers tend to not adopt self-matching. There is no theoretical understanding of when and why these patterns result.

We aim to develop insights on when to expect different patterns of self-matching decisions for multichannel retailers in a given category: all self-match, some self-match while others don't, and none self-match. To this end, we examine the strategic use of a self-matching pricing policy by multichannel retailers across a variety of competitive settings: a monopoly; a duopoly with two competing multichannel retailers; and a mixed duopoly in which a multichannel retailer competes with an e-tailer. More specifically, we address the following research questions:

1. What strategic mechanisms underpin the decision to implement a self-matching policy?
2. When do multichannel retailers choose to self-match in equilibrium? How do customer and product characteristics, and the nature of competition, influence a retailer to self-match?
3. How does self-matching affect the prices charged online and in-store?
4. Are retailers always better or worse off having access to self-matching as a strategic tool?

To investigate these questions, we develop a model that allows us to capture the effects of self-matching on consumer and retailer decisions. We allow for consumer heterogeneity along

[3]See Appendix B.7 for examples of current self-matching policies from retailer websites.

a number of important dimensions. These dimensions include consumers' channel preferences, their stage in the decision-making process, and their preference across retailers. With regard to channel preferences, we allow for "store-only" consumers who have a strong preference to purchase in-store where they can "touch and feel" merchandise and obtain the product instantly. In contrast, "channel-agnostic consumers" don't have a strong preference for which channel they purchase from. We also distinguish between consumers who know the exact product they want to purchase ("Decided") and consumers who only recognize the need to purchase from a category and require a visit to a store to shop around and find the specific version or model that best fits their needs ("Undecided"). Finally, consumers have horizontal taste (or brand) preferences across retailers.

The retailers offer unique products of similar value and are located at the ends of a Hotelling linear city, with consumer location on the line indicating retailer preference. Retailers first choose a self-matching pricing policy and subsequently set price levels for both store and online channels simultaneously. We analyze the subgame perfect equilibria of the game.

Our analysis reveals several underlying mechanisms that affect the profitability of self-matching in equilibrium. The combination of these effects and the tradeoffs between them are novel and, to the best of our knowledge, have not been previously identified. The first effect, termed *channel arbitrage*, is negative and reduces profits, whereas the other effects of *decision-stage discrimination* and *online competition dampening* increase profits. Thus, the overall profit implications of implementing a self-matching pricing policy depend on the existence and magnitude of these effects.

Consider the pricing incentives faced by a multichannel retailer. In the store channel, retailers face two types of consumers - those who research the product online before choosing a store (decided consumers) and those who visit their preferred store first to identify the product (undecided consumers). The retailers would like to charge a higher price to the latter type, but are unable to do so because both types purchase in the store channel. Furthermore, consumers who shop online tend to be informed of the online prices at both retailers, which leads to more competitive pricing in the online channel than in-store.

Now, consider the strategic impact of self-matching policies. With self-matching, consumers who research the product online but purchase in-store redeem the lower online price, and we refer to this profit-reducing effect of self-matching as *channel arbitrage*. However, consumers who visit their preferred store first without searching online are unable to obtain evidence of a lower price for their preferred product before arriving at the store. These consumers pay the store price even when a self-matching policy is in effect. We term this the *decision-stage discrimination effect*, where a self-matching policy allows the retailer to charge different prices to store consumers based on their decision-stage. This mechanism involves a unique type of price discrimination that has not received attention.

If only one retailer self-matches, decided consumers can redeem the lower online price at the self-matching retailer's store but only have access to the higher price at the rival's store. This leads the self-matching retailer to set a higher online price to mitigate the negative impact of channel arbitrage. The rival follows suit due to strategic complementarity of prices and sets a higher online price as well, thus softening online competition. We refer to this profit-increasing mechanism as the *online competition dampening effect*, which serves as a commitment device to prevent online prices from being lowered to the competitive level. It occurs only when one of the retailers offers to self-match. If both retailers self-match, consumers have access to the online prices at both stores, and intense competition in the online channel ensues.

We further investigate the profitability of the self-matching policy. Our analysis helps show that self-matching is not necessarily harmful. In fact, *both retailers can be better off by offering to self-match* when the positive online competition dampening and decision-stage discrimination effects dominate the negative channel arbitrage effect.

We investigate several relevant model extensions. First, we examine how the presence of consumers equipped with "smart" devices, who can retrieve online price information when in the store, affects retailers' incentives to implement a self-matching policy. Intuitively, when more consumers are able to retrieve the lower online price, the negative channel arbitrage effect is more pronounced. However, we find that the presence of "smart" consumers can allow retailers to benefit even more from online competition dampening by charging higher online prices. In another extension, we examine the effects of self-matching in a mixed duopoly, in which a multichannel retailer competes with an online-only retailer (or pure e-tailer). We find that competition is dampened in the online channel when the multichannel retailer chooses to self-match, allowing both retailers to benefit.

We review the literature (Section 3.2), describe the model (Section 3.3), and proceed to analyze equilibrium strategies and outcomes (Section 3.4). We then consider several extensions of the model (Section 3.5) and finally conclude by discussing managerial and empirical implications as well as future research (Section 3.6).

3.2 Literature Review

We draw from two separate streams of past research. The first is focused on multichannel retailing, and the second, on competitive price-matching in a single channel. Research in multichannel retailing has typically assumed that retailers either set the same or different prices across channels, without examining the incentives to adopt a self-matching policy. Liu *et al.* (2006), for example, study a bricks-and-mortar retailer's decision to open an online arm, assuming either price consistency or different prices across channels. Zhang (2009) considers separate prices per channel and studies the retailer's decision to operate an online arm and advertise store prices. Ofek *et al.* (2011) study retailers' incentives to offer store sales assistance when also operating an online channel, allowing for either identical or different pricing across channels. Aside from ignoring self-matching pricing policies, this literature has not considered or modeled heterogeneity in consumers' decision-making processes, which plays an important role in their channel choice in practice.

A developing literature examines the effects of free-riding (Shin, 2007) or show-rooming on retailer competition (Mehra *et al.*, 2013), including suggestions that store to online price-matching may serve as a strategy to combat show-rooming. However, there has not been a careful modeling and evaluation of whether and when such policies can be effective; particularly in a competitive context.

Competitive price-matching is an area that has been well-studied. This literature has generally focused on retailers' incentives to match the prices of their competitors in a single channel setting, typically bricks-and-mortar stores. Salop (1986) argued that when retailers price-match each other, this leads to higher prices than otherwise, as they no longer have an incentive to engage in price competition; implying a form of tacit collusion (Zhang, 1995). However, competitive price-matching has also been found to intensify competition because it encourages consumer search (Chen *et al.*, 2001). Other research in competitive price-matching has explored the impact of hassle costs (Hviid and Shaffer, 1999), its role as a signaling mechanism for certain aspects of a retailer's product or service (Moorthy and Winter, 2006; Moorthy and Zhang, 2006), the interaction with product assortment decisions (Coughlan and Shaffer, 2009), and the impact of product availability (Nalca *et al.*, 2010).

By contrast, *self-matching* pricing policies represent a phenomenon only relevant for multi-channel retailers, and recent retailing trends make self-matching an important issue to study. First, the nature of competition is evolving in many categories, from retailers carrying the same products from multiple brands to manufacturers that establish their own retail stores e.g., Apple, Microsoft and Samsung. Second, many retailers are moving towards establishing strong private label brands or building exclusive product lines to avoid price wars with competitors (Bustillo and Lawton, 2009; Mattioli, 2011). For instance, about 56% of the products sold by health supply retailer GNC are exclusive or GNC-branded, and electronics retailers like Brookstone and Best Buy are also increasingly focusing on private label products. These trends accentuate the relevance of self-matching relative to competitive price-matching as the product assortments retailers carry become more differentiated.

The mechanisms underlying self-matching are also connected to the broad literature on price discrimination. Cooper (1986) examines "most-favored-customer" pricing in a two-period model, where retailers commit to giving consumers who purchase in the first period the difference between the first and second period prices if the latter price is lower (a form of inter-temporal self-matching). The author shows that this policy may increase retailer profits as it reduces the incentive to lower prices in the second period for both retailers. This effect shares similarities with the online competition dampening effect we identify, whereby a retailer reduces its own incentive to price lower online by inducing channel arbitrage through self-matching. However, whereas both retailers can offer and benefit from a "most-favored-customer" policy in the inter-temporal setting, the online competition dampening effect can only exist if one retailer offers to self-match. If both retailers self-match, they re-ignite competition in the online market for decided consumers and nullify the effect. Cross-channel price-matching is thus driven by different strategic incentives.

Thisse and Vives (1988); Holmes (1989) and Corts (1998) consider cases when price discrimination may lead to lower profits for competing retailers in equilibrium. In a similar way, retailers may be compelled to self-match in equilibrium even though they would have been better off had self-matching not been an option. In our context, on the one hand, a self-matching policy acts as a commitment not to price discriminate decided consumers across channels, which can potentially lead to greater profits for both retailers because this creates an incentive to increase the online price to mitigate the effect. On the other hand, a self-matching policy enables price discrimination between undecided and decided consumers who shop in-store. Depending on the relative sizes of these segments, self-matching policies may emerge in equilibrium and lead to greater or lower profits for both retailers.

Desai and Purohit (2004) consider a competitive setting where consumers may haggle over price with retailers. Some form of haggling may occur in the self-matching setting if retailers are not explicit about their self-matching policies and consumers must wrangle with managers to obtain a self-match. This interaction may induce additional costs for consumers and for retailers when processing a self-matching policy. In our analysis, we focus on the case when retailers explicitly announce their self-matching policies and illustrate how self-matching may emerge in equilibrium in the absence of consumer haggling. In an extension, we consider the implications of a retailer processing cost when a consumer redeems a self-match.

3.3 Model

3.3.1 Retailers

Two competing retailers in the same category are situated at the endpoints of a unit consumer interval, or linear city (Hotelling, 1929). The retailers offer unique and non-overlapping sets of products. Since the products carried by retailers are different, they do not have the option of offering competitive price-matching guarantees. For example, both Gap and Aeropostale sell apparel and operate in the same categories, but the items themselves are not the same and reflect the dedicated designs and logos of each of these retailers.

We model a two-stage game in which the retailers must first decide on self-matching policies and then on prices in each channel. We denote by $SM_i = 0$ the decision of retailer i not to self-match and by $SM_i = 1$ the decision to self-match, leading to four possible subgames - $(0,0), (1,1), (1,0)$ and $(0,1)$ representing (SM_1, SM_2). In each subgame, p_j^k denotes the price set by retailer $j \in \{1,2\}$ in channel $k \in \{on, s\}$, where on stands for the online or internet channel and

s stands for the physical store channel. With self-matching, consumers who retrieve the match pay the lowest of the two channel prices. In the equilibrium analysis that follows, we find that retailers never set lower prices in-store than online. Hence, the only relevant matching policy to focus on is the store-to-online self-match. All retailer costs are assumed to be zero.

3.3.2 Consumers

In order to capture important features of the consumer shopping process in multichannel environments, we model consumers as being heterogeneous along multiple dimensions.

Retailer Brand Preferences : Consumers vary in their preferences for a retailer's product, e.g., a consumer might prefer Macy's clothing lines to those offered at JCPenney's. This aspect of heterogeneity is captured by allowing consumers to be distributed uniformly along a unit line segment in the preference space, $x \sim U[0, 1]$. A consumer at preference location x incurs a "misfit cost" θx when purchasing from retailer 1 and a cost $\theta(1 - x)$ when purchasing from retailer 2. Note that the parameter θ does *not* involve transportation costs, rather it represents horizontal retailer-consumer "misfit" costs, which are the same across the online and store channels. Misfit costs reflect heterogeneity in taste over differentiated products of similar value, e.g., the collection of suits at Banana Republic compared to those at J.Crew.

Channel Preferences : *Channel-agnostic* (A) consumers do not have an inherent preference for either channel and, for a given retailer, would purchase from whichever channel has the lower price for the product they want. On the other hand, *Store-only* (S) consumers find the online channel insufficient, e.g., due to waiting times for online purchase, risks associated with online purchases (such as product defects), etc. These consumers purchase only in the store, although they might research products online and obtain online price information after deciding on the specific product they want to purchase. We assume that channel-agnostic consumers form a fraction η of the market while store-only consumers form a fraction $1 - \eta$.

Decision Stage : Consumers differ in their decision-making process (DMP) stage, a particularly important aspect of multichannel shopping (Neslin *et al.*, 2006; Mulpuru, 2010; Mohammed, 2013). *Undecided* (U) consumers (of proportion β, where $0 < \beta < 1$) need to undertake a shopping trip to the store because they don't have a clear idea of the product they wish to purchase. *Decided* (D) consumers (of proportion $1 - \beta$) are certain about the product they wish to buy and can thus costlessly search for price information from home. Undecided consumers first visit a retailer's store, selecting the store closest to their preference location to discover an appropriate product that fits their needs. After determining fit, they may purchase the product in-store or return home to purchase online (from either retailer), depending on their channel preference. Consumers obtain a value v from purchasing their preferred product.

Categories like apparel, fashion, furniture and sporting goods are likely to feature more undecided consumers, as styles and sizes of products are important factors that change frequently. Since undecided consumers do not know which product they want before visiting a store, they do not have at their disposal all prices while at the store, since keeping track of a large number of products, models and versions even within a category would be impractical. Undecided consumers in the model do not infer how self-matching influences retailer prices before they visit a store. They are unaware of the exact product they wish to purchase beforehand, thus limiting their ability to infer prices under different self-matching configurations. More generally, the product category is sufficiently large and varied to make forming accurate expectations of prices more costly than simply visiting the preferred store.

We set the travel cost for a consumer's first shopping trip to be zero and assume that additional trips are sufficiently costly. Note that if consumers have no cost to visit multiple stores in person, then we obtain a trivial specification where there is no distinction between decided and undecided consumers who shop in-store. Throughout the paper, we focus on the more interesting case where additional shopping trips are costly enough so that SU consumers do not shop across multiple physical stores (see Appendix for formal conditions on the travel cost). However, in all cases, decided consumers (both SD and AD) research products and prices in advance.[4]

Table 3.1 depicts the different consumer segments included in the model. We denote the four segments of consumers as SU, AU, SD, and AD, depending on their channel preference and

[4]In the baseline model, consumers cannot access online prices after arriving in the store, although we examine this possibility in Section 3.5, by modeling a segment of consumers with mobile internet access.

Table 3.1: *Consumer Segments and Proportions*

Channel Preference \ Decision Stage	Undecided (U)	Decided (D)
Store (S)	(SU): $(1-\eta)\beta$	(SD): $(1-\eta)(1-\beta)$
Agnostic (A)	(AU): $\eta\beta$	(AD): $\eta(1-\beta)$

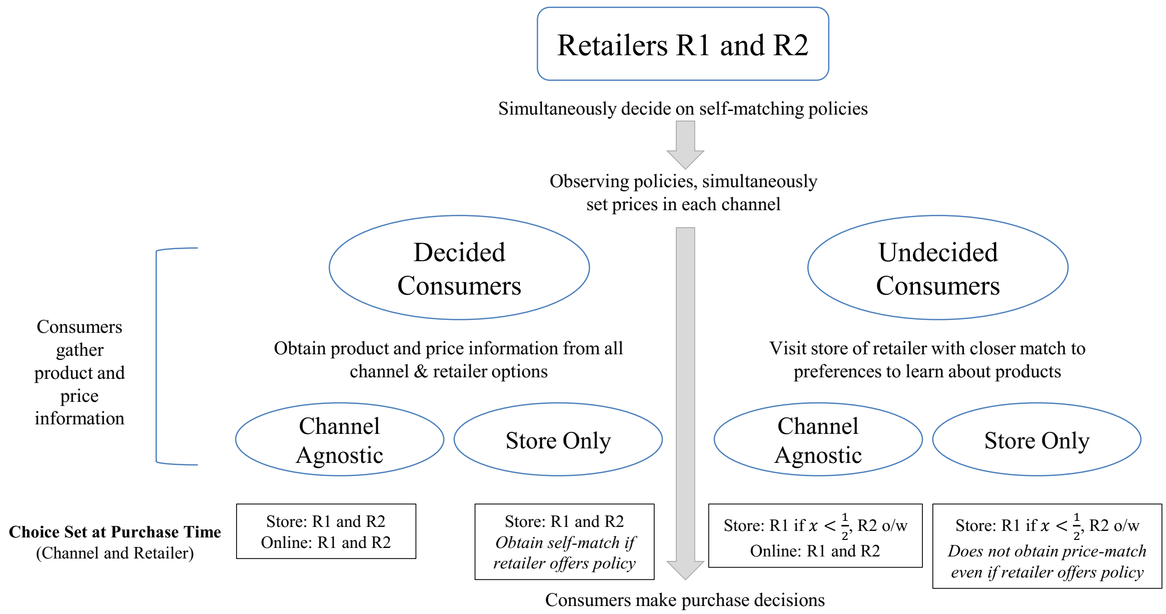

Figure 3.1: *Sequence of Events in Model*

decision-stage; the size of each segment is indicated in the corresponding cell of the table. Each of the four segments is uniformly distributed on a Hotelling linear city of unit length, such that consumer location on the line determines retailer preference. In Appendix B.2, we present an example to illustrate the decision process of a consumer from each segment. Table 3.2 details the notation used throughout the paper.

Sequence of Events

Figure 3.1 details the sequence of events. First, retailers simultaneously decide upon a self-matching pricing strategy. Then, after observing each other's self-matching decisions, they determine the price levels in each channel. Consumers, depending on their type (decided or

Table 3.2: *Summary of Notation*

Notation	Definition
p_j^s	In-store price for retailer j
p_j^{on}	Online price for retailer j
SM_j	Retailer j's self-matching decision
$\Pi_j^{SM_i,SM_j}$	Retailer j's total profit in subgame (SM_i, SM_j)
Consumer	
v	Consumer valuation of the product
θ	Retailer differentiation
$1-\beta$	Fraction of decided segment of consumers
β	Fraction of undecided segment of consumers
η	Fraction of channel-agnostic consumers
$1-\eta$	Fraction of store-only consumers
$\theta \cdot x$ for $x \in [0,1]$	Measure of retailer preference for consumer at location x
u_j^k	Utility for purchasing from retailer j in channel k

undecided, store or channel-agnostic, and horizontal preference), decide on which channel and retailer to shop at. Decided consumers, who know the online price prior to visiting the store, can ask for a price-match if the online price is lower *and* the retailer has chosen to self-match. Finally, consumers make purchase decisions and retailer profits are realized.

3.3.3 Consumer Utility

We now specify the utility consumers derive under different self-matching scenarios. Recall that decided consumers know all prices across both retailers and channels before they make a purchase decision. Channel-agnostic undecided consumers (AU) have the option of visiting a store to learn what they want and then returning home to make an online purchase, whereas store undecided consumers (SU) either purchase in the store they first visit or make no purchase. Consumers obtain zero utility when they don't make a purchase.

Consider the case when neither retailer self-matches, i.e., $(SM_1, SM_2) = (0,0)$. For a consumer who knows the product she wishes to purchase, the utility for each retailer and channel option is as follows:

$$u_1^{on} = v - p_1^{on} - \theta x, \qquad u_1^s = v - p_1^s - \theta x, \tag{3.1}$$
$$u_2^{on} = v - p_2^{on} - \theta(1-x), \qquad u_2^s = v - p_2^s - \theta(1-x),$$

where v is the value of the product, p_1^k and p_2^k are the prices set by retailers 1 and 2 respectively,

θ measures the degree of consumer preferences for retailers, and $x \in [0,1]$ is the consumer's location (in preference space) relative to retailer 1's location. Retailers are located at $x = 0$ and $x = 1$.

Whereas these utilities apply to all consumers, not all segments have access to all purchase options. Figure 3.1 details the choice set available to each segment. For example, the channel-agnostic decided consumer (AD) has access to all 4 options, whereas the store-only undecided consumer (SU) only has the option of purchasing from his preferred store (say, store of retailer 1). Thus, consumer heterogeneity results in different choice sets available to each segment.

Undecided consumers (both SU and AU), who don't know which specific product they need, first visit the retailer closer to their preference location (i.e., visit retailer 1 if $x < \frac{1}{2}$ and retailer 2 otherwise). After their shopping trip, the store-only undecided segment (SU) must decide whether to buy the product that fits their needs or make no purchase at the store; hence only the corresponding u^s-expression in (3.1) is relevant for such a consumer. Channel-agnostic undecided consumers (AU) can either purchase in the store they first visit and pay the store price, or return home and make an online purchase from either retailer; the utility expressions u_1^s, u_1^{on}, u_2^{on} are thus relevant for AU consumers who prefer retailer 1, and u_2^s, u_1^{on}, u_2^{on} are relevant for AU consumers who prefer retailer 2.

The Impact of Self-Matching Prices

We now examine how self-matching practices by retailers impact consumer utilities. Decided consumers know *all* prices for the specific product they want, and if they shop at the store offering self-matching, they can come armed with the online price and request a price-match. Thus, decided consumers can obtain a price-match in-store whereas undecided consumers cannot do so.

When both retailers offer a self-matching pricing policy, i.e., under $(SM_1, SM_2) = (1,1)$, a consumer located at $x \in [0,1]$ who *can* obtain a self-match faces the following utilities:

$$u_1^{on} = v - p_1^{on} - \theta x, \qquad u_1^s = v - \min(p_1^s, p_1^{on}) - \theta x \tag{3.2}$$
$$u_2^{on} = v - p_2^{on} - \theta(1-x), \qquad u_2^s = v - \min(p_2^s, p_2^{on}) - \theta(1-x).$$

Consumers who *cannot* obtain a price self-match (i.e., undecided consumers) continue to face the same utilities as specified in equation (3.1). Note that although the utility expressions remain the same, retailers may set different prices under different self-matching scenarios. Hence, the equilibrium utilities experienced by consumers will typically differ depending on retailer self-matching policies.

Next, consider consumers' channel preferences. Channel-agnostic decided consumers (AD) have no particular preference for any channel and would choose the lower-priced channel option. Store decided consumers (SD) choose one of the stores, based on their preferences and prices. However, now they can obtain the lower online price if the retailer offers a self-matching policy. Thus, the expressions for u_j^s for $j = 1, 2$ are different in (3.2) compared to (3.1). Undecided consumers (both AU and SU) do not know which product they want until they visit the store. They face the same utilities under $(1, 1)$ as under $(0, 0)$ since they cannot redeem matching policies when they visit a retailer's store without making an additional costly set of trips: back home to determine online prices and then back to a store to make the purchase.

Utilities in the asymmetric subgame $(1, 0)$, where only retailer 1 offers to self-match prices, are defined similar to the $(0, 0)$ case, with only u_1^s changing for decided consumers, who can obtain a self-match only from retailer 1 but not retailer 2:[5]

$$u_1^s = v - \min(p_1^s, p_1^{on}) - \theta x.$$

3.4 Analysis

We begin our analysis by considering the benchmark monopoly case, then the multichannel duopoly setting. All proofs are in the Appendix and threshold values and constants are defined there as well. Note that in all cases, we will derive conditions for the market to be covered in the proof, and our discussion in the text will focus on the region of coverage in equilibrium.[6] For exposition, we assume that $0 < \beta < \beta'$ where β' is defined in Appendix B.1. $\hat{}$ denotes equilibrium strategies and outcomes, e.g. $\hat{p}_1^s$.

[5] The other asymmetric equilibrium $(0, 1)$ is obtained by relabeling the retailers.

[6] Formally, we require bounds on v and β, which are detailed in the appendix. In Appendix B.4 we explore a setting where the market is not fully covered.

3.4.1 Benchmark Monopoly: A Single Entity Owns Both Retailers

Consider a monopolist that jointly maximizes the profits of two multichannel retailers by choosing a self-matching policy and setting prices.[7] The following holds:

Proposition 1 *A monopolist cannot increase profits by self-matching prices across channels.*

The monopolist will price to extract the greatest surplus from each channel. Since both undecided and decided consumers are present in both channels, the prices charged will be the same in both and equal to the monopoly price, regardless of whether or not the monopolist offers a self-matching policy. The monopolist thus obtains no additional profit when offering the policy and will not offer it when it entails a minimal implementation cost.

3.4.2 Multichannel Duopoly

We now consider the case of two competing multichannel retailers who make decisions according to the timeline in Figure 3.1. We examine each of the possible self-matching policy subgames and conclude with a result highlighting the conditions under which self-matching emerges in equilibrium.

For notational convenience, we define the function $\Phi(p_1^k, p_2^k; \theta) := \frac{1}{2} + \frac{(p_2^k - p_1^k)}{2\theta}$ to represent the proportion of demand obtained by retailer 1 from a specific segment of consumers who face prices p_1^k and p_2^k from the two retailers.

No Self-Matching - (0,0): In the $(0, 0)$ subgame where neither retailer self-matches, store-only consumers (both SD and SU segments) purchase from the store channel and pay the store price. Channel-agnostic consumers (both AD and AU segments) can purchase from either retailer's online channel. AD consumers will begin their search process online, whereas AU consumers will first visit their preferred retailer to browse products, but then return home to purchase online after they discover the type of product they wish to purchase. Retailers compete for these two consumer segments in the online channel.

[7]Note that qualitatively similar results hold if the retailer is located at the center of the city.

SU consumers visit the retailer closest in preference to them to learn about products. Recall that these consumers do not purchase online and do not switch stores because of travel costs associated with a second store visit. In a sense, these consumers can be thought of as having just one store as their purchase option in addition to the no-purchase option. Each retailer effectively has a subset ($\frac{\beta}{2}$) of such consumers.

On the other hand, SD consumers, who know the product they want and are informed of all prices, have a choice of which retailer's store to purchase from. These consumers purchase in-store, but make a decision on which store to visit after factoring in their retailer preferences and prices. Thus, there is intense competition among retailers for this segment, since by reducing store price, a retailer can attract more SD consumers.

The profit functions of retailers 1 and 2 can be written as:

$$\Pi_1^{0,0} = \underbrace{\eta\Phi(p_1^{on}, p_2^{on})p_1^{on}}_{\text{Channel-Agnostic Decided and Undecided}} + (1-\eta)\left(\underbrace{(1-\beta)\Phi(p_1^s, p_2^s)}_{\text{Store-Only Decided}} + \underbrace{\frac{\beta}{2}}_{\text{Store-Only Undecided}}\right) p_1^s,$$

$$\Pi_2^{0,0} = \eta\left(1-\Phi(p_1^{on}, p_2^{on})\right) p_2^{on} + (1-\eta)\left((1-\beta)(1-\Phi(p_1^s, p_2^s)) + \frac{\beta}{2}\right) p_2^s.$$

In the $(0,0)$ case, the situation online is similar to retailers competing in a horizontally differentiated market comprised of only channel-agnostic consumers. The resulting equilibrium prices, therefore, reflect the strength of consumers' preferences for retailers, with $\hat{p}_1^{on} = \hat{p}_2^{on} = \theta$. We will refer to a price of θ as the "competitive" price level to reflect the fact that this would be the price charged in a standard Hotelling duopoly model with one retail channel.

Next, we turn attention to the store channel, where we obtain symmetric equilibrium prices of

$$\hat{p}_1^s = \hat{p}_2^s = \begin{cases} v - \frac{\theta}{2}, & \frac{v}{\theta} \leq \frac{1}{2} + \frac{1}{1-\beta} \\ \frac{\theta}{1-\beta}, & \frac{v}{\theta} > \frac{1}{2} + \frac{1}{1-\beta} \end{cases} \tag{3.3}$$

There are a few useful observations to be made here. First, for $\frac{v}{\theta} \leq \frac{1}{2} + \frac{1}{1-\beta}$, retailers serve the entire market even though they charge the monopoly price in-store. This is possible because of the existence of SU consumers: retailers prefer to charge the monopoly price to extract all surplus from SU consumers if the ratio of product value to retailer differentiation is sufficiently low. Second, if $\frac{v}{\theta} > \frac{1}{2} + \frac{1}{1-\beta}$, then retailers charge a store price of $\frac{\theta}{1-\beta}$, which is larger than the competitive price of θ. When $\frac{v}{\theta}$ is sufficiently large, retailers can no longer maintain monopoly prices in-store and prefer to compete for SD consumers. However, they can still charge a higher price in-store than online. The existence of SU consumers enables retailers to charge more in-store than online, and the retailers extract surplus from this segment.

Symmetric Self-Matching - (1,1): In this case, both retailers implement a self-matching policy. The first and most obvious result of self-matching is the *channel arbitrage effect*, and the intuition here is straightforward. Recall that SD consumers shop in-store and pay $\hat{p}_j^s$ as in equation (3.3) absent a self-matching policy. However, with a self-match they pay the lower online price while shopping in-store, resulting in less profit for the retailer due to arbitrage.

Although this arbitrage intuition is correct, it is incomplete in determining whether in equilibrium a self-matching pricing policy will be adopted. When a multichannel retailer chooses to self-match, there emerges an important distinction between the store-only decided (SD) and undecided (SU) consumers. Whereas the SD consumers are able to obtain a price-match, SU consumers only know which product they desire *after* visiting a store. Since they lack evidence of a lower online price, they always pay the store price. Thus, even though the two segments of store consumers both obtain the product in-store, they effectively pay different prices. Self-matching enables to retailer to price discriminate consumers based on their decision-stage.

Retailers' profits in this self-matching setting are thus:

$$
\begin{aligned}
\Pi_1^{1,1} &= \underbrace{(1-\beta(1-\eta))\,\Phi_1(p_1^{on},p_2^{on})p_1^{on}}_{\text{Channel-Agnostic \& Store-Only Decided}} + \underbrace{(1-\eta)\frac{\beta}{2}p_1^s}_{\text{Store-Only Undecided}}, \\
\Pi_2^{1,1} &= (1-\beta(1-\eta))\,(1-\Phi_1(p_1^{on},p_2^{on}))p_2^{on} + (1-\eta)\frac{\beta}{2}p_2^s.
\end{aligned}
$$

In the pricing sub-game, retailers set equilibrium online prices $\hat{p}_1^{on} = \hat{p}_2^{on} = \theta$, since there is no force to prevent online prices from dropping to their competitive level. However, retailers set store prices $\hat{p}_1^s = \hat{p}_2^s = \left(v - \frac{\theta}{2}\right)$ to extract surplus from their respective "captive" sets of SU customers, who pay the store price.

Figures 3.2(a) and 3.2(b) compare pricing and purchase outcomes across the $(0,0)$ and $(1,1)$ subgames. The light-shaded regions cover the segments that pay the store price, whereas the darker regions show the segments that pay the online price. In the $(0,0)$ subgame, both SU and

SD consumers purchase in-store and pay the same store price, whereas AU and AD consumers purchase online and pay the online price. In the $(1,1)$ subgame, SU and SD consumers both purchase in-store, but SD consumers now pay the online price. Thus, self-matching allows a retailer to simultaneously segment and price discriminate consumers across decision-stages and de-segment decided consumers across channels. We refer to the ability to extract additional surplus from SU consumers through self-matching as the *decision-stage discrimination effect*.

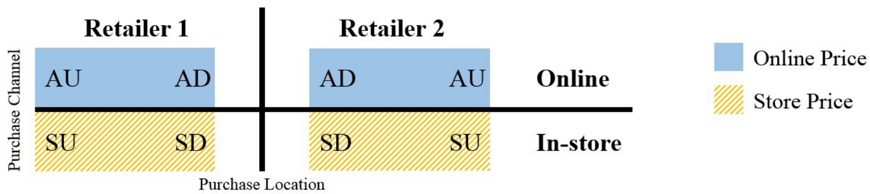

(a) No Self-Matching - (0,0)

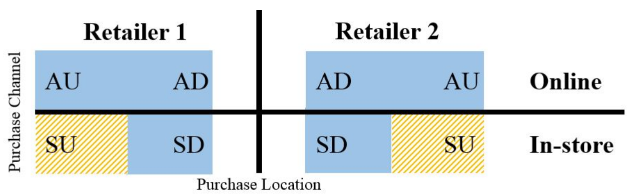

(b) Symmetric Self-Matching - (1,1)

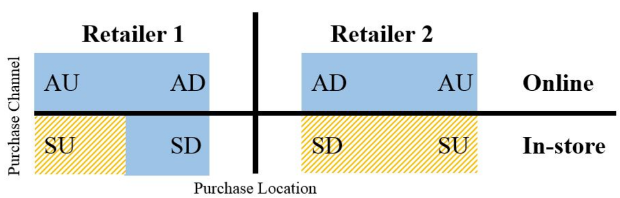

(c) Asymmetric Self-Matching - (1,0)

Figure 3.2: *The Effects of Self-Matching*

Asymmetric Self-Matching - (1,0): Next, we explore the case where retailer 1 offers a self-matching policy while retailer 2 does not, i.e., the $(1,0)$ self-matching subgame. By symmetry (or relabeling), similar results follow in the $(0,1)$ subgame. Observe that SD consumers who visit retailer 1's store can purchase there and pay the lower of the online and store price, i.e., $\min(p_1^{on}, p_1^s)$. However, if an SD consumer visits retailer 2's store instead, she faces a price of p_2^s and cannot obtain the online price in-store (since retailer 2 does not self-match). Moreover, by offering a self-matching policy, and as long as its online price satisfies $p_1^{on} < p_2^s$, retailer 1 attracts some SD consumers who are closer in preference to the competing retailer 2 but who choose to visit retailer 1's store in anticipation of paying the lower online price through a self-match.

For SD consumers, under $p_1^{on} < p_1^s$, the store price of retailer 1 is irrelevant (since they retrieve the price-match), and the retailer can set a store price level to capture the highest possible surplus from the SU consumers who are closer to its location. Thus, decision-stage price discrimination persists in the asymmetric subgame.

We obtain the following profit functions:

$$\Pi_1^{1,0} = \underbrace{\eta\Phi_1(p_1^{on}, p_2^{on})p_1^{on}}_{\text{Channel-Agnostic}} + (1-\eta)\left(\underbrace{(1-\beta)\Phi_1(p_1^{on}, p_2^{s})p_1^{on}}_{\text{Store-Only Decided}} + \underbrace{\frac{\beta}{2}p_1^s}_{\text{Store-Only Undecided}}\right),$$

$$\Pi_2^{1,0} = \eta(1-\Phi_1(p_1^{on}, p_2^{on}))p_2^{on} + (1-\eta)\left((1-\beta)(1-\Phi_1(p_1^{on}, p_2^{s})) + \frac{\beta}{2}\right)p_2^s.$$

Solving for the second-stage pricing subgame, we find that the price levels chosen by the retailers are *higher* than the Hotelling competitive price of θ *in both channels* and critically depend on the ratio of product value v to the retailer preference parameter θ as follows. For $\frac{v}{\theta} \leq \left(\frac{4}{3} + \frac{1}{6(1-\beta(1-\eta))} + \frac{\beta}{2(1-\beta)}\right)$, both retailers will extract all surplus from SU consumers and set prices

$$\hat{p}_1^{on} = \theta + \frac{(1-\beta)(1-\eta)(2v-3\theta)}{4(1-\beta(1-\eta))-\eta}, \qquad \hat{p}_2^{on} = \theta + \frac{(1-\beta)(1-\eta)(2v-3\theta)}{8(1-\beta(1-\eta))-2\eta},$$
$$\hat{p}_1^s = \hat{p}_2^s = v - \frac{\theta}{2}.$$

For $\frac{v}{\theta} > \left(\frac{4}{3} + \frac{1}{6(1-\beta(1-\eta))} + \frac{\beta}{2(1-\beta)}\right)$, we find that retailers set prices

$$\hat{p}_1^{on} = \theta\left(\frac{2}{3} + \frac{1}{3(1-\beta(1-\eta))}\right), \qquad \hat{p}_2^{on} = \theta\left(\frac{5}{6} + \frac{1}{6(1-\beta(1-\eta))}\right),$$
$$\hat{p}_1^s = v - \frac{\theta}{2}, \qquad \hat{p}_2^s = \hat{p}_2^{on} + \frac{\beta\theta}{2(1-\beta)}.$$

Interestingly, equilibrium online prices in the asymmetric self-matching $(1,0)$ case are greater than those set in the no self-matching $(0,0)$ case and the symmetric self-matching $(1,1)$ case. The intuition follows from the idea that while the self-matching retailer 1 loses profit from the SD segment of consumers who can invoke the price self-match, the policy effectively acts like a "commitment device" to prevent online prices from going all the way down to the competitive level. More importantly, this results in channel-agnostic AD and AU consumers paying a higher price (relative to the competitive online price of θ they were paying under no self-matching or symmetric self-matching). Thus, self-matching has a positive effect on profits through this third mechanism, which we term the *online competition dampening effect*.

Note that the situation in the asymmetric $(1, 0)$ case differs from the case when both retailers self-match. Under $(1, 1)$, SD consumers can redeem the online price at both retailers' stores, which forces online prices down to their competitive level θ. By contrast, Figure 3.2(c) illustrates how in the asymmetric $(1, 0)$ case, SD consumers can only redeem the match from retailer 1. Retailer 2 will price higher in-store relative to retailer 1's online price, as both SD consumers and its captive segment of SU consumers pay its store price, whereas retailer 1 fully segments out its store consumers through the self-matching policy which is invoked by its SD consumers (while retailer 1's SU consumers continue to pay its store price). Consequently, retailer 1 does not set its online price as low as θ and chooses a higher online price to mitigate channel arbitrage. Since online prices are strategic complements across retailers, retailer 2's best response is to increase its online price as well. This results in online prices at both retailers being higher than the competitive level, leading to online competition dampening.

The results detailed in this section are based on the pricing subgame, taking the self-matching policies as given. The pricing equilibria depend on the magnitudes of the three effects induced by self-matching. Table 3.3 presents a summary of the effects we have identified. It is helpful to observe that the negative channel arbitrage effect and the positive decision-stage discrimination effect always occur for a self-matching retailer, while online competition dampening occurs only when one self-matches but the rival does not. We now examine the full equilibrium results of the game, beginning with the self-matching strategy choices.

3.4.3 Self-Matching Policy Equilibria in a Multichannel Duopoly

For a self-matching policy configuration to emerge in equilibrium, it must be the case that neither retailer would be better off by unilaterally deviating to offer a different policy. Proposition 2 details the equilibrium conditions and the resulting choices of self-matching policies. Across all regions of the parameter space, we restrict focus to Pareto-dominant equilibria.

Table 3.3: *Effects of Self-Matching for Retailer 1 in a Multichannel Duopoly*

Effect	Relevant Subgames
(−) **Channel Arbitrage**: SD consumers redeem lower online price in the store, reducing profits from SD.	(1,0) (1,1)
(+) **Decision-Stage Discrimination**: Retailer avoids competing for SD segment on store prices; instead letting them obtain lower online prices. This allows higher store prices to captive SU segment.	(1,0) (1,1)
(+) **Online Competition Dampening**: Retailer charges higher online price to mitigate arbitrage, increasing profit from AD, AU and SD segments.	(1,0)

Proposition 2 *In a duopoly featuring two multichannel retailers, self-matching policies are determined by the following mutually exclusive regions:*

Asymmetric Equilibrium $(1,0)$**:** *One retailer will offer to self-match its prices while the other will not when product values are relatively low or retailer differentiation is high.*

Symmetric non-matching equilibrium $(0,0)$**:** *Neither retailer will self-match its prices when product values and retailer differentiation are at intermediate levels.*

Symmetric matching equilibrium $(1,1)$**:** *Both retailers will self-match prices when product values are high or retailer differentiation is low.*

The above result indicates that all three types of joint strategies can emerge in equilibrium depending on the nature of the product and competitive interaction. To understand the intuition behind the emergence of the different equilibria, it is critical to examine how the focal retailer's *best response function* evolves as the ratio of product value to retailer differentiation $\left(\frac{v}{\theta}\right)$ changes. We translate best response functions into equilibria in Figure 3.3. The topmost arrow depicts retailer 1's best response if retailer 2 does not self-match. The middle arrow depicts retailer 1's best response if retailer 2 self-matches. The dominant effects for retailer 1 are listed underneath the arrows in bold. The bottom arrow shows the emergent Pareto-dominant equilibria. The best response of the focal retailer depends on the competitor's self-matching strategy as well as the three effects we have previously described – *channel arbitrage, online competition dampening, and decision-stage discrimination*.

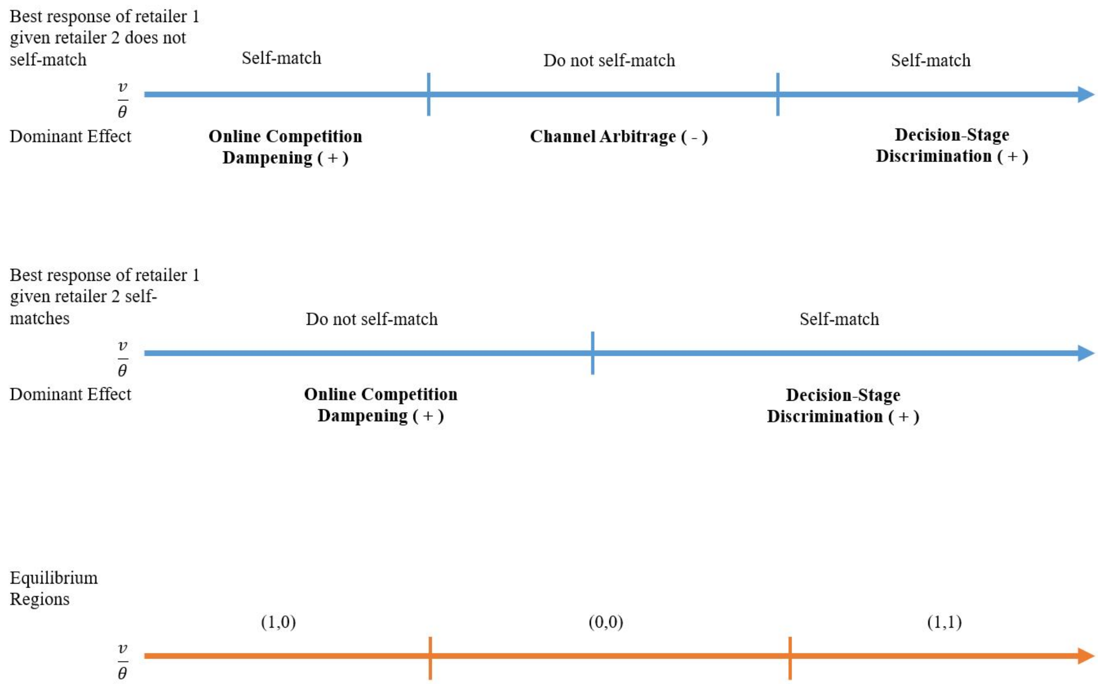

Figure 3.3: *Retailer Best Responses*

Retailer 1's Best Response to Retailer 2 Not Self-Matching: For low $\frac{v}{\theta}$, the retailer's store price $\left(v - \frac{\theta}{2}\right)$ is close to its competitive online price θ because there is not much additional surplus the retailer can extract from its captive SU consumers by pricing higher in-store. As a result, effects that have an impact on the store channel, i.e., both the negative channel arbitrage effect and the positive decision-stage discrimination effect are negligible. However, online prices can increase with self-matching due to the online competition dampening effect. This leads retailer 1 to offer a self-matching policy to take advantage of the additional profits from the online channel.

As $\frac{v}{\theta}$ increases, so does the difference in prices across channels. For intermediate values of $\frac{v}{\theta}$, the retailer can extract more surplus from SU consumers, driving it to price higher in-store even if it does not self-match, making the impact of decision-stage discrimination small. Since a self-matching policy allows SD consumers to redeem the lower online price, the channel arbitrage effect increases. As competition in the online channel is more intense than in-store, the positive online competition dampening effect can no longer overcome the negative channel arbitrage effect. Consequently, the channel arbitrage effect dominates the other two effects, and the retailer no longer finds it profitable to self-match as a best response.

At high $\frac{v}{\theta}$ levels, retailer 1 is compelled to compete more intensely for SD consumers closer to retailer 2 in preference, resulting in a store price of $\frac{\theta}{1-\beta}$ that no longer grows in v. Thus, the negative impact due to channel arbitrage is limited. However, if the retailer were to offer a self-matching policy, decision-stage discrimination would allow it to charge the monopoly price $\left(v - \frac{\theta}{2}\right)$ to captive SU consumers, which increases as $\frac{v}{\theta}$ increases. This creates a strong positive impact on profits, resulting in retailer 1 choosing to offer self-matching.

Retailer 1's Best Response to Retailer 2 Self-Matching: We now turn our attention to the case when retailer 2 decides to offer a self-matching policy. Recall that when both retailers self-match, the online competition dampening effect ceases to exist. Since retailer 2 is self-matching, its actions will result in online competition dampening only if retailer 1 does *not* self-match. This creates an incentive for retailer 1 to refrain from self-matching at low values of $\frac{v}{\theta}$. If retailer 1 does not self-match, its decided store consumers can obtain retailer 2's lower online prices, so retailer 1 has an incentive to lower its store price to retain them. Since the store prices are set low, the benefits of decision-stage discrimination are also reduced, and retailer 1 cannot extract a substantial amount of surplus from SU consumers. As $\frac{v}{\theta}$ increases, the benefit of decision-stage discrimination grows because the retailer can extract greater surplus from SU consumers if it can charge them a different price than SD consumers. This leads retailer 1 to adopt self-matching for high values of $\frac{v}{\theta}$.

Strategic Substitutes or Complements: We integrate the best responses to obtain equilibrium strategies and focus on whether self-matching strategies across retailers are strategic complements or substitutes. We find from the best responses that at low product values, or at high levels of retailer differentiation, the self-matching strategies act like *strategic substitutes*, so that a retailer will choose the strategy opposite to that of its competitor. As $\frac{v}{\theta}$ increases to an intermediate level, we obtain a symmetric equilibrium where no retailer self-matches and strategies are *strategic complements*. Finally, when $\frac{v}{\theta}$ is above a high threshold, the strong impact of decision-stage discrimination leads to self-matching *as a dominant strategy* regardless of what the competitor chooses. Figure 3.4 shows the equilibrium regions that emerge in the $\frac{v}{\theta} \leftrightarrow \beta$ space for a fixed value of $\eta \in (0, 1)$ based on Proposition 2.

We turn our attention to how the equilibrium regions are affected by β and η by stating the

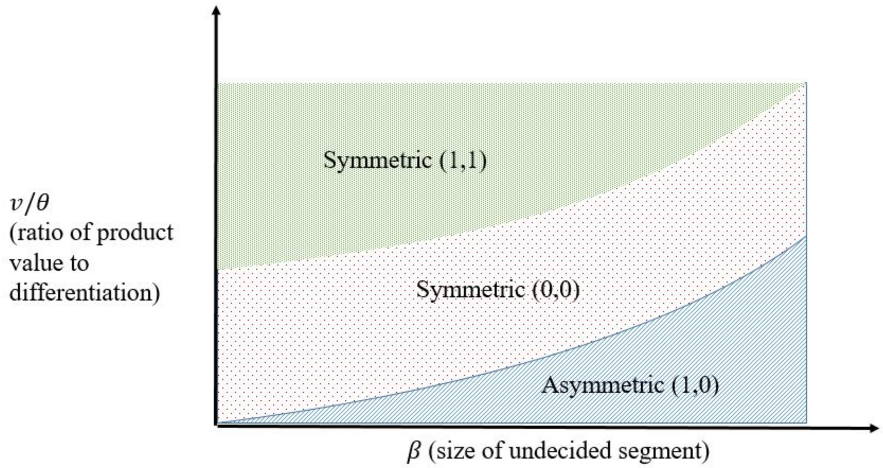

Figure 3.4: *Equilibrium Regions*

following Corollary.

Corollary 1 *An increase in the fraction of undecided consumers β will grow the asymmetric equilibrium region and shrink the symmetric equilibrium regions.*

To understand the intuition for Corollary 1, consider the case of focal retailer 1's best response when retailer 2 does not self-match. According to the Corollary, retailer 1 has more of an incentive to offer a self-matching policy as β increases, implying that the $\frac{v}{\theta}$-region for which we can sustain the $(1,0)$ equilibrium expands. This happens because as the fraction of undecided consumers increases, the retailers stand to gain more from online competition dampening (because of the greater fraction AU consumers). Thus, if retailer 1 self-matches, retailer 2 will refrain from doing so (because when both self-match, the online competition dampening effect is nullified).

The next Corollary examines the effect of η on the equilibrium regions.

Corollary 2 *An increase in the fraction of channel-agnostic consumers η will grow the region where retailers choose not to self-match.*

The profitability of self-matching depends largely on the existence of SU consumers. When the fraction of store-only consumers decreases (i.e., η grows), retailers can no longer benefit as much from decision-stage discrimination and online competition dampening. The profitability of offering a self-matching policy decreases as η increases.

3.4.4 The Profitability of Self-Matching

We have thus far analyzed how retailers decide whether to adopt self-matching policies and characterized the strategies that can be sustained in equilibrium. Here, we examine the profit impact of having self-matching available as a strategic option. The key issue we seek to understand is whether retailers are compelled by competitive forces to adopt self-matching, even though it might make them worse off and result in lower equilibrium profits than were self-matching not an option. The result in Proposition 3 below addresses this issue.

Proposition 3 *The profit implications of the self-matching, compared to the baseline case where self-matching is not available as an option, are as follows:*

(a) In the asymmetric equilibrium $(1,0)$*: the retailer offering to self-match earns greater profits, but the competing retailer earns lower profits.*

(b) In the symmetric self-matching equilibrium $(1,1)$*: both retailers earn higher profits when product valuation is high or retailer differentiation is low. Otherwise, they both earn lower profits.*

At low values of $\frac{v}{\theta}$, the profit impact of self-matching is mixed, with the self-matching retailer obtaining higher profits.

We find that at intermediate levels of $\frac{v}{\theta}$, self-matching occurs in equilibrium and reduces profits for both retailers because of the lower positive impact of decision-stage discrimination and the increasing negative impact of channel arbitrage. This interaction results in a situation where both retailers would have been better off had self-matching not been an option. However, when $\frac{v}{\theta}$ is high, we find that both retailers choose to self-match and earn higher profits by doing so. This occurs because at high $\frac{v}{\theta}$, decision-stage discrimination overtakes the negative impact of channel arbitrage. Overall, we find that the availability of self-matching as a strategy has the potential to enhance profits for at least one retailer and can also do so for both retailers for a range of parameters, highlighting the importance of self-matching as a strategic option.

3.5 Extensions

The base model analyzed in Section 3.4 focused on developing an understanding of the mechanisms underlying the effectiveness of self-matching and the conditions for retailers to implement the policy in equilibrium. Here, we have two main objectives. First, we aim to examine additional settings that are of relevance to retailers as they contemplate whether to offer a self-matching pricing policy. Second, we relax a few key assumptions in the baseline model, with a view towards increasing the range of applicability of the findings. Proofs for results here are provided in Appendix B.3 and Appendix B.4.

3.5.1 Impact of "Smart-Device" Enabled Consumers

The baseline model characterized undecided consumers as not knowing what specific product they want until they visit a store to evaluate which item from the many available options best fits their needs. They could not invoke a self-matching policy because they were in-store at the time of their final decision, and there was no way for them to access the internet to produce evidence of a lower online price.

Here, we recognize the increasing importance of mobile devices to alter this dynamic and examine the implications for self-matching policies. Retail TouchPoints (Fiorletta, 2013) notes that, *"Amplified price transparency due to the instant availability of information via the web and mobile devices has encouraged retailers to rethink their omnichannel pricing strategies."* Intuitively, one might expect that the greater the proportion of consumers who carry smart devices and take the trouble to check online when in-store, the less profitable self-matching should be (because of the increased threat of cross-channel arbitrage). We show that this need not be the case.

Suppose that a fraction μ $(0 < \mu < 1)$ of consumers has access to the internet while shopping in-store. We refer to these consumers as "smart" to reflect the notion that with the aid of internet-enabled smartphone devices these consumers can easily obtain online price information while in-store. Store-only undecided smart consumers will invoke a self-matching policy if the online price offered by a retailer is lower than its store price.

Channel-agnostic undecided smart consumers will behave as we have already modeled in the baseline model. An increase in smart consumers can be understood as increasing the fraction of store-only undecided (SU) consumers who redeem the online price. However, these consumers can only purchase from the retailer they first visit, in contrast to store-only decided (SD) consumers who have the option of buying from other retailers. Effectively, each retailer obtains an additional subsegment of store undecided consumers who pay the online price. To see how the existence of smart consumers impacts retailers' strategies, consider the profits retailers earn if they both offer to self-match:

$$\Pi_1^{1,1} = (1-\beta(1-\eta))\,\Phi_1(p_1^{on}, p_2^{on})p_1^{on} + (1-\eta)\frac{\beta}{2}\left((1-\mu)p_1^s + \mu p_1^{on}\right),$$
$$\Pi_2^{1,1} = (1-\beta(1-\eta))\,(1-\Phi_1(p_1^{on}, p_2^{on}))p_2^{on} + (1-\eta)\frac{\beta}{2}\left((1-\mu)p_2^s + \mu p_2^{on}\right).$$

Solving for the equilibrium reveals that retailers will set $p_1^{on} = p_2^{on} = \theta(1-\mu) + \frac{\mu\theta}{1-\beta(1-\eta)}$ and $\hat{p}_1^s = \hat{p}_2^s = v - \frac{\theta}{2}$. Note that the online prices are increasing in μ. We detail how smart consumers impact retailers' equilibrium incentives to self-match in Proposition 4.

Proposition 4 *In a duopoly with two multichannel retailers, where some consumers can use a smart device in-store to obtain online price information:*

(a) As the fraction of smart consumers increases, the region of self-matching equilibria grows for low product values but shrinks for high product values.

(b) Retailer profits can increase in the fraction of smart consumers.

At low product values, holding fixed the other model parameters, more smart consumers enhance the online competition dampening effect, which allows retailers to price higher online when offering to self-match. On the other hand, the conditions for symmetric self-matching policies to emerge in equilibrium for high product values become more stringent as μ grows. Namely, as $\mu \to 1$, the symmetric self-matching region for high v shrinks in size to zero. This happens because the existence of smart consumers greatly erodes the positive decision-stage discrimination effect of self-matching, as there are less SU consumers that will still pay the high store price, while more consumers pay the lower online price; thereby reducing retailers' incentives to self-match.

Thus, and somewhat counterintuitively, the presence of smart consumers need not decrease the profitability of a self-matching retailer (see Appendix for details of the profit enhancing case). On the contrary, smart consumers can enable retailers to charge higher online prices, increasing the profitability of self-matching policies. This suggests that given current technology trends, horizontally differentiated retailers would find it worthwhile to more carefully examine whether self-matching is an appropriate strategic option.

3.5.2 Mixed Duopoly: Multichannel Retailer and E-tailer

We consider the case of a multichannel retailer facing a pure online e-tailer, i.e., a "mixed duopoly market." This market structure is becoming more important for a number of multichannel retailers, e.g., several retailers find that Amazon and potentially other e-tailers are their primary rivals. Past research has considered the strategic implications of direct sellers, such as e-tailers, competing with traditional retail channels (Balasubramanian, 1998). However, self-matching policies have not been examined in this setting as competitors are assumed to operate a single channel.

We denote the focal multichannel retailer as retailer 1 and the online-only e-tailer as retailer 2. In this setting, only retailer 1 can offer a self-matching policy in stage 1 of the game. Subsequently, both retailers set prices and compete for demand per the timeline in Figure 3.1.

First, consider the case when the multichannel retailer does not self-match its prices. Store-only consumers can only consider retailer 1's store channel and are captive to this retailer, whereas channel-agnostic consumers have the option of shopping across the two retailers' online sites. Profits for both retailers can be expressed as follows.

$$\begin{aligned} \Pi_1^{0,0} &= \underbrace{\eta\Phi_1(p_1^{on}, p_2^{on})p_1^{on}}_{\text{Channel-Agnostic Decided \& Undecided}} + \underbrace{(1-\eta)p_1^s}_{\text{Store-Only Decided \& Undecided}}, \\ \Pi_2^{0,0} &= \underbrace{\eta(1-\Phi_1(p_1^{on}, p_2^{on}))p_2^{on}}_{\text{Channel-Agnostic Decided \& Undecided}}. \end{aligned} \tag{3.4}$$

Retailer 1 serves as an effective monopolist for store-only consumers (both SD and SU segments), who comprise a combined segment of size $1-\eta$, and will attempt to extract surplus from them by setting a store price of $\hat{p}_1^s = v - \theta$. Note that in contrast to the multichannel duopoly, store-only decided consumers do not drive prices down in the mixed duopoly case because the e-tailer does not have a store that serves as a competitive option. Both retailers compete online for the channel-agnostic consumers, who form a segment of size η. We allow for channel-agnostic undecided consumers located closer in preference to retailer 2, the e-tailer, to browse the product category at retailer 1's store and then purchase online from the e-tailer. The equilibrium online prices are at the competitive level, with $\hat{p}_1^{on} = \hat{p}_2^{on} = \theta$.

Next, consider the $(1,0)$ subgame where the multichannel retailer offers a self-matching policy. SD consumers can now retrieve the multichannel retailer's online price in-store. Consider the retailers' profits as given below:

$$\begin{aligned} \Pi_1^{1,0} &= \underbrace{\eta\Phi(p_1^{on}, p_2^{on})p_1^{on}}_{\text{Channel-Agnostic Decided \& Undecided}} + \underbrace{(1-\eta)(1-\beta)p_1^{on}}_{\text{Store-Only Decided}} + \underbrace{(1-\eta)\beta p_1^{s}}_{\text{Store-Only Undecided}}, \\ \Pi_2^{1,0} &= \eta(1-\Phi(p_1^{on}, p_2^{on}))p_2^{on}. \end{aligned} \tag{3.5}$$

As in the no self-matching case, competition online places downward pressure on the price levels p_1^{on} and p_2^{on} in the $(1,0)$ case. However, a portion of store consumers, namely the $(1-\eta)(1-\beta)$-sized SD segment, now receive the online price by invoking the self-match policy instead of paying the store price. As in the multichannel duopoly case, retailer 1 thus faces a *channel arbitrage effect* when it allows consumers to obtain a price-match. We might intuitively expect self-matching to be unprofitable – especially since the SD consumers, regardless of retailer preferences, can not defect to the e-tailer due to their preference for the store channel. However, once again, the *online competition dampening effect* can act to increase profitability when the multichannel retailer chooses to self-match. The following result reflects the net impact of these effects.

Proposition 5 *In a mixed duopoly featuring a multichannel retailer and a pure e-tailer, the multichannel retailer chooses to adopt a self-matching policy when product value is relatively low or retailer differentiation is high. Otherwise, the retailer will not adopt a self-matching policy.*

The intuition for Proposition 5 follows naturally from the implications of the online competition dampening effect in the asymmetric case of the multichannel duopoly scenario. When retailer 1 decides to self-match, there is a cross-channel arbitrage externality. In a bid to reduce the negative impact of channel arbitrage, the multichannel retailer that implements a self-match has an incentive to raise its online price relative to the no self-matching case. Strategic complementarity in prices leads both retailers to set higher online prices than under no self-matching. Note that there is no decision-stage discrimination effect. SU consumers pay the same price regardless of whether or not there is self-matching because the e-tailer has no store to induce competition for the SD segment and lower the multichannel retailer's store price.

The trade-off between the channel arbitrage and competition dampening effects depends on $\frac{v}{\theta}$. For low enough v (or high θ), the self-matching multichannel retailer prices similarly across channels and thus channel arbitrage is low. In this case, the online competition dampening effect dominates and grows in $\frac{v}{\theta}$. However, the negative channel arbitrage effect also increases in $\frac{v}{\theta}$ (as SD consumers redeem the lower online price) and eventually dominates the online competition dampening effect. As a result, self-matching emerges as an equilibrium outcome only for low values of $\frac{v}{\theta}$.

Turning attention to the profit impact of self-matching on the e-tailer, we find the following:

Corollary 3 *In a mixed duopoly, the e-tailer makes higher profits when the multichannel retailer uses a self-matching policy.*

Thus, a self-matching policy has a positive externality on the e-tailer due to reduced competition in the online channel, which allows the e-tailer to increase prices. This holds even though the e-tailer's price level is lower than that of the multichannel retailer, since the latter internalizes a higher benefit of raising its online price because of the positive impact on its store channel.

3.5.3 Additional Analyses

Incomplete Market Coverage

To relax the assumption that all markets are fully covered, we focus on a scenario where retailers compete in a linear city and also face markets that are not fully covered.[8] Specifically, retailers are located at $x = 0$ and $x = 1$ on a Hotelling line of length $\frac{1}{5}\left(3 + \frac{6v}{\theta}\right)$ such that the distance between retailers is still equal to 1 but retailers face additional monopoly segments of consumers outside of the unit interval.

We conduct equilibrium analysis for high levels of $\frac{v}{\theta}$ and find that retailers will choose to offer self-matching policies. This result coincides with Proposition 2, where we found that symmetric self-matching emerges in equilibrium at high levels of $\frac{v}{\theta}$. However, in contrast to the profitability results in Proposition 3, we find that retailers earn lower profits when self-matching than when self-matching was not available as a strategic option. This is because retailers now have an

[8] If $\frac{v}{\theta}$ is sufficiently low in the main model, retailers no longer compete and act as monopolists. The results of Proposition 1 apply and no retailer will choose to self-match.

incentive to keep store prices low, even when self-matching, to attract consumers located outside of the unit interval. As a result, the decision-stage discrimination effect is reduced, and the profitability of self-matching suffers.

Retailer Processing Costs

Retailers may incur a processing cost for consumers who redeem a self-matching policy, for example, the staff time verifying the evidence and entering it into the system. The idea here is somewhat analogous to the idea of hassle costs developed in Desai and Purohit (2004). In fully covered markets, an increase in retailer processing costs will reduce the profitability of self-matching by, in effect, increasing the magnitude of the channel arbitrage effect in the asymmetric subgame. As a result, self-matching is more difficult to sustain in equilibrium and the region labeled $(1,1)$ in Figure 3.4 shrinks as retailer processing costs grow. In Appendix B.4, we illustrate the impact a small but non-zero retailer cost of servicing consumers who redeem a self-matching policy has on retailer prices and profits to highlight the greater channel arbitrage effect and the reduced profitability of self-matching.

Different Structure of Consumer Heterogeneity in a Monopoly Setting

Proposition 1 establishes that self-matching will never be adopted by a monopolist and provides a benchmark for the duopoly analysis. However, the monopolist may choose to self-match in models that allow for a different structure of consumer heterogeneity. To illustrate how self-matching may be profitable for a monopolist, in Appendix B.4, we develop an alternative model where consumers exhibit heterogeneity in their travel costs and product valuations. All consumers are at first undecided and must visit the retailer's store to identify their preferred product. Consumers have heterogeneous product valuations that are perfectly correlated with their travel costs, i.e., consumers with a high product valuation have a high travel cost, and consumers with a low product valuation have a low travel cost. A self-matching policy may enable the monopolist to charge a higher store price and sell to store-only consumers with a high travel cost and a high product valuation as these consumers may find it costly to visit the store multiple times to redeem a self-matching policy.

3.6 Discussion, Limitations, and Conclusion

The self-matching pricing policy has become an important strategic aspect of multichannel retailing and is used in a variety of markets, including consumer electronics, discount retail, and office supplies. Our paper makes the first attempt to model this strategic pricing policy and investigate how a company's self-matching decision is determined by consumer behavior and the competitive landscape.

Retailers in our model choose whether to offer a self-matching pricing policy in the first stage and then set price levels in the second stage. The retailers' products are horizontally differentiated, with consumers having heterogeneous preferences over retailers. We further allow for consumer heterogeneity along two additional dimensions, decision-stage and channel preference. Thus, we explicitly capture a wide variety of decision-making processes for consumers enabled by the multichannel setting.

The analysis illustrates how retailers in a multichannel setting face downward price pressure in-store from competition induced by the presence of store-only decided consumers. By self-matching, a retailer relinquishes its ability to charge different prices to decided consumers across channels (de-segmentation). Channel de-segmentation induces channel arbitrage, but produces another effect - it can act as a commitment device to increase online prices when only one retailer chooses to self-match. We refer to this as the online competition dampening effect. Self-matching also enables the retailer to charge store-only undecided consumers a higher store price, which we call the decision-stage discrimination effect, and this can result in both retailers self-matching.

Self-matching is thus profitable when the positive effects of online competition dampening and decision-stage discrimination overcome the negative effects of channel arbitrage. We further find that the profitability of self-matching is determined by product value (relative to retailer differentiation), as well as consumer heterogeneity across different dimensions, such as decision-stage and channel preference.

Beyond the baseline model, we consider several extensions, one of which explicitly models a setting with smart-device enabled ("smart") consumers, who can look up online prices while in-store, an increasingly prevalent phenomenon. We find that self-matching can increase retailer profitability as the proportion of smart consumers increases. This consumer trend may prove to be an important issue for retailers to consider when making pricing policy decisions going forward.

Our model yields results that are empirically testable. First, retailers offering to self-match will have a larger online to store price discrepancy relative to those that do not self-match. Second, we should find asymmetric self-matching equilibrium configurations in markets with relatively low-valued products (or highly differentiated retailers). Third, as the penetration of smart-devices among consumers increases, online prices set by retailers offering to self-match are expected to rise.

Although we believe this to be the first research to rigorously examine the idea of self-matching as a pricing strategy, the present paper has several limitations that future research could consider addressing. First, we do not model *competitive* price-matching policies. Such policies have been extensively studied in the literature, and our focus is on retailers with unique product assortments, where competitive price-matching does not play a role. It would be interesting to examine whether self-matching complements or substitutes competitive matching policies in settings where competing retailers sell identical products. Second, by assuming sufficiently large consumer travel costs for store visits beyond the first, we ensure that retailers can price-discriminate their captive consumers who find it too costly to search additional stores for product information. Although we incorporated a variety of consumer decision-making processes and preference dimensions, it would be useful to consider a richer model of consumer search; for example, where consumers could visit a retailer's store, then decide whether to visit a second based on expectations of price as well as the benefits they may obtain. Such an effort would connect with the search literature, and it would be useful to examine whether self-matching then leads to more search and larger consideration sets in the spirit of Diamond (1971) and Liu and Dukes (2013). Third, the dimensions of consumer heterogeneity might be correlated, e.g., consumers who prefer store shopping may also be more undecided. While we don't expect this to change our primary findings, careful modeling of these dependencies might reveal additional effects. Finally, while we expect the mechanisms detailed here to apply to the case where there are *ex-ante* differences among retailers (based on costs or loyalty) beyond horizontal differentiation, there may be additional insights obtained in modeling the more general case.

Broadly, our findings suggest that although a self-matching policy may at first appear to be an unprofitable but necessary evil, it has more subtle and positive competitive implications. Indeed, self-matching can be profitably used as a strategic lever and can result in higher profits for all retailers in the industry. Multichannel retailers should therefore treat their self-matching decisions as an important element of their overall cross-channel strategy, taking into account the products they sell and market/consumer characteristics, as well as the competitive landscape.

References

ANDREWS, D. W. and SOARES, G. (2010). Inference for parameters defined by moment inequalities using generalized moment selection. *Econometrica*, **78** (1), 119–157.

BALASUBRAMANIAN, S. (1998). Mail versus mall: A strategic analysis of competition between direct marketers and conventional retailers. *Marketing Science*, **17** (3), 181–195.

BAYE, M. R., KOVENOCK, D. and DE VRIES, C. G. (1993). Rigging the lobbying process: an application of the all-pay auction. *American Economic Review*, pp. 289–294.

—, — and — (1994). The solution to the tullock rent-seeking game when $r > 2$: Mixed-strategy equilibria and mean dissipation rates. *Public Choice*, **81** (3-4), 363–380.

BEGGS, S., CARDELL, S. and HAUSMAN, J. (1981). Assessing the potential demand for electric cars. *Journal of Econometrics*, **17** (1), 1–19.

BERRY, S., KHWAJA, A., KUMAR, V., MUSALEM, A., WILBUR, K. C., ALLENBY, G., ANAND, B., CHINTAGUNTA, P., HANEMANN, W. M., JEZIORSKI, P. and MELE, A. (2014). Structural models of complementary choices. *Marketing Letters*, **25** (3), 245–256.

BORKOVSKY, R. N., ELLICKSON, P. B., GORDON, B. R., AGUIRREGABIRIA, V., GARDETE, P., GRIECO, P., GURECKIS, T., HO, T.-H., MATHEVET, L. and SWEETING, A. (2015). Multiplicity of equilibria and information structures in empirical games: challenges and prospects. *Marketing Letters*, **26** (2), 115–125.

BOUDREAU, K. J., LACETERA, N. and LAKHANI, K. R. (2011). Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Science*, **57** (5), 843–863.

—, LAKHANI, K. R. and MENIETTI, M. (2016). Performance responses to competition across skill levels in rank-order tournaments: Field evidence and implications for tournament design. *RAND Journal of Economics*, **47** (1), 140–165.

BUREAU OF LABOR STATISTICS, U. S. (2016). Occupational employment statistics.

BUSTILLO, M. and LAWTON, C. (2009). Best buy expands private-label brands. *The Wall Street Journal*.

BYERS, J. W., MITZENMACHER, M. and ZERVAS, G. (2012a). Daily deals: Prediction, social diffusion, and reputational ramifications. In *Proceedings of the 5th ACM International Conference on Web Search and Data Mining*, ACM, pp. 543–552.

—, — and — (2012b). The groupon effect on yelp ratings: a root cause analysis. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, ACM, pp. 248–265.

Chapman, R. G. and Staelin, R. (1982). Exploiting rank ordered choice set data within the stochastic utility model. *Journal of Marketing Research*, pp. 288–301.

Che, Y.-K. and Gale, I. (2003). Optimal design of research contests. *American Economic Review*, **93** (3), 646–671.

— and Gale, I. L. (1998). Caps on political lobbying. *American Economic Review*, pp. 643–651.

Chen, Y., Narasimhan, C. and Zhang, Z. (2001). Consumer heterogeneity and competitive price-matching guarantees. *Marketing Science*, **20** (3), 300–314.

Chernozhukov, V., Hong, H. and Tamer, E. (2007). Estimation and confidence regions for parameter sets in econometric models1. *Econometrica*, **75** (5), 1243–1284.

Cooper, T. (1986). Most-favored-customer pricing and tacit collusion. *RAND Journal of Economics*, **17** (3), 377–388.

Corts, K. (1998). Third-degree price discrimination in oligopoly: all-out competition and strategic commitment. *RAND Journal of Economics*, **29** (2), 306–323.

Coughlan, A. and Shaffer, G. (2009). Price-matching guarantees, retail competition, and product-line assortment. *Marketing Science*, **28** (3), 580–588.

Dahan, E. and Mendelson, H. (2001). An extreme-value model of concept testing. *Management Science*, **47** (1), 102–116.

Desai, P. S. and Purohit, D. (2004). "let me talk to my manager": Haggling in a competitive environment. *Marketing Science*, **23** (2), 219–233.

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory*, **3** (2), 156–168.

Dohmen, T. and Falk, A. (2011). Performance pay and multidimensional sorting: Productivity, preferences, and gender. *American Economic Review*, pp. 556–590.

Edelman, B., Jaffe, S. and Kominers, S. D. (2014). To groupon or not to groupon: The profitability of deep discounts. *Working Paper*.

— and Kominers, P. (2011). Consumer protection in online discount voucher sales.

Eizenberg, A. (2014). Upstream innovation and product variety in the us home pc market. *Review of Economic Studies*, **81** (3), 1003–1045.

Eriksson, T., Teyssier, S. and Villeval, M.-C. (2009). Self-selection and the efficiency of tournaments. *Economic Inquiry*, **47** (3), 530–548.

Fiorletta, A. (2013). Pricing challenges in an omni-channel world. *Retail Touchpoints*.

Fullerton, R. L. and McAfee, R. P. (1999). Auctioning entry into tournaments. *Journal of Political Economy*, **107** (3), 573–605.

Girotra, K., Terwiesch, C. and Ulrich, K. T. (2010). Idea generation and the quality of the best idea. *Management Science*, **56** (4), 591–605.

Grau, J. (2012). Multichannel pricing strategy. *eMarketer*.

Grewal, D., Janakiraman, R., Kalyanam, K., Kannan, P., Ratchford, B., Song, R. and Tolerico, S. (2010). Strategic online and offline retail pricing: A review and research agenda. *Journal of Interactive Marketing*, **24** (2), 138–154.

Gross, D. (2016). Performance feedback in competitive product development. *Working Paper*.

Gross, D. P. (2014). Creativity under fire: The effects of competition on innovation and the creative process. *Working Paper*.

Gupta, S., Lehmann, D. R. and Stuart, J. A. (2004). Valuing customers. *Journal of Marketing Research*, **41** (1), 7–18.

Holmes, T. (1989). The effects of third-degree price discrimination in oligopoly. *American Economic Review*, **79** (1), 244–250.

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

Hviid, M. and Shaffer, G. (1999). Hassle costs: The achilles' heel of price-matching guarantees. *Journal of Economics & Management Strategy*, **8** (4), 489–521.

Inman, J. and McAlister, L. (1994). Do coupon expiration dates affect consumer behavior? *Journal of Marketing Research*, pp. 423–428.

Ishii, J. (2008). Compatibility, competition, and investment in network industries: ATM networks in the banking industry. *Working Paper*.

Kalra, A. and Shi, M. (2001). Designing optimal sales contests: A theoretical perspective. *Marketing Science*, **20** (2), 170–193.

Kirby, J. (2013). Creative that cracks the code. *Harvard Business Review*, **91** (3), 86–89.

Klaassen, A. (2011). Harley-Davidson breaks consumer-created work from Victors & Spoils. *Ad Age*.

Lazear, E. P. and Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, **89** (5), 841–864.

Lee, R. S. and Pakes, A. (2009). Multiple equilibria and selection by learning in an applied setting. *Economics Letters*, **104** (1), 13–16.

Liu, L. and Dukes, A. (2013). Consideration set formation with multiproduct firms: The case of within-firm and across-firm evaluation costs. *Management Science*, **59** (8), 1871–1886.

Liu, Y., Gupta, S. and Zhang, Z. (2006). Note on self-restraint as an online entry-deterrence strategy. *Management Science*, **52** (11), 1799–1809.

Loch, C. H., Terwiesch, C. and Thomke, S. (2001). Parallel and sequential testing of design alternatives. *Management Science*, **47** (5), 663–678.

Luo, X., Andrews, M., Song, Y. and Aspara, J. (2014). Group-buying deal popularity. *Journal of Marketing*, **78** (2), 20–33.

Mattioli, D. (2011). Retailers try to thwart price apps. *The Wall Street Journal*.

Mehra, A., Kumar, S. and Raju, J. (2013). 'Showrooming' and the competition between store and online retailers. *Working Paper*.

Mehta, N., Rajiv, S. and Srinivasan, K. (2004). Role of forgetting in memory-based choice decisions: A structural model. *Quantitative Marketing and Economics*, **2** (2), 107–140.

Mohammed, R. (2013). Is Target's price matching policy a mistake? *HBR Blog Network*.

Moldovanu, B. and Sela, A. (2001). The optimal allocation of prizes in contests. *American Economic Review*, **91** (3), 542–558.

Moorthy, S. and Winter, R. (2006). Price-matching guarantees. *RAND Journal of Economics*, **37** (2), 449–465.

— and Zhang, X. (2006). Price matching by vertically differentiated retailers: Theory and evidence. *Journal of Marketing Research*, **43** (2), 156–167.

Mulpuru, S. (2010). The state of retailing online 2010: Key metrics and multichannel and global strategies. *Forrester Research*.

— (2012). US cross-channel retail forecast, 2011 to 2016. *Forrester Research*.

— (2013). US cross-channel retail forecast, 2012 to 2017. *Forrester Research*.

Nalca, A., Boyaci, T. and Ray, S. (2010). Competitive price-matching guarantees under imperfect store availability. *Quantitative Marketing and Economics*, **8** (3), 275–300.

Neslin, S., Grewal, D., Leghorn, R., Shankar, V., Teerling, M., Thomas, J. and Verhoef, P. (2006). Challenges and opportunities in multichannel customer management. *Journal of Service Research*, **9** (2), 95–112.

Ofek, E., Katona, Z. and Sarvary, M. (2011). Bricks and clicks: The impact of product returns on the strategies of multichannel retailers. *Marketing Science*, **30** (1), 42–60.

Offenberg, J. P. (2007). Markets: Gift cards. *Journal of Economic Perspectives*, **21** (2), 227–238.

Pakes, A., Porter, J., Ho, K. and Ishii, J. (2015). Moment inequalities and their application. *Econometrica*, **83** (1), 315–334.

Reich, G. (2014). Divide and conquer: Recursive likelihood function integration for dynamic discrete choice models with serially correlated errors. *Working Paper*.

Retail Systems Research (2013). Tough love: An in-depth look at retail pricing practices.

Sahni, N. S. (2015). Effect of temporal spacing between advertising exposures: Evidence from online field experiments. *Quantitative Marketing and Economics*, **13** (3), 203–247.

Salop, S. (1986). Practices that (credibly) facilitate oligopoly co-ordination. In *New Developments in the Analysis of Market Structure*, The MIT Press, p. 265.

Shin, J. (2007). How does free riding on customer service affect competition? *Marketing Science*, **26** (4), 488–503.

Shu, S. B. and Gneezy, A. (2010). Procrastination of enjoyable experiences. *Journal of Marketing Research*, **47** (5), 933–944.

Silk, T. and Janiszewski, C. (2008). Managing mail-in rebate promotions. *Working Paper*.

Sisak, D. (2009). Multiple-prize contests–the optimal allocation of prizes. *Journal of Economic Surveys*, **23** (1), 82–114.

Stein, W. E. (2002). Asymmetric rent-seeking with more than two contestants. *Public Choice*, **113** (3-4), 325–336.

Szymanski, S. and Valletti, T. M. (2005). Incentive effects of second prizes. *European Journal of Political Economy*, **21** (2), 467–481.

Tasoff, J. and Letzler, R. (2014). Everyone believes in redemption: Nudges and overoptimism in costly task completion. *Journal of Economic Behavior & Organization*.

Taubinsky, D. (2013). From intentions to actions: A model and experimental evidence of inattentive choice. *Working Paper*.

Taylor, C. R. (1995). Digging for golden carrots: an analysis of research tournaments. *American Economic Review*, pp. 872–890.

Terwiesch, C. and Xu, Y. (2008). Innovation contests, open innovation, and multiagent problem solving. *Management Science*, **54** (9), 1529–1543.

Thisse, J. and Vives, X. (1988). On the strategic choice of spatial price policy. *American Economic Review*, **78** (1), 122–137.

Toubia, O. (2006). Idea generation, creativity, and incentives. *Marketing Science*, **25** (5), 411–425.

Tullock, G. (1980). Efficient rent seeking. In *Efficient Rent-Seeking*, Springer, pp. 3–16.

Verhoef, P. C., Kannan, P. and Inman, J. J. (2015). From multi-channel retailing to omni-channel retailing: Introduction to the special issue on multi-channel retailing. *Journal of Retailing*, **91** (2), 174–181.

Wollman, T. (2014). Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *Working Paper*.

Wooten, J. O. and Ulrich, K. T. (2015a). Idea generation and the role of feedback: Evidence from field experiments with innovation tournaments. *Working Paper*.

— and — (2015b). The impact of visibility in innovation tournaments: Evidence from field experiments. *Working Paper*.

Yoganarasimhan, H. (2016). Estimation of beauty contest auctions. *Marketing Science*, **35** (1), 27–54.

Zhang, X. (2009). Retailers' multichannel and price advertising strategies. *Marketing Science*, **28** (6), 1080–1094.

Zhang, Z. (1995). Price-matching policy and the principle of minimum differentiation. *Journal of Industrial Economics*, pp. 287–299.

Appendix A

Appendix to Chapter 1

A.1 Do Participants Submit Their Best Idea First?

I estimate the difference between the rating of the first submission made by a participant and the ratings of her subsequent submissions, controlling for the total number of submissions made by the participant. For the set of entrants, I estimate the regression

$$\text{Rating}_{sit} = \alpha + \beta \text{First\_Submission}_{sit} + \epsilon_{sit}, \tag{A.1}$$

where Rating_{sit} is the rating of submission s made participant i in contest t, $\text{First\_Submission}_{sit}$ is an indicator for whether or not submission s was participant i's first submission, and ϵ_{sit} is an error term. Regression A.1 is estimated for all cases where participants made d submissions, where $d \in \{2, 3, 4, 5\}$. Table A.1 shows the resulting estimates of β. I find evidence that whenever participants make 5 submissions, the rating assigned to the first submissions tends to be higher than the rating assigned to subsequent submissions, although the difference is not economically significant. For all other categories, I fail to find a significant effect of submission order on submission rating. Similar results hold when I compare the rating of the first submission to the last submission and when I use a linear or logarithmic function of submission order instead of an indicator for first submission in Regression A.1. I conclude that there is limited evidence of a relationship between submission order and rating.

Table A.1: *Regression of Rating on Submission Order*

Total Number of Submissions	2	3	4	5
Relative Rating of First Submission	0.008	0.014	0.010	0.015
	(0.009)	(0.08)	(0.009)	(0.005)
R^2	0.000	0.000	0.000	0.000
Observations	11650	13761	13364	54250

Note: An observation is a submission. Standard errors in parentheses.

A.2 Deriving an Upper Bound on Marginal Costs

In this section, I reproduce the proof presented in Pakes *et al.* (2015), adapted to my notation and setting, to show that $m^U(\theta) \geq 0$. For clarity, I drop the t subscript and focus on a single contest.

First, let $\Delta r_i = -\Delta r_i^*(d_i+1, d_i; \theta)$ and use order-statistic notation to rank participants by ν_i and Δr_i, so that $\nu_{(1)} \leq \nu_{(2)} \leq ... \leq \nu_{(I)}$ and $\Delta r_{(1)} \leq \Delta r_{(2)} \leq ... \leq \Delta r_{(I)}$. Then, define the sets $L = \{i : d_i > 0\}$, $L_\nu = \{i : \nu_i \leq \nu_{(n)}\}$, $U = \{i : \Delta r_i \geq \Delta r_{(n+1)}\}$, and $U_\nu = \{i : \nu_i \leq \nu_{(I-n)}\}$, where I is the total number of participants, and n is the number of entrants. Let the change in expected profits from making $d_i - 1$ to d_i submissions for $i \in L$ be

$$\Delta\pi_i(d_i, d_i - 1) = \Delta r_i(d_i, d_i - 1; \theta) - \nu_i + \omega_{id_i,d_i-1}$$

and similarly, let the change in expected profits from making one additional submission be

$$\Delta\pi_i(d_i + 1, d_i) = \Delta r_i^*(d_i + 1, d_i; \theta) - \nu_i + \omega^*_{id_i+1,d_i}$$

where $\omega^*_{id_i+1,d_i} = \omega_{id_i+1,d_i}$ if $d_i < 5$ and $\omega^*_{id_i+1,d_i} = 0$ otherwise. Then, we have that

$$\begin{aligned}
&\frac{1}{I}\sum_{i\in L}\Delta r_i(d_i, d_i-1;\theta) - \frac{1}{I}\sum_{i\in U}\Delta r_i^*(d_i+1, d_i;\theta) \\
&\qquad \geq \frac{1}{I}\sum_{i\in L}\Delta r_i(d_i, d_i-1;\theta) - \frac{1}{I}\sum_{i\in U_\nu}\Delta r_i^*(d_i+1, d_i;\theta) \\
&\qquad = \frac{1}{I}\sum_{i\in L}\left(E[\Delta\pi_i(d_i, d_i-1)|\mathcal{J}_i] + \nu_i - \omega_{id_i,d_i-1}\right) \\
&\qquad\quad - \frac{1}{I}\sum_{i\in U_\nu}\left(E[\Delta\pi_i(d_i+1, d_i)|\mathcal{J}_i] + \nu_i - \omega^*_{id_i+1,d_i}\right) \\
&\qquad \geq \frac{1}{I}\left(\sum_{i\in L}\nu_i - \sum_{i\in U_\nu}\nu_i\right) - \frac{1}{I}\left(\sum_{i\in L}\omega_{id_i,d_i-1} - \sum_{i\in U_\nu}\omega^*_{id_i+1,d_i}\right),
\end{aligned}$$

where the first inequality follows from the definition of the set U. The second inequality follows from the assumption that participants take the optimal action given their information sets. Note that

$$\frac{1}{I}\left(\sum_{i\in L}\nu_i - \sum_{i\in U_\nu}\nu_i\right) \geq \frac{1}{I}\left(\sum_{i\in L_\nu}\nu_i - \sum_{i\in U_\nu}\nu_i\right) = \frac{1}{I}\left(\sum_{i=1}^{n}\nu_{(i)} - \sum_{i=1}^{I-n}\nu_{(i)}\right).$$

The distributional assumption on ν_i (Assumption 4) ensures that

$$E\left[\frac{1}{I}\left(\sum_{i=1}^{n}\nu_{(i)} - \sum_{i=1}^{I-n}\nu_{(i)}\right)\right] \geq 0.$$

Furthermore,

$$E\left[\frac{1}{I}\sum_{i\in L}\omega_{id_i,d_i-1}\right] = \frac{1}{I}\sum_{i=1}^{I}E\left[1\{d_i>0\}\omega_{id_i,d_i-1}\right] = \frac{1}{I}\sum_{i=1}^{I}E\left[1\{d_i>0\}E[\omega_{id_i,d_i-1}|\mathcal{J}_i]\right] = 0.$$

Expectational errors are mean-zero for entrants because the action d_i is an element of the participant's information set. I also require the following assumption:

Assumption 6 $E\left[\frac{1}{I}\sum_{i\in U_\nu}\omega^*_{id_i+1,d_i}\right] \geq 0.$

In other words, participants in U_ν cannot consistently underestimate their expected marginal returns. Note that this applies only to participants in U_ν with $d_i < 5$, as otherwise, $\omega^*_{id_i+1,d_i} = 0$.

As a result,

$$E\left[\frac{1}{I}\sum_{i\in L}\Delta r_i(d_i, d_i-1;\theta) - \frac{1}{I}\sum_{i\in U}\Delta r^*_i(d_i+1, d_i;\theta)\right] \geq 0.$$

A.3 Counterfactual Simulation Procedure

A.3.1 Complete Information

To simulate counterfactual contest designs under complete information, I draw sample parameters from the identified set and use iterated best response to obtain equilibrium strategies. I make the assumption that first-stage ability estimates are obtained without error. The following steps can be used to obtain counterfactual equilibrium outcomes for a contest t:

1. Uniformly sample θ^s from the identified set of average cost parameters.

2. At the sampled parameter, obtain bounds on the cost draw for each participant. Note that if θ^s were the true parameter, then by revealed preference, $\nu_{it} \geq \Delta r^*_{it}(d_{it}+1, d_{it};\theta^s)$ at the observed submission decisions, where $\Delta r^*_{it}(d_{it}+1, d_{it};\theta^s)$ is evaluated at the sampled

parameter θ^s. Similarly, $\nu_{it} \leq \Delta r_{it}(d_{it}, d_{it}-1; \theta^s)$ if participant i submitted at least once to contest t. Otherwise, I use $\nu_{it} \leq \max_{j=1,...,I_t} \{-\Delta r^*_{jt}(d_{jt}+1, d_{jt}; \theta^s)\}$ as an upper bound. For each participant, obtain a lower bound ν^{Ls}_{it} and an upper bound ν^{Us}_{it}.

3. Uniformly sample ν^s_{it} from the interval $\left[\nu^{Ls}_{it}, \nu^{Us}_{it}\right]$ for each participant to obtain a cost draw that is consistent with the observed behavior and the estimated parameters.

4. Compute equilibrium actions according to the following procedure:

   (a) For each participant $i = 1, ..., I_t$, choose a random starting action $d^s_{it} \in \{0, 1, ..., D\}$, where D is the submission limit.

   (b) Loop through participants, updating participant i's action according to

   $$d^s_{it} = \arg\max_{d_{it}} \left[R_t(d_{it}, d^s_{-it}; X_i, X_{-it}) - (\theta^s_1 + \theta^s_2 d_{it} + \nu^s_{it})d_{it}\right]$$

   for $d_{it} \in \{0, 1, ..., D\}$, where D is the counterfactual submission limit and $R_t(.)$ is the counterfactual contest expected returns function.

   (c) Repeat 4b until the updating procedure no longer changes participant actions. This rest-point is a Nash Equilibrium of the contest game.

5. Calculate contest outcome metric V^s_t at the equilibrium actions, the parameter vector θ^s and the cost draws $\{\nu^s_{it}\}^{I_t}_{i=1}$.

Steps 1-5 are repeated S times. In Figures 1.5-1.7, I report the lower bound on the counterfactual outcome as $V^L_t = \min(V^s_t)$ and the upper bound as $V^U_t = \max(V^s_t)$. To obtain the average outcome across contests, as shown in Table 1.9, I use $V^L = \frac{1}{T}\sum_{t=1}^{T} V^L_t$ for the lower bound and $V^U = \frac{1}{T}\sum_{t=1}^{T} V^U_t$ for the upper bound.

A.3.2 Incomplete Information

The procedure described in Appendix A.3.1 can be modified as follows to incorporate incomplete information. First, in Step 2, all instances of $R_t(d_{it}, d_{-it}; X_i, X_{-it})$ must be replaced with $ER_{it}(d_{it})$, which can be obtained using the procedure described in Section 1.8.1. Then, the resulting cost intervals $\left[\nu^{Ls}_{it}, \nu^{Us}_{it}\right]$ will take into account that participants had incomplete information when choosing their actions. Second, Step 4 must be modified to capture the change in the density of the number of competitors, participant actions, and characteristics when the structure of the contest changes. Formally, Step 4 can be modified as follows, assuming that $\left[\nu_{it}^{Ls}, \nu_{it}^{Us}\right]$ have been obtained for all participants in the contests in $\mathcal{W}$ in a previous step.

4. Compute equilibrium actions according to the following procedure:

   (a) For each participant $i = 1, ..., I_t$, set d_{it}^s to the participant's observed action.

   (b) At iteration $k+1$, loop through participants and contests, updating participant i's action in contest t according to

   $$d_{it}^{sk+1} = \arg\max_{d_{it}} \left[ER_{it}^k(d_{it}) - (\theta_1^s + \theta_2^s d_{it} + \nu_{it}^s)d_{it}\right]$$

   where

   $$ER_{it}^k(d_{it}) = \frac{1}{B}\sum_{b=1}^{B} R_b(d_{it}, d_{-j_b b}^{sk}; X_i, X_{-j_b b})$$

   for $d_{it} \in \{0, 1, ..., D\}$, where D is the counterfactual submission limit, $R_b(.)$ is the counterfactual contest expected returns function, and j_b denotes a random participant in contest b as in Section 1.8.1.

   (c) Repeat 4b until $d_{it}^{sk+1} = d_{it}^{sk}$ for all $t \in \mathcal{W}$ and all i in contest t. This rest point is an equilibrium of the incomplete information contest game.

In general, the procedure will only recover one of many possible equilibria. However, I find that when multiple equilibria do exist, the outcome metrics do not differ significantly across equilibria. Lee and Pakes (2009) obtain similar results in their analysis of counterfactual equilibria in the model of Ishii (2008).

Appendix B

Appendix to Chapter 3

B.1 Proofs of Propositions

Proof of Proposition 1

First, we determine both interior and corner solutions without self-matching. At price p, consumer demand is $D(p) = \min\left(2\frac{v-p}{\theta}, 1\right)$. For an interior optimal price, we have the FOC $\frac{\partial pD(p)}{\partial p}\big|_{p^*} = 0 = -\frac{p}{\theta} + \frac{v-p}{\theta} \implies \hat{p} = \frac{v}{2}$. The condition for an interior solution is $v < \theta$. When we have a corner solution, i.e. under $v > \theta$, the monopolist sets a price of $\hat{p} = v - \frac{\theta}{2}$. In the rest of the proof, we focus on the case when the markets are covered, i.e. $v > \theta$.

The monopolist's profit is determined as follows:

$$\Pi_1^{SM=0} = (1-\beta)\left(\eta p_1^{on} + (1-\eta)p_1^s\right) + \beta(\eta p_1^{on} + (1-\eta)p_1^s),$$
$$\Pi_1^{SM=1} = (1-\beta)\left(\eta p_1^{on} + (1-\eta)\min(p_1^{on}, p_1^s)\right) + \beta(\eta p_1^{on} + (1-\eta)p_1^s),$$

so the demand from all segments is equal to 1.[1]

To solve for prices, the consumer located farthest from the monopolist must be indifferent

[1]To see this, consider the case of a AD consumer. A consumer located at x purchases online if $v - p_1^{on} - \theta\left|x - \frac{1}{2}\right| \geq 0$. In other words, consumers located at $x \geq \frac{1}{2} - \frac{v-p_1^{on}}{\theta}$ for $x \leq \frac{1}{2}$ and at $x \leq \frac{1}{2} + \frac{v-p_1^{on}}{\theta}$ for $x \geq \frac{1}{2}$ purchase and the remainder do not, leading to a total demand of $2\left(\frac{v-p_1^{on}}{\theta}\right)$ for the monopolist from AD consumers. For an interior solution to exist (the market is not completely served) it must be the case that $2\left(\frac{v-p_1^{on}}{\theta}\right) < 1$, or $p_1^{on} > v - \frac{\theta}{2}$. However, solving the optimization problem for the monopolist (maximizing $2\left(\frac{v-p_1^{on}}{\theta}\right)p_1^{on}$) will yield a price of $\frac{v}{2}$, and $\frac{v}{2} > v - \frac{\theta}{2}$ if only if $v < \theta$. A similar logic follows for the other segments. Hence, the condition $v > \theta$ ensures that the monopolist serves the entire market.

between purchasing or not. This yields $v - p_1^s - \frac{\theta}{2} = 0$ and $v - p_1^{on} - \frac{\theta}{2} = 0$ in the case of no self-matching policy. Prices are then $\hat{p}_1^s = \hat{p}_1^{on} = v - \frac{\theta}{2}$. Similarly, for when the retailer self-matches, we solve $v - p_1^s - \frac{\theta}{2} = 0$, $v - \min(p_1^{on}, p_1^s) - \frac{\theta}{2} = 0$ and $v - p_1^{on} - \frac{\theta}{2} = 0$. Regardless of whether or not SD consumers choose to redeem the self-matching policy, the multichannel retailer will set identical prices across channels, equal to those had it not self-matched: $\hat{p}_1^s = \hat{p}_1^{on} = v - \frac{\theta}{2}$. As the profits under the two conditions are equal, the monopolist will always prefer $SM = 0$, which weakly dominates $SM = 1$.

Below, we refer to m as the cost of undertaking a second shopping trip for store-only undecided consumers. We derive bounds on m that ensure the consumer behavior specified in our assumptions.

Proof of Proposition 2

First, we consider each subgame separately. Then, we compare the profits from each subgame to derive the bounds for the equilibrium results. The following constraints need to be imposed:

- $v > \frac{3\theta}{2}$ ensures that all markets are fully covered,
- $\beta < 5/8$ ensures that no retailer sets such a high online price to earn zero demand from decided consumers in the $(1, 0)$ subgame,
- $v < \theta \left(\frac{1}{\beta} + \frac{1}{4(1-\beta)} - \frac{\eta}{36(1-\beta(1-\eta))^2} + \frac{7-11\eta}{36(1-\beta(1-\eta))} + \frac{7}{18} \right)$ ensures that no retailer wants to price exclusively for its captive segment of store-only undecided consumers and forgo all demand for store-only decided consumers,
- $m > v - \frac{3\theta}{2}$ ensures that no store-only undecided consumers switch stores after their first visit, and that no consumer returns home and visits the store a second time to redeem a self-matching policy.

We focus on the case when $\beta > 0$, so that there are at least some undecided consumers.

No Matching - (0,0)

Channel-agnostic consumers will purchase online. Store-only consumers will buy in-store. All consumers will pay the price set in the channel they buy from. The retailers will earn profits

$$\begin{aligned}\Pi_1^{0,0} &= \eta\Phi_1(p_1^{on}, p_2^{on})p_1^{on} + (1-\eta)\left((1-\beta)\Phi_1(p_1^s, p_2^s) + \frac{\beta}{2}\right)p_1^s,\\ \Pi_2^{0,0} &= \eta(1-\Phi_1(p_1^{on}, p_2^{on}))p_2^{on} + (1-\eta)\left((1-\beta)(1-\Phi_1(p_1^s, p_2^s)) + \frac{\beta}{2}\right)p_2^s\end{aligned}$$

We solve for the first order conditions $\frac{\partial \Pi_j^{0,0}}{\partial p_j^s} = 0$ and $\frac{\partial \Pi_j^{0,0}}{\partial p_j^{on}} = 0$ for $j \in \{1, 2\}$ and check for corner solutions. We find an interior solution with equilibrium prices at $\hat{p}_1^{on} = \hat{p}_2^{on} = \theta$ and $\hat{p}_1^s = \hat{p}_2^s = \frac{\theta}{1-\beta}$ for $\frac{v}{\theta} > \frac{1}{2} + \frac{1}{1-\beta}$, and a corner solution in store prices with $\hat{p}_1^s = \hat{p}_2^s = v - \frac{\theta}{2}$ for $\frac{v}{\theta} \leq \frac{1}{2} + \frac{1}{1-\beta}$. The store price is larger than the online price in all cases as retailers have an incentive to price higher for their captive segment of store-only consumers. The binding condition for an interior solution requires that all SU consumers purchase in equilibrium. For retailer 1, this can be written as $v - p_1^s - \frac{\theta}{2} > 0$ (the utility for the SU consumer farthest away from store 1 is greater than zero). When this condition fails (i.e. $\frac{v}{\theta} \leq \frac{1}{2} + \frac{1}{1-\beta}$), we have a corner solution where retailers set local monopoly prices $v - \frac{\theta}{2}$ in-store. No other constraints apply and there are no other corner solutions. The equilibrium profits earned by retailers are

$$\Pi_1^{0,0} = \Pi_2^{0,0} = \begin{cases} \frac{1}{4}\left[2v(1-\eta) - \theta(1-3\eta)\right], & \frac{v}{\theta} \leq \frac{1}{2} + \frac{1}{1-\beta} \\ \frac{\theta}{2}\left(1 + \frac{\beta(1-\eta)}{1-\beta}\right), & \frac{v}{\theta} > \frac{1}{2} + \frac{1}{1-\beta} \end{cases}$$

Symmetric Self-Matching - (1,1)

Channel-agnostic consumers will purchase online and pay the online price. Store-only decided consumers will buy in-store but will redeem the online price of the store they purchase from. Store-only undecided consumers will buy in the store they first visit and will pay the store price. The retailers will earn profits:

$$\Pi_1^{1,1} = (1-\beta(1-\eta))\Phi(p_1^{on}, p_2^{on})p_1^{on} + (1-\eta)\frac{\beta}{2}p_1^s,$$
$$\Pi_2^{1,1} = (1-\beta(1-\eta))(1-\Phi(p_1^{on}, p_2^{on}))p_2^{on} + (1-\eta)\frac{\beta}{2}p_2^s,$$

and set prices $\hat{p}_1^{on} = \hat{p}_2^{on} = \theta$ online and $\hat{p}_1^s = \hat{p}_2^s = v - \theta/2$ in-store. The online price is the familiar competitive price θ and is an interior solution to the first-order conditions $\frac{\partial \Pi_j^{1,1}}{\partial p_j^{on}} = 0$ for $j \in \{1,2\}$. Differentiating with respect to store prices yields $\frac{\partial \Pi_j^{1,1}}{\partial p_j^s} = (1-\eta)\frac{\beta}{2} > 0$, implying a corner solution. The retailers will set the highest store price they can, ensuring all SU consumers purchase, which is $v - \frac{\theta}{2}$. There are no other corner solutions. The equilibrium profits earned by retailers are

$$\Pi_1^{1,1} = \Pi_2^{1,1} = \frac{1}{2}\left[\theta(1-\beta(1-\eta)) + \beta(1-\eta)\left(v - \frac{\theta}{2}\right)\right].$$

Asymmetric Self-Matching - (1,0)

Channel-agnostic consumers will purchase online and pay the online price. Store-only decided consumer will buy in-store but will redeem the online price (as it will be lower) if they buy from the self-matching retailer. They will pay the store price if they buy from the non-self-matching retailer. Store-only undecided consumers will buy from the store they first visit and pay the store price. The retailers profits are then:

$$\Pi_1^{1,0} = \eta\Phi(p_1^{on}, p_2^{on})p_1^{on} + (1-\eta)\left((1-\beta)\Phi(p_1^{on}, p_2^s)p_1^{on} + \frac{\beta}{2}p_1^s\right),$$
$$\Pi_2^{1,0} = \eta(1-\Phi(p_1^{on}, p_2^{on}))p_2^{on} + (1-\eta)\left((1-\beta)(1-\Phi(p_1^{on}, p_2^s)) + \frac{\beta}{2}\right)p_2^s.$$

The first order conditions can be written as:

$$\frac{\partial \Pi_2^{1,0}}{\partial p_2^s} = 0, \frac{\partial \Pi_j^{1,0}}{\partial p_j^{on}} = 0 \text{ for } j \in \{1,2\}, \text{ and } \frac{\partial \Pi_1^{1,0}}{\partial p_1^s} = (1-\eta)\frac{\beta}{2} > 0.$$

In equilibrium, there is an interior solution for online prices and for the store price of retailer 2 and a corner solution for the store price of retailer 1 for large v. the retailers set online prices $\hat{p}_1^{on} = \theta\left(\frac{2}{3} + \frac{1}{3(1-\beta(1-\eta))}\right)$ and $\hat{p}_2^{on} = \theta\left(\frac{5}{6} + \frac{1}{6(1-\beta(1-\eta))}\right)$ and store prices $\hat{p}_1^s = v - \frac{\theta}{2}$ and $\hat{p}_2^s = \hat{p}_2^{on} + \frac{\beta\theta}{2(1-\beta)}$ for $\frac{v}{\theta} > \left(\frac{4}{3} + \frac{1}{6(1-\beta(1-\eta))} + \frac{\beta}{2(1-\beta)}\right)$.

Otherwise, if v is small, we have a corner solution for p_2^s which yields prices $\hat{p}_1^{on} = \theta + \frac{(1-\beta)(1-\eta)(2v-3\theta)}{4(1-\beta(1-\eta))-\eta}$, $\hat{p}_2^{on} = \theta + \frac{(1-\beta)(1-\eta)(2v-3\theta)}{8(1-\beta(1-\eta))-2\eta}$ online and $\hat{p}_1^s = \hat{p}_2^s = v - \frac{\theta}{2}$ in-store. The binding threshold on v for an interior solution requires that all of retailer 2's SU consumers purchase in equilibrium. In other words, $v - p_2^s - \frac{\theta}{2} > 0$. Substituting the interior solution equilibrium store price for retailer 2 into the inequality shows that the corner solution holds for $\frac{v}{\theta} \leq \left(\frac{4}{3} + \frac{1}{6(1-\beta(1-\eta))} + \frac{\beta}{2(1-\beta)}\right)$.

The equilibrium profits earned by retailers are

$$\Pi_1^{1,0} = \frac{v\beta}{2}(1-\eta) + \frac{\theta}{9}\left[4(1-\beta) - \frac{\beta}{4}(1-17\eta) + \frac{1}{2(1-\beta(1-\eta))}\right],$$
$$\Pi_2^{1,0} = \frac{\theta}{72}\left[\frac{9\beta^2(1-\eta)}{1-\beta} + \beta(34-5\eta) + 29(1-\beta) + \frac{7}{1-\beta(1-\eta)}\right],$$

for $\frac{v}{\theta} > \left(\frac{4}{3} + \frac{1}{6(1-\beta(1-\eta))} + \frac{\beta}{2(1-\beta)}\right)$. The expression for $\frac{v}{\theta} \leq \left(\frac{4}{3} + \frac{1}{6(1-\beta(1-\eta))} + \frac{\beta}{2(1-\beta)}\right)$ can be obtained similarly by substituting equilibrium prices into the profit functions and is available upon request.

Equilibrium Analysis

A self-matching configuration is an SPNE if no retailer has the incentive to unilaterally deviate. Equivalently, for $(0,0)$ to be an SPNE, retailer 1 must not have the incentive to deviate to $(1,0)$. For $(1,1)$ to be an SPNE, retailer 2 must not have the incentive to deviate to $(1,0)$. For $(1,0)$ or $(0,1)$ to be an SPNE the self-matching retailer must not prefer $(0,0)$ and the non-self-matching retailer must not prefer $(1,1)$. By comparing the profits at the equilibrium prices defined above, we can construct equilibrium regions.

Let $\beta' = \frac{27-\sqrt{25\eta^2+448\eta+256}}{32(1-\eta)} + \frac{5}{32}$. The results in Proposition 2 focus on the region where $\beta < \beta'$ for clarity of exposition. In this proof, we provide an extended analysis, including the region where $\beta \geq \beta'$. Define

- $z_1 = \min\left(\frac{27-17\beta}{18(1-\beta)} + \frac{\beta}{9(1-\beta)(1-\beta(1-\eta))}, \frac{7}{36(1-\beta(1-\eta))} + \frac{1}{4(1-\beta)} + \frac{25}{18}, \frac{3\eta}{8\beta(1-\eta)-8-\eta} + \frac{5}{2}\right)$,
- $z_2 = \max\left(\frac{27-17\beta}{18(1-\beta)} + \frac{\beta}{9(1-\beta)(1-\beta(1-\eta))}, \frac{7}{36(1-\beta(1-\eta))} + \frac{1}{4(1-\beta)} + \frac{25}{18}\right)$,
- $z_3 = \frac{1}{1-\beta} - \frac{1}{9(1-\beta(1-\eta))} + \frac{17}{18}$.

We calculate equilibrium profits and the applicable thresholds under all subgames. Then, for $\frac{v}{\theta} < z_1$, the incremental profit from self-matching for retailer 1 is positive: $\Pi_1^{1,0} - \Pi_1^{0,0} > 0$, and retailer 2 prefers not to deviate as $\Pi_2^{1,1} - \Pi_2^{1,0} < 0$, so $(1,0)$ and $(0,1)$ are SPNE. For $z_1 < \frac{v}{\theta} < z_2$, $(0,0)$ is the unique equilibrium for $\beta < \beta'$, as $\Pi_1^{1,0} - \Pi_1^{0,0} < 0$ while $\Pi_2^{1,1} - \Pi_2^{1,0} < 0$, and $(1,1)$

is the unique equilibrium for $\beta > \beta'$ as $\Pi_1^{1,0} - \Pi_1^{0,0} > 0$ while $\Pi_2^{1,1} - \Pi_2^{1,0} > 0$. For $z_2 < \frac{v}{\theta} < z_3$, both $(0,0)$ and $(1,1)$ are SPNE as $\Pi_1^{1,0} - \Pi_1^{0,0} < 0$ while $\Pi_2^{1,1} - \Pi_2^{1,0} > 0$, so no retailer prefers to unilaterally deviate from either symmetric setup. For $\frac{v}{\theta} > z_3$, $(1,1)$ is the unique SPNE as $\Pi_1^{1,0} - \Pi_1^{0,0} > 0$ while $\Pi_2^{1,1} - \Pi_2^{1,0} > 0$.

Note that for sufficiently large β, retailers prefer to offer symmetric self-matching policies at intermediate v. This is because as β grows, retailer 1 has an incentive to match for lower v given that retailer 2 also matches. As the critical threshold of v becomes lower, it may cross the threshold at which the other retailer no longer prefers to match, yielding an equilibrium where both retailers match for intermediate v.

To summarize, for low β, as v increases, there will first be an asymmetric solution, then $(0,0)$, then both $(0,0)$ and $(1,1)$, and then uniquely $(1,1)$. For high β, as v increases, there will first be an asymmetric solution, then $(1,1)$, then both $(0,0)$ and $(1,1)$, and then uniquely $(1,1)$.

Proof of Proposition 3

A comparison of profits in the $(1,0)$ subgame reveals that $\Pi_{1,0}^1 > \Pi_{1,0}^2$ everywhere. A comparison of profits earned by retailer 1 in the $(1,1)$ subgame and in the $(0,0)$ subgame reveals that $\Pi_{1,1}^1 > \Pi_{0,0}^1$ if $\frac{v}{\theta} > \frac{3}{2} + \frac{1}{1-\beta} = z_4$, which is strictly greater than z_3.

B.2 Illustration of Consumer Behavior

To better understand the behavior of the different consumer segments enabled by our model, consider the following example.

Mr. Sam Dawkins is an SD consumer shopping for a new suit. He knows exactly the suit type he wants to buy. He has spent a lot of time researching specific products online and identified item numbers at Macy's and JCPenney's website. However, he needs the suit immediately and cannot wait to order online because he has an interview coming up. The products are available at both JCPenney and Macy's stores and he ponders which store he should visit, given the self-matching policy as well as prices.

Mr. Stanley Underwood is an SU consumer shopping for a new suit. He wants to go to the store to touch and feel the products, discover whether he wants a navy or a gray suit, and also measure his sizes. He chooses to go to Macy's since that retailer is generally closer to his preference than the competing JCPenney. Stanley prefers to obtain the suit immediately and cannot afford to make a second trip to the store or wait for the product to be delivered. After he discovers the exact suit he likes at Macy's, he either buys it at the store or chooses not to purchase.

Mr. Art Davos is an AD consumer who knows the suit he wants to buy is gray with pinstripes. He has identified products across both retailers and researches the prices online and in the stores. He is closer in preference to JCPenney, so he values the suits there higher, but will make his final decision by looking at both retailers and channels, with no particular preference for the online or the store channel.

Mr. Andy Urbany is an AU consumer who does not know the particular suit he wants and is not in a hurry to buy. He is closer to Macy's in preference, so he visits the store and discovers that he wants a navy suit with a peak lapel. He then goes home because he expects lower prices online, compares both Macy's and JCPenney for prices, and buys from the website of the retailer that offers him the best deal.

B.3 Proofs for Extensions

Proof of Proposition 4

The proof of Proposition 4 proceeds just as in Proposition 2, except with an extra parameter μ representing fraction of "smart" consumers, or consumers who can costlessly search for online information while in-store. The μ segment will be relevant for store-only undecided consumers, as these will only be able to claim a self-matching policy if they have access to the internet in-store. The remaining store-only undecided consumers will not be able to claim a self-matching policy and will have to pay the store price. We require the following restrictions:

- $v > \frac{3\theta}{2} + \frac{\beta\mu\theta(1-\eta)}{1-\beta(1-\eta)}$ ensures that all markets are fully covered,
- $\beta < 1 - \frac{3}{2(4-\mu)}$ ensures that no retailer sets such a high online price to earn zero demand from decided consumers in the $(1,0)$ subgame,
- $v < \theta\left(\frac{1-\beta}{\beta} + \frac{4\mu+16\eta-4\mu\eta+2}{12\eta} + \frac{3\beta\eta}{12(1-\beta)\eta} - \frac{(1+2\mu)^2(1-\beta)^2(1-\eta)^2}{36\eta(1-\beta(1-\eta))^2} + \frac{(1+2\mu)(1-\beta)(1-\eta)(2\mu+11\eta-2\mu\eta-5)}{36\eta(1-\beta(1-\eta))}\right)$

ensures that no retailer wants to price exclusively for its captive segment of store-only undecided consumers and forego all demand for store-only decided consumers,

- $m > v - \frac{3\theta}{2} + \mu\theta - \frac{\mu\theta}{1-\beta(1-\eta)}$ ensures that no store-only undecided consumers switch stores after their first visit, and that no consumer returns home and visits the store a second time to redeem a self-matching policy.

No Retailers Self-Match - (0,0)

The equilibrium prices under $(0,0)$ emerge just as in Proposition 2, as mobile consumers behave just as the rest of the consumers.

One Retailer Self-Matches - (1,0)

Store-only undecided consumers who are mobile will redeem the self-matching policy if they first visit the store that offers the policy. Profits are

$$\Pi_1^{1,0} = \eta\Phi(p_1^{on}, p_2^{on})p_1^{on} + (1-\eta)\left((1-\beta)\Phi(p_1^{on}, p_2^{s})p_1^{on} + \frac{\beta}{2}((1-\mu)p_1^{s} + \mu p_1^{on})\right),$$
$$\Pi_2^{1,0} = \eta(1-\Phi(p_1^{on}, p_2^{on}))p_2^{on} + (1-\eta)\left((1-\beta)(1-\Phi(p_1^{on}, p_2^{s})) + \frac{\beta}{2}\right)p_2^{s}.$$

In equilibrium, the retailers set online prices $\hat{p}_1^{on} = \theta\left(\frac{1+2\mu}{3(1-\beta(1-\eta))} + \frac{2(1-\mu)}{3}\right)$ and $\hat{p}_2^{on} = \theta\left(\frac{1+2\mu}{6(1-\beta(1-\eta))} + \frac{5-2\mu}{6}\right)$, and store prices $\hat{p}_1^{s} = v - \frac{\theta}{2}$ and $\hat{p}_2^{s} = \hat{p}_2^{on} + \frac{\beta\theta}{2(1-\beta)}$ for $\frac{v}{\theta} > \left(\frac{1+2\mu}{6(1-\beta(1-\eta))} + \frac{1}{2(1-\beta)} - \frac{\mu}{3} + \frac{5}{6}\right)$. Otherwise, $\hat{p}_1^{on} = \frac{v}{2} + \frac{\theta}{4} - \frac{\theta\mu}{2} - \frac{8\theta\mu+9\theta\eta-6v\eta-2\theta\eta\mu}{4(4\beta(1-\eta)+\eta-4)}$, $\hat{p}_2^{on} = \hat{p}_1^{on} - \frac{\theta(\beta(6\eta-4+\mu\eta)+4-3\eta)}{4\beta(1-\eta)+\eta-4}$ and $\hat{p}_1^{s} = \hat{p}_2^{s} = v - \frac{\theta}{2}$. The first order conditions and the binding constraint are just as in the proof for (1,0) in Proposition 2, except for the addition of an extra parameter μ.

Both Retailers Self-Match - (1,1)

The retailers will earn profits

$$\Pi_1^{1,1} = (1-\beta(1-\eta))\Phi(p_1^{on}, p_2^{on})p_1^{on} + (1-\eta)\frac{\beta}{2}((1-\mu)p_1^{s} + \mu p_1^{on}),$$
$$\Pi_2^{1,1} = (1-\beta(1-\eta))(1-\Phi(p_1^{on}, p_2^{on}))p_2^{on} + (1-\eta)\frac{\beta}{2}((1-\mu)p_2^{s} + \mu p_2^{on}).$$

and set prices $p_1^{on} = p_2^{on} = \theta(1-\mu) + \frac{\mu\theta}{1-\beta(1-\eta)}$ online and $p_1^s = p_2^s = v - \frac{\theta}{2}$ in-store.

Equilibrium Analysis

To prove existence of the result, we provide an example with $\eta = \frac{1}{5}$ and $\beta = \frac{1}{3}$. Let

$$y_0 = \frac{2385\mu - 41(529\mu^2 + 1408\mu + 1936)^{\frac{1}{2}} + 7810}{4004}, y_1 = \frac{32\mu + 331}{198} + \frac{4}{11(2+\mu)},$$
$$y_2 = \frac{64\mu + 395}{198} + \frac{3}{88(1-\mu)}, y_3 = \frac{32\mu + 427}{198} + \frac{3}{22(1-\mu)}.$$

Comparing profits when $\frac{v}{\theta} < y_0$, $(1,1)$ is the unique SPNE. For $y_0 < \frac{v}{\theta} < y_1$, $(1,0)$ and $(0,1)$ are SPNE. For $y_1 < \frac{v}{\theta} < y_2$, $(0,0)$ is the unique equilibrium. For $y_2 < \frac{v}{\theta} < y_3$, both $(0,0)$ and $(1,1)$ are SPNE. For $\frac{v}{\theta} > y_3$, $(1,1)$ is the unique SPNE. To prove the associated Proposition, note that y_0, y_1, y_2 and y_3 are all increasing in μ, so that holding constant $\frac{v}{\theta}$, an increase in mobile consumers shrinks the equilibrium region that admits self-matching policies.

Increasing profits with mobile consumers In the $(1,1)$ equilibrium for large v, $\eta = \frac{1}{5}$ and $\beta = \frac{1}{3}$, the retailers' profits are increasing in μ if $\frac{v}{\theta} < \frac{5}{2} + \frac{8\mu}{11}$, which is possible if $\mu < 0.83$. Furthermore, the retailers' profits are larger than when $\mu = 0$ if $\frac{v}{\theta} < 5/2 + 4\mu/11$, which is possible if $\mu < 0.72$. This shows that retailer profits may increase as the fraction of mobile consumers increases.

Proof of Proposition 5

Suppose that a multichannel retailer competes with an online-only e-tailer. Assume $v > 2\theta$ to ensure that all markets are fully covered. Assume $v < 4\theta$ and $\beta > \frac{5}{2} - \frac{3}{2(1-\eta)}$ to ensure that the multichannel retailer has positive online sales. Under $(0,0)$ the retailers earn profits

$$\Pi_1^{0,0} = \eta\Phi(p_1^{on}, p_2^{on})p_1^{on} + (1-\eta)p_1^s, \qquad \Pi_2^{0,0} = \eta(1 - \Phi(p_1^{on}, p_2^{on}))p_2^{on}.$$

Taking the FOCs with respect to the prices, we solve:

$$\frac{\partial \Pi_1^{0,0}}{\partial p_1^{on}} = \eta \left[\Phi(p_1^{on}, p_2^{on}) + p_1^{on}\frac{\partial \Phi(p_1^{on}, p_2^{on})}{\partial p_1^{on}}\right] = 0$$
$$\frac{\partial \Pi_1^{0,0}}{\partial p_1^s} = (1-\eta) > 0, \text{implying a corner solution.}$$

We obtain the corresponding FOCs for retailer 2 and solve for the equilibrium corresponding to the best responses of both retailers. All channel-agnostic consumers will purchase online, whereas store-only consumers will buy from the multichannel retailer's store. The retailers will set competitive prices online $\hat{p}_1^{on} = \hat{p}_2^{on} = \theta$, and retailer 1 will set monopoly price in-store $\hat{p}_1^s = (v - \theta)$. That is, we obtain an interior solution for online pricing, but a corner solution for the store price where the multichannel retailer maximizes profits from all captive SU consumers.

Under the $(1,0)$ subgame of competition between a self-matching multichannel retailer with an e-tailer, retailers earn profits

$$\Pi_1^{1,0} = \eta\Phi(p_1^{on}, p_2^{on})p_1^{on} + (1-\eta)((1-\beta)p_1^{on} + \beta p_1^s),$$
$$\Pi_2^{1,0} = \eta(1 - \Phi(p_1^{on}, p_2^{on}))p_2^{on}.$$

As under $(0,0)$, channel-agnostic consumers will purchase online and store-only consumers will purchase from the multichannel retailer's store. Store-only decided consumers redeem the matching policy and pay the online price, whereas store-only undecided consumers fail to redeem the policy and pay the store price. The retailers will set prices $\hat{p}_1^{on} = \theta + \frac{4\theta(1-\beta)(1-\eta)}{3\eta}$, $\hat{p}_2^{on} = \theta + \frac{2\theta(1-\beta)(1-\eta)}{3\eta}$, $\hat{p}_1^s = v - \theta$ for $v > 2\theta + \frac{4\theta(1-\beta)(1-\eta)}{3\eta}$. Once again, there is an interior solution in online pricing for v sufficiently large and a corner monopoly solution for the store price. The threshold for v is derived from the condition that in equilibrium $p_1^{on} < v - \theta$ for an interior solution. That is, the online price charged by retailer 1 cannot exceed the monopoly price for SD consumers, or equivalently, $v - p_1^{on} - \theta > 0$, ensuring that the SD consumer farthest away from store 1 purchases in equilibrium for the market to remain covered. For $v \leq 2\theta + \frac{4\theta(1-\beta)(1-\eta)}{3\eta}$, this condition fails, and retailers will set prices $\hat{p}_1^{on} = v - \theta$, $\hat{p}_2^{on} = v/2$, $\hat{p}_1^s = v - \theta$ which corresponds to a corner solution.

Now we substitute prices into profits for the appropriate v and identify the parameter ranges for which $\Pi_{1,0}^1 > \Pi_{0,0}^1$ to see when retailer 1 would prefer to self-match. Suppose that $\beta > \frac{7}{4} - \frac{3}{4(1-\eta)}$. Then, $\Pi_{1,0}^1 > \Pi_{0,0}^1$ if $v < z_1 = \theta\left(\frac{7}{3} + \frac{8(1-\beta)(1-\eta)}{9\eta}\right)$. Otherwise, if $\beta \leq \frac{7}{4} - \frac{3}{4(1-\eta)}$, then $\Pi_{1,0}^1 > \Pi_{0,0}^1$ if $v < z_2 = 3\theta$. Hence, there exists a $z_0 = \min\{z_1, z_2\}$, such that for $v < z_0$, the multichannel retailer will prefer to self-match.

Proof of Corollary 3

A comparison of the e-tailer's profits, $\Pi_2^{1,0} - \Pi_2^{0,0}$ reveals that it earns greater profits when the multichannel retailer offers a self-matching policy. To see this, note that the e-tailer's price under $(1,0)$ is $\hat{p}_2^{on} = \theta + \frac{2\theta(1-\beta)(1-\eta)}{3\eta}$ which is greater than θ, the price it would charge under $(0,0)$. Also in $(1,0)$, the e-tailer's price is less than $\hat{p}_1^{on} = \theta + \frac{4\theta(1-\beta)(1-\eta)}{3\eta}$, the online price charged by the multichannel retailer. Under $(1,0)$ the e-tailer sets a higher price and earns a greater fraction of demand than under $(0,0)$. As a result, its profits are greater.

Precisely, suppose $v \leq 2\theta + \frac{4\theta(1-\beta)(1-\eta)}{3\eta}$. Then $\Pi_2^{1,0} = \frac{v^2\eta}{8\theta}$ and $\Pi_2^{0,0} = \frac{\theta\eta}{2}$. The difference $\Pi_2^{1,0} - \Pi_2^{0,0} = \frac{(v^2-4\theta^2)\eta}{8\theta}$ which is positive when $v > 2\theta$, which is the lower bound required for markets to be fully covered. Now, suppose $v > 2\theta + \frac{4\theta(1-\beta)(1-\eta)}{3\eta}$. Then $\Pi_2^{1,0} = \frac{\theta(2\beta(1-\eta)-2-\eta)^2}{18\eta}$ and $\Pi_2^{0,0} = \frac{\theta\eta}{2}$. The difference $\Pi_2^{1,0} - \Pi_2^{0,0} = \frac{2\theta(1-\beta)(1-\eta)(1+2\eta-\beta(1-\eta))}{9\eta}$ is greater than zero whenever $\frac{\beta}{1-\beta} > -\frac{1+2\eta}{3\eta}$, which is always the case as $\beta > 0$. Hence, the e-tailer always makes higher profits when the multichannel retailer matches.

B.4 Additional Analyses

Incomplete Market Coverage

Suppose that retailers are located on a Hotelling line of length $\frac{1}{5}\left(3 + \frac{6v}{\theta}\right)$, such that the distance between retailers is still equal to 1. Consumers located to the left of retailer 1 will purchase only from retailer 1 if their utility from doing so is non-negative, and similarly, consumers located to the right of retailer 2 will only purchase from retailer 2 if their utility from doing so is non-negative. In the analysis that follows, we focus on the case when $\frac{v}{\theta} > \frac{1}{2} + \frac{2}{3-\beta}$, which ensures that retailers compete on the unit segment but maintain monopoly power over consumers outside of the unit segment.

The profit functions of retailers 1 and 2 in the subgame where no retailer self-matches can be written as:

$$\Pi_1^{0,0} = \underbrace{\frac{1}{2}\left(\Phi(p_1^{on}, p_2^{on}) + \frac{v - p_1^{on}}{\theta}\right) p_1^{on}}_{\text{Channel-Agnostic Decided and Undecided}} + \underbrace{\frac{1-\beta}{2}\left(\Phi(p_1^{s}, p_2^{s}) + \frac{v - p_1^{s}}{\theta}\right) p_1^{s}}_{\text{Store-Only Decided}} + \underbrace{\frac{\beta}{2}\left(\frac{1}{2} + \frac{v - p_1^{s}}{\theta}\right) p_1^{s}}_{\text{Store-Only Undecided}},$$

$$\Pi_2^{0,0} = \frac{1}{2}\left(1 - \Phi(p_1^{on}, p_2^{on}) + \frac{v - p_2^{on}}{\theta}\right) p_2^{on} + \frac{1-\beta}{2}\left(1 - \Phi(p_1^{s}, p_2^{s}) + \frac{v - p_2^{s}}{\theta}\right) p_2^{s} + \frac{\beta}{2}\left(\frac{1}{2} + \frac{v - p_2^{s}}{\theta}\right) p_2^{s}.$$

Effectively, each retailer obtains an additional segment of size $\frac{v-p}{\theta}$ in each channel. We find that retailers set prices $\hat{p}_1^{on} = \hat{p}_2^{on} = \frac{2v+\theta}{5}$ in the online channel and $\hat{p}_1^{s} = \hat{p}_2^{s} = \frac{2v+\theta}{5-\beta}$ in-store. In contrast to §3.4.2, both online and store prices now depend on v, reflecting the retailer's incentive to attract consumers located outside of the unit segment.

If both retailers offer self-matching policies, they obtain profits:

$$\Pi_1^{1,1} = \underbrace{\left(1 - \frac{\beta}{2}\right)\left(\Phi_1(p_1^{on}, p_2^{on}) + \frac{v - p_1^{on}}{\theta}\right) p_1^{on}}_{\text{Channel-Agnostic \& Store-Only Decided}} + \underbrace{\frac{\beta}{2}\left(\frac{1}{2} + \frac{v - p_1^{s}}{\theta}\right) p_1^{s}}_{\text{Store-Only Undecided}},$$

$$\Pi_2^{1,1} = \left(1 - \frac{\beta}{2}\right)\left(1 - \Phi_1(p_1^{on}, p_2^{on}) + \frac{v - p_2^{on}}{\theta}\right) p_2^{on} + \frac{\beta}{2}\left(\frac{1}{2} + \frac{v - p_2^{s}}{\theta}\right) p_2^{s},$$

and set prices $\hat{p}_1^{on} = \hat{p}_2^{on} = \frac{2v+\theta}{5}$ online and $\hat{p}_1^{s} = \hat{p}_2^{s} = \frac{1}{2}\left(v + \frac{\theta}{2}\right)$ in-store. As in §3.4.2, the retailers set the same price in the online channel as they would have under no self-matching. The store price is larger than under no self-matching as the retailer can now extract more surplus from its captive segment of store undecided consumers.

Under asymmetric self-matching, retailers obtain profits:

$$\Pi_1^{1,0} = \underbrace{\frac{1}{2}\left(\Phi_1(p_1^{on}, p_2^{on}) + \frac{v - p_1^{on}}{\theta}\right) p_1^{on}}_{\text{Channel-Agnostic}} + \underbrace{\frac{1-\beta}{2}\left(\Phi_1(p_1^{on}, p_2^{s}) + \frac{v - p_1^{on}}{\theta}\right) p_1^{on}}_{\text{Store-Only Decided}} + \underbrace{\frac{\beta}{2}\left(\frac{1}{2} + \frac{v - p_1^{s}}{\theta}\right) p_1^{s}}_{\text{Store-Only Undecided}},$$

$$\Pi_2^{1,0} = \frac{1}{2}\left(1 - \Phi_1(p_1^{on}, p_2^{on}) + \frac{v - p_2^{on}}{\theta}\right) p_2^{on} + \frac{1-\beta}{2}\left(1 - \Phi_1(p_1^{on}, p_2^{s}) + \frac{v - p_2^{s}}{\theta}\right) p_2^{s} + \frac{\beta}{2}\left(\frac{1}{2} + \frac{v - p_2^{s}}{\theta}\right) p_2^{s}.$$

Retailers set online prices

$$\hat{p}_1^{on} = \frac{2(2v+\theta)(3\beta^2 - 17\beta + 21)}{(10-3\beta)(21-11\beta)} \qquad \hat{p}_2^{on} = \frac{(2v+\theta)(13\beta^2 - 69\beta + 84)}{2(10-3\beta)(21-11\beta)},$$

and store prices

$$\hat{p}_1^{s} = \frac{1}{2}\left(v + \frac{\theta}{2}\right) \qquad \hat{p}_2^{s} = \frac{(2v+\theta)(6\beta^2 - 55\beta + 84)}{2(10-3\beta)(21-11\beta)}.$$

In equilibrium, retailers will choose to offer self-matching policies as $\Pi_1^{1,0} > \Pi_1^{0,0}$ and $\Pi_2^{1,1} > \Pi_2^{1,0}$. However, retailers earn lower profits when self-matching than had they not self-matched as $\Pi_1^{1,1} < \Pi_1^{0,0}$.

Retailer Processing Costs

Suppose that the retailer incurs a positive but small cost c when serving consumers who redeem a self-matching policy. We show how this cost can reduce the profitability of a self-matching policy and the size of the parameter region where self-matching can be sustained in equilibrium. For simplicity, we focus on the region where $\frac{v}{\theta} > \frac{1}{2} + \frac{1}{1-\beta}$.

First, consider the case when both retailers offer self-matching policies. Profits can be written as

$$\Pi_1^{1,1} = \underbrace{\eta\Phi_1(p_1^{on}, p_2^{on})p_1^{on}}_{\text{Channel-Agnostic Decided and Undecided}} + \underbrace{(1-\eta)(1-\beta)\Phi_1(p_1^{on}, p_2^{on})(p_1^{on}-c)}_{\text{Store-Only Decided}} + \underbrace{(1-\eta)\frac{\beta}{2}p_1^s}_{\text{Store-Only Undecided}},$$

$$\Pi_2^{1,1} = \eta(1-\Phi_1(p_1^{on}, p_2^{on}))p_2^{on} + (1-\eta)(1-\beta)(1-\Phi_1(p_1^{on}, p_2^{on}))(p_2^{on}-c) + (1-\eta)\frac{\beta}{2}p_2^s.$$

The retailers suffer a cost c from each sale to SD consumers but may offset this by charging a higher online price. We find that retailers set prices $\hat{p}_1^s = \hat{p}_2^s = v - \frac{\theta}{2}$ in-store, as in the main model, but set prices $\hat{p}_1^{on} = \hat{p}_2^{on} = \theta + \frac{c(1-\beta)(1-\eta)}{1-\beta(1-\eta)}$ in the online channel, which are higher that the price of θ set in the main model absent service costs. Channel arbitrage increases and profits from SD consumers fall, but online competition dampening increases. Overall, we find that retailer profits remain unchanged as service costs are passed on to AD and AU consumers.

Now, consider the case when only one of the retailers offers to self-match. Profits can be written as

$$\Pi_1^{1,0} = \underbrace{\eta\Phi_1(p_1^{on}, p_2^{on})p_1^{on}}_{\text{Channel-Agnostic Decided and Undecided}} + \underbrace{(1-\eta)(1-\beta)\Phi_1(p_1^{on}, p_2^{s})(p_1^{on}-c)}_{\text{Store-Only Decided}} + \underbrace{(1-\eta)\frac{\beta}{2}p_1^s}_{\text{Store-Only Undecided}},$$

$$\Pi_2^{1,0} = \eta(1-\Phi_1(p_1^{on}, p_2^{on}))p_2^{on} + (1-\eta)(1-\beta)(1-\Phi_1(p_1^{on}, p_2^{s}))p_2^{s} + (1-\eta)\frac{\beta}{2}p_2^s.$$

The retailers set prices $\hat{p}_1^s = v - \frac{\theta}{2}$, $\hat{p}_2^s = \frac{1}{6}\left(2(\theta+c) + \frac{\theta-2\eta c}{1-\beta(1-\eta)} + \frac{3\theta}{1-\beta}\right)$ in-store and $\hat{p}_1^{on} = \frac{1}{3}\left(2(\theta+c) + \frac{\theta-2\eta c}{1-\beta(1-\eta)}\right)$, $\hat{p}_2^{on} = \frac{1}{6}\left(5\theta + 2c + \frac{\theta-2\eta c}{1-\beta(1-\eta)}\right)$ online. All prices other than $\hat{p}_1^s$ appear higher than had the self-matching retailer not incurred a processing cost for SD consumers. As in the $(1,1)$ subgame, retailer 1 increases its online price to mitigate channel arbitrage. In response, retailer 2 increases both its in-store and online prices.

Note that retailer profits remain unchanged when neither retailer offers a self-matching policy. A comparison of profits reveals that the difference $\Pi_1^{1,0} - \Pi_1^{0,0}$ is larger if $c = 0$, suggesting that in the presence of a processing cost, retailers are less inclined to deviate from symmetric non-matching at a broader range of model parameters. Similarly, the difference $\Pi_2^{1,0} - \Pi_2^{1,1}$ is larger when $c > 0$, suggesting that retailers have less of an incentive to maintain a symmetric self-matching equilibrium. As a result, in the presence of a small but non-zero cost for serving self-matching consumers, we expect the equilibrium region labeled $(0, 0)$ in Figure 3.4 to increase in size.

In certain cases, retailers may not be explicit about their self-matching policies, which may induce consumers to haggle with managers. Haggling results in additional costs for consumers and retailers (Desai and Purohit, 2004). On the consumer side, as long as the costs of haggling are not large enough to discourage consumers from redeeming the self-matching policy altogether, and markets remain fully covered, store-only decided consumers will purchase from their preferred store channel and redeem self-matching policies as in the main model.

Different Structure of Consumer Heterogeneity in a Monopoly Setting

A segment of high value, high travel cost consumers of size δ incurs a cost s_H to visit the store channel and has a product valuation of v_H. A segment of size $1-\delta$ has a zero travel cost and a low product valuation of v_L. Suppose that $v_L > \frac{\theta(2-\delta)}{1-\delta}$ and $v_L + s_H < v_H \leq \frac{\theta}{2} + \frac{2v_L - \theta}{2\delta} + s_H$ so that all markets are fully covered and v_H is not so large that the monopolist chooses to attract only high valuation consumers. We also require that $v_H \leq v_L + 2s_H$ so that store-only consumers with high valuations and travel costs do not return home, collect evidence of a lower online price, and make a second trip to the store to redeem a self-matching policy if the retailer offers to self-match. The monopolist's profits are

$$\Pi_1^{SM=0} = (1-\delta)\left(\eta p_1^s + (1-\eta)p_1^s\right) + \delta(\eta p_1^s + (1-\eta)p_1^s),$$

$$\Pi_1^{SM=1} = (1-\delta)\left(\eta p_1^{on} + (1-\eta)\min(p_1^{on}, p_1^s)\right) + \delta(\eta p_1^{on} + (1-\eta)p_1^s),$$

As in the proof for Proposition 1, the monopolist will set prices so that the consumer located farthest from the monopolist will be indifferent between purchasing or not. If the retailer does not offer to self-match, it sets $\hat{p}_1^s = v_L - \frac{\theta}{2}$ and any online price no less than $\hat{p}_1^s$. All consumers visit the store to discover their preferred product and purchase from the store channel. Consumers with a high product valuation pay the same price as consumers with a low product valuation.

The monopolist does not increase prices to extract surplus from high valuation consumers as v_L is sufficiently large $\left(v_L > \frac{\theta(2-\delta)}{1-\delta}\right)$ to maintain full market coverage for low valuation consumers and v_H is not large enough $\left(v_H \leq \frac{\theta}{2} + \frac{2v_L-\theta}{2\delta} + s_H\right)$ to justify selling only to high valuation consumers. If the retailer chooses to self-match, it sets prices $\hat{p}_1^{on} = v_L - \frac{\theta}{2}$ and $\hat{p}_1^{s} = v_H - s_H - \frac{\theta}{2}$. Store-only high valuation, high travel cost consumers do not redeem the self-matching policy as the cost of making a second trip exceeds the savings. The retailer can extract more surplus from high valuation consumers by self-matching. All markets remain fully covered and the retailer earns higher profits than had it not offered a self-matching policy.

Store Consumers Purchasing Online

Throughout the analysis, we assumed that store consumers have high channel-misfit costs when shopping online, which is why they never shop via that channel. We now relax this assumption so that the store-only consumers face a cost of s when shopping through the online channel and may consider purchasing online if the price is sufficiently low. We restrict attention to the case when these misfit costs are small compared to the product value, so that there is a non-trivial tradeoff.[2]

In this case, the channels become more substitutable in the eyes of store-only consumers. However, retailers can still extract greater surplus from SU consumers by self-matching. We find that both retailers offer self-matching in equilibrium. This result points to the robustness of Proposition 2, as in both cases high product value results in symmetric self-matching policies in equilibrium, i.e., $(1, 1)$.

We require that $v > s + \frac{\theta}{2} + \frac{\theta}{1-\beta(1-\eta)}$ for the analysis of the equilibria below. In all cases, profit expressions as a function of prices are identical to those in the proof for Proposition 2. The only difference is that now v is sufficiently large so that retailers can no longer price at a monopoly level for store consumers, or they may switch to the online channel - which the retailers would find unprofitable given the more competitive nature of the online channel.

[2]With costs above a threshold, in equilibrium consumers would not consider buying online, which could be ensured with a misfit cost close to v.

No Retailers Self-Match - (0,0)

Solutions to the first order conditions yield that in this subgame the two retailers will price $\hat{p}_1^{on} = \hat{p}_2^{on} = \theta$ and $\hat{p}_1^s = \hat{p}_2^s = \frac{\theta}{1-\beta}$ for $s > \frac{\theta\beta}{1-\beta}$ (an interior solution), and $\hat{p}_1^{on} = \hat{p}_2^{on} = \frac{\theta - s(1-\eta)(1-\beta)}{1-\beta(1-\eta)}$ and $\hat{p}_1^s = \hat{p}_2^s = \hat{p}_1^{on} + s$ for $s \leq \frac{\theta\beta}{1-\beta}$ (a corner solution). Note that for $s \leq \frac{\theta\beta}{1-\beta}$ we have a corner solution, where the binding constraint $p_j^s \leq p_j^{on} + s$ that ensures store consumers purchase in-store and not online binds for both retailers $j \in \{1,2\}$. Neither retailer wants store consumers to purchase online as the online channel is more competitive because it does not contain a captive segment of consumers. The equilibrium profits earned by retailers are

$$\Pi_1^{0,0} = \Pi_2^{0,0} = \begin{cases} \frac{s\beta\eta(1-\eta)+\theta}{2(1-\beta(1-\eta))}, s \leq \frac{\theta\beta}{1-\beta} \\ \frac{\theta}{2}\left(1 + \frac{\beta(1-\eta)}{1-\beta}\right), s > \frac{\theta\beta}{1-\beta} \end{cases}$$

Both Retailers Self-Match - (1,1)

Under symmetric self-matching, the retailers will set $\hat{p}_1^{on} = \hat{p}_2^{on} = \frac{\theta}{1-\beta(1-\eta)}$ (interior solution) online and $\hat{p}_1^s = \hat{p}_2^s = \hat{p}_1^{on} + s$ (corner solution) in-store.

The equilibrium profits earned by retailers are

$$\Pi_1^{1,1} = \Pi_2^{1,1} = \frac{1}{2}\left[s\beta(1-\eta) + \frac{\theta}{1-\beta(1-\eta)}\right].$$

One Retailer Self-Matches - (1,0)

In the asymmetric subgame, the retailers will price $\hat{p}_1^{on} = \frac{\theta}{1-\beta(1-\eta)}$, $\hat{p}_2^{on} = \frac{\theta}{2} + \frac{\theta}{2(1-\beta(1-\eta))}$ online (an interior solution), and $\hat{p}_1^s = \hat{p}_1^{on} + s$ (a corner solution), $\hat{p}_2^s = \frac{\theta}{2(1-\beta)} + \frac{\theta}{2(1-\beta(1-\eta))}$ (an interior solution) in-store for $s > \frac{\theta\beta}{2(1-\beta)}$. For $s \leq \frac{\theta\beta}{2(1-\beta)}$, retailer 2 will set $\hat{p}_2^{on} = \frac{\theta - s(1-\eta)(1-\beta)}{1-\beta(1-\eta)}$ and $\hat{p}_2^s = \frac{\theta+\eta s}{1-\beta(1-\eta)}$ (a corner solution). For s sufficiently small, the relevant constraint is $p_2^s \leq p_2^{on} + s$ for retailer 2. The equilibrium profits earned by retailers are

$$\Pi_1^{1,0} = \frac{1}{2}\left[s\beta(1-\eta) + \frac{\theta}{1-\beta(1-\eta)}\right],$$

$$\Pi_2^{1,0} = \begin{cases} \frac{1}{2(1-\beta(1-\eta))}\left[s\beta\eta(1-\eta) + \theta - \frac{s^2(1-\beta)\eta(1-\eta)}{\theta}\right], & s \leq \frac{\theta\beta}{2(1-\beta)} \\ \frac{\theta}{8}\left[1 + \frac{\beta(1-\eta)}{1-\beta} + \frac{3}{1-\beta(1-\eta)}\right], & s > \frac{\theta\beta}{2(1-\beta)} \end{cases}$$

Equilibrium Analysis

A comparison of equilibrium profits yields that $\Pi_1^{1,0} - \Pi_1^{0,0} > 0$ and $\Pi_1^{1,1} - \Pi_1^{0,1} > 0$, so $(1, 1)$ is the unique equilibrium solution.

B.5 Consumer Survey Across Product Categories

We conducted a survey across $N = 499$ individuals in the US using Amazon's Mechanical Turk (mTurk) service to identify the degree of consumer heterogeneity across a wide range of product categories.[3]

Our model and analysis depend on consumer value for a product (v), retailer differentiation (θ), and the dimensions of consumer heterogeneity leading to multiple segments: decided (1) versus undecided (β), and store-only ($1 - \eta$) versus channel-agnostic (η). For each category, consumers were asked whether they were typically decided or undecided when purchasing a product in that category, whether they preferred to purchase in-store or online, and the average value of a product in the category (as measured by price).

Surveying Heterogeneity

We detail how model constructs are operationalized in our survey below.

(1) **Value**: To operationalize the value of the product from a retailer, we asked consumers to estimate how much they spent on a typical single item in this product category. We then computed the weighted average of the responses, taking the mid-point of the range.

(2) **Undecided versus Decided Consumers**: To operationalize the decision-stage discrimination criterion, we computed the proportion of undecided consumers by taking the ratio of individuals who chose "Undecided" to the total number of individuals surveyed.

(3) **Store versus Channel-Agnostic Consumers**: To operationalize consumer preference across channels, we asked consumers whether they would prefer to shop in-store or online for each product category.

[3]We filtered out users who participated but who did not pass a number of validation checks including multiple attention checks and minimum time to complete the survey for a final sample of $N = 430$ individual responses.

Table B.1: *Self-Matching Outcomes*

Market	Outcome
Pet Supply	$(0,1)$
Apparel	$(0,0)$
Department (Low)	$(0,0)$
Department (Upscale)	$(1,1)$
Office Supply	$(1,1)$
Home Improvement	$(1,1)$
Electronics	$(1,1)$

First, in Table B.1, we detail the equilibrium outcomes in the markets studied here. Figure B.1 illustrates a plot of consumer heterogeneity across multiple product categories.In Panel (a), we find the data plotted along the scaled product value (ratio of value to differentiation or $\frac{v}{\theta}$).

There are a few observations that deserve attention here. First, from Panel (a), the proportion of undecided consumers displays a considerable range, from $\approx$ 20% for pet supply products, to almost 60% for apparel. Second, from Panel (b), we find that consumers also display a range of channel preferences, from home improvement, where a large majority prefer to shop in store, to electronics where about the half the consumers prefer to shop in store. Third, when we plot the channel preference versus the proportion of undecided consumers, we find that there is significant variation along both dimensions, with home improvement and pet supply featuring more consumers who prefer store purchases, despite being more decided than undecided. The apparel product category, on the other hand, is characterized by significantly more undecided consumers ($\approx$ 60), and demonstrates more in store preference for shopping.

Based on the survey, we find evidence of significant heterogeneity in consumer preferences and behavior across a wide range of product categories, lending credibility to our model tenets. These findings can help explain why we observe such variation in self-matching pricing policies in the marketplace.

Connection to market outcomes

Next, we discuss how findings from the survey integrated with the insights from our theoretical model are supported by the observed self-matching policies of firms. Here, we focus attention of panel (a), and suggest the reader to refer to Figure 3.4, which illustrates the equilibrium outcomes corresponding to market characteristics.

First, we observe that the Pet Supply market, which we find to have few undecided consumers and low relative value, reflects an asymmetric (1,0) outcome, consistent to what the model predicts. Second, the apparel and low-end department stores, characterized by intermediate relative value and medium to high levels of undecided consumers demonstrates a no-self-matching, or (0,0) equilibrium in practice, again consistent with model predictions. Finally, we examine the markets with all firms self-matching, i.e. electronics, upscale department stores and home improvement stores, and office supply stores. We find (except for office supply products) that they have a high relative value, and intermediate proportion of undecided consumers. However, this might reflect that for office supply firms, they view the primary competition as coming online from Amazon. In such a case, our mixed duopoly model would predict that even low values could support self-matching strategies.

B.6 Survey Questions

Consumers were surveyed about their preferences and decision-making process regarding purchases in the following product categories.[4]

(1) Pet Supply

(2) Office Supply

(3) Electronics

(4) Home Improvement

(5) Apparel

(6) Department Store (Mid-Low End)

(7) Department (Upscale)

The survey questions are detailed below:

[4]We focused on retailers who were mostly focused on specific product categories and avoided broad multi-category retailers.

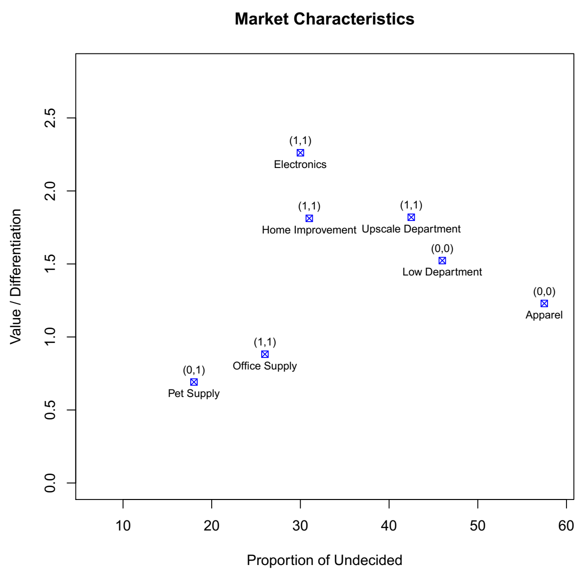

(a) Value / Decision Stage

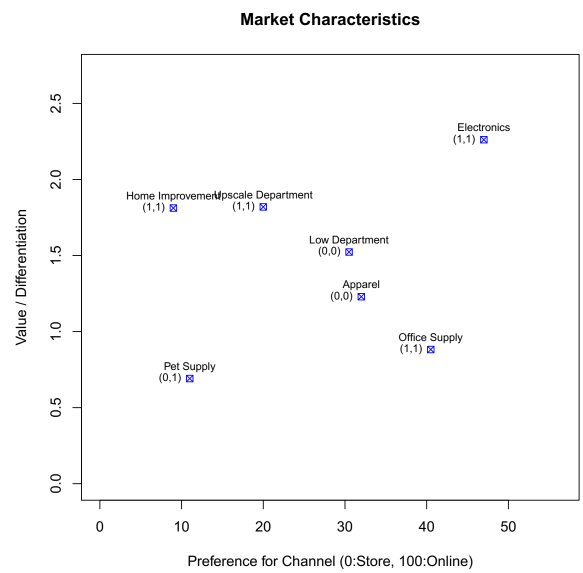

(b) Value / Channel

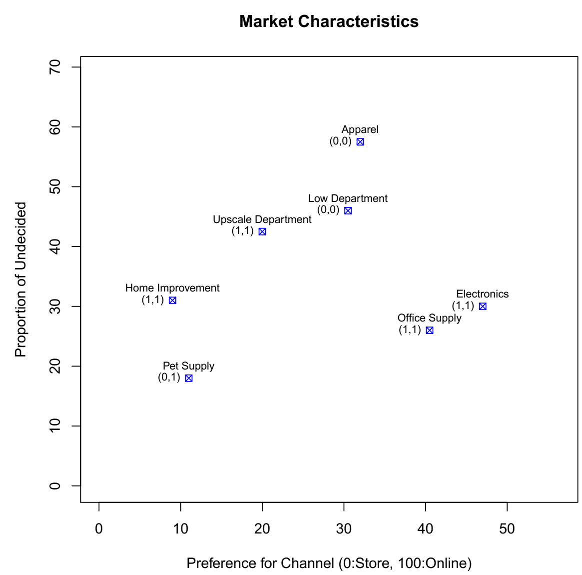

(c) Channel / Decision Stage

Figure B.1: *Consumer Heterogeneity Across Product Categories*

(1) **Online / Undecided** When you are shopping at each of the retail categories below, do you prefer to shop online, offline, or are you equally likely to purchase either online / offline?

Please indicate using the sliders below, where "0" indicates a strong preference for shopping "Offline" (at a physical store) and "100" a strong preference for shopping "Online" and 50 for no preference for either channel.

(2) **Decided / Undecided** Please indicate on a scale of 0-100 how "decided" or "undecided" you are about the specific items you will buy when planning to shop in the following retail categories.

This should reflect your level of uncertainty for the goods you might buy BEFORE you go shopping in each of the store categories below. And where:

(a) Decided: I know the specific product I want to purchase

(b) Neutral

(c) Undecided: I need to visit the physical store to figure out what product I want to purchase

(3) **Product Value / Price**[5] When shopping at each of the retail categories below, think of a typical single item that you might purchase. How expensive is that typical item (range: $0 - $500)? Please indicate the price of a typical SINGLE item in DOLLARS for each of the categories.

B.7 Self-Matching Policies in Practice

Below we list the self-matching policies of several popular retailers. We obtained these from retailers' websites on 1/14/2016 and verified by calling store locations to inquire about matching the website price (if lower).

Self-Matching Retailers

Best Buy: "We match BestBuy.com prices on in-store purchases"

- Also matches online and local competitors.

[5]We use typical price as a measure of product value.

Sears: "If you find a lower price on an identical brand and model number from another Sears branded non-outlet retail format or website, Sears will match that price for up to 7 days after the date of your purchase."

- Also matches online and local competitors.

Staples: "If you purchase an item from Staples and tell us within 14 days that you found that item at a lower price in our stores or at staples.com®, we'll refund the difference."

- Also matches Amazon.com and any retailer who sells products in both retail stores and online under the same brand name.

Office Depot: "If you find a lower price on a new identical item on OfficeDepot.com or OfficeMax.com at the time of purchase or within 14 days of your purchase, show us the lower price and Office Depot or OfficeMax stores will match the price or refund you the difference."

- Also matches Amazon.com and any retailer who sells products in both retail stores and online under the same brand name.

Sports Authority: "For online prices: Provide store associate with mobile device or printout of the competitor's website or sportsauthority.com"

- Also matches online and local competitors.

Toys"R"Us: "We will match Toysrus.com and Babiesrus.com online pricing in our stores"

- Also matches online and local competitors.

Petsmart: Website price will be honored in store.

- Obtained from customer service at 203-937-2749.

Lowe's: Website price will be honored in store.

- Obtained from customer service at 1-800-445-6937.

Home Depot: Website price will be honored in store.

- Obtained from customer service at 1-800-466-3337.

Retailers Who Do Not Self-Match

JCPenney: "All online and mobile pricing, promotions, advertisements, or offers, including from jcp, are excluded from our price matching policy."

- Matches local competitors.

Macy's: "macys.com and Macy's stores operate separately. This means that the products and prices offered at each may be different."

- Does not match competitors.

Urban Outfitters: "While merchandise offered on-line at UrbanOutfitters.com will usually be priced the same as merchandise offered at our affiliate Urban Outfitters stores, in some cases, Urban Outfitters stores may have different prices or promotional events at different times."

- Does not state whether or not it matches competitors.

Petco: "...Petco and Unleashed by Petco stores do not match the prices of unleashedbypetco.com, petco.com or other online sellers and/or websites."

- Does not match competitor websites but does match local competitor stores.