Search for niches or generics? Product designs with heterogeneous search costs

AU Pak Hung, HSU Ke Cheng

July 21, 2024

Abstract

Building on Wolinsky's (1986) sequential search model with differentiated products, this paper introduces search cost heterogeneity to analyze how firms leverage observable product designs to influence consumer search behavior. In equilibrium, consumers' search patterns form a partitioned structure: those with low search costs tend to seek niche products, while consumers facing higher search costs gravitate toward standardized products. In the case of two product designs, a reduction in search costs may lead to higher prices for both if the endogenous shift in consumer search modes—the composition effect—dominates the search-encouraging effect. Additionally, we demonstrate that prices respond asymmetrically to decreases in search costs based on the relative position within the distribution where the shock occurs. Specifically, improvements in search technology benefiting consumers with initially low search costs tend to have anti-competitive effects, whereas those benefiting higher-cost searchers can be pro-competitive.

1 Introduction

Improvements in search technology have shifted aggregate consumption toward specialty products, a pattern known as the long-tail effect (Brynjolfsson et al., 2006; Anderson, 2006). As Anderson (2006) observes, the Pareto principle has increasingly given way to a "98 percent rule," under which nearly every available product finds buyers. In his words, "Apple said that every one of the 1 million tracks in iTunes had sold at least once... Netflix reckoned that 95 percent of its 25,000 DVDs are rented at least once a quarter." This negative relationship between search friction and sales concentration arises from the interaction between demand and supply. On the demand side, better search tools enable consumers to explore, evaluate, and purchase specialized products that would otherwise be difficult to find. On the supply side, anticipating this expanded demand for niches, firms find it worthwhile to offer products tailored to narrower tastes, which broadens the assortment available in the market. Controlling for these supply-side changes in product assortment, Brynjolfsson et al. (2011) show that sales remain less concentrated when search costs fall, indicating that lower search costs also directly shift consumers' search behavior toward niche products.

This paper models that demand channel and asks whether the negative relationship between search costs and product nicheness holds at the individual level, that is, whether consumers with lower search costs search for niche products while consumers with higher search costs search for generic products. We then study how improvements in search technology affect consumers' choices between niches and generics and the resulting equilibrium prices. In our model, ex-ante identical firms choose product designs ranked in dispersive order, and consumers differ in their search costs. A generic design induces a concentrated match-value distribution, whereas a niche design induces a dispersed distribution with a fatter right tail. Product designs are observable before search begins, so firms can use design choices to affect consumers' search decisions, and consumers can choose which design segment to search rather than sampling randomly across the entire market. This observability is natural in practice. Before entering Starbucks to try a seasonal release such as a salted caramel latte, a customer knows she will get something familiar and broadly acceptable, even if it is not her favorite. By contrast, a buyer visiting a specialty coffee shop for Kopi Luwak knows she may encounter a polarizing flavor that is either a great match or a poor one. On an e-commerce platform, shoppers looking for mainstream products can sort listings by sales volume before reading item descriptions, while shoppers looking for niche products can filter by a specific feature.

Because consumers' search strategies depend on firms' design and pricing choices while firms' design choices depend on how consumers with different search costs sort across segments, the two sides of the market are closely intertwined. Even so, the model delivers a clean equilibrium characterization and sharp comparative statics. When consumers' search costs are sufficiently varied, multiple product designs arise endogenously in equilibrium, and firms offering more niche designs charge higher prices. Consumers' search strategies take a monotone partitional form. Consumers in the lowest search-cost interval search the most niche segment, those in the next lowest interval search the next niche segment, and consumers with higher search costs search progressively broader segments. Intuitively, low-search-cost consumers can afford to be picky and keep searching until they find a high match value, so they are drawn to the thicker right tail of a niche design, whereas high-search-cost consumers mainly want to avoid a poor match and prefer the concentrated distribution of a standardized product. To test this prediction, we use more than ten million user-item observations from Alibaba's Tianchi platform. Using a shopper's total number of clicks as

an inverse proxy for search cost and an item's total clicks across all users as a measure of popularity, we find that more active searchers tend to click on less mainstream items, consistent with our partitional sorting result.

In our comparative statics, we show that a reduction in search costs can cause firms across all product designs to raise their prices. To see why, we decompose the effect of lower search costs into an intensive margin and an extensive margin. On the intensive margin is the standard search-encouraging effect. Holding the set of shoppers in a segment fixed, lower search costs make consumers more demanding and lengthen search, which intensifies price competition and pushes prices down. On the extensive margin is a compositional effect that works in the opposite direction. At a lower search cost, the pickiest consumers in a broader segment switch to a neighboring, more niche segment. Because these switching consumers had the lowest search costs in the broad segment they leave, but have the highest search costs in the niche segment they join, the average search cost rises in both segments. Facing a customer pool with a higher average search cost, every firm, whether offering a broad or a niche design, has an incentive to raise its price. We identify sufficient conditions under which the compositional effect on the extensive margin outweighs the search-encouraging effect on the intensive margin, so that a search-cost reduction raises prices across the market.

This decomposition also explains why prices respond differently depending on whether search costs fall for low-cost or high-cost shoppers. Consider a retail market in which tech-savvy consumers browse online with low search costs, whereas less tech-savvy consumers visit brick-and-mortar stores with high search costs. Suppose search costs fall only for the low-cost group, for instance because e-commerce platforms introduce better attribute filters, price-comparison plugins, or AI search tools. Lower search costs in the niche segment make niche products more attractive and draw the lowest-cost shoppers away from the standardized segment. Standardized firms see no direct reduction in the search costs of their remaining customers, yet they lose their most price-sensitive shoppers, so the average search cost in the standardized segment rises and standardized prices go up. Niche prices can also rise if enough higher-cost shoppers switch into the niche segment. By contrast, suppose that search costs fall only for the high-cost group, for instance through clearer in-store signage, better shelf organization, or improved transit access to physical stores. Competition intensifies directly in the standardized segment and lowers standardized prices. That price drop attracts intermediate-cost shoppers away from the niche segment and into the standardized segment. Here the switching of shoppers reinforces competition in both markets. The niche segment loses its highest-cost shoppers, while the standardized segment gains shoppers who have lower search costs than its existing customers, so average search costs and equilibrium prices fall in both segments. As a result, search-cost reductions that benefit low-cost shoppers tend to be anti-competitive, whereas those that benefit high-cost shoppers are pro-competitive.

Literature Review

Our paper builds on the sequential search model with differentiated products pioneered by Wolinsky (1986) and Anderson and Renault (1999). Within this framework, Bar-Isaac et al. (2012) and Larson (2013) endogenize product design by allowing firms to choose parameters that govern the dispersion of consumers' taste shocks.[1] Bar-Isaac et al. (2012) show that equilibrium prices and profits can rise as search costs fall because firms shift endogenously from broad to niche designs, and the resulting increase in product dispersion

[1] Similar insights appear earlier in Kuksov (2004), where product design is modeled as location choice, and Cachon et al. (2008), where market expansion from a broader assortment offsets stronger price competition under lower search costs.

outweighs the search-encouraging effect. Their work also accounts for the coexistence of the superstar and long-tail effects as search costs decline, since longer search durations benefit high-quality firms while intermediate-quality firms switch from broad to niche designs.

Although both papers study product design, our model differs from Bar-Isaac et al. (2012) in two main ways. First, product designs in our model are observable before search, allowing consumers to narrow their consideration sets and conduct design-directed search. Second, firms in our model are ex-ante symmetric rather than vertically differentiated, and consumers vary in their search costs. Firms therefore adopt different designs in equilibrium not because of ex-ante quality differences, but because they target different search-cost segments. In technical terms, a firm's price and design choices depend on the position of its marginal consumers in the valuation distribution, and differences in that position are driven by search-cost heterogeneity rather than quality heterogeneity.

Models of search-cost heterogeneity date back to Varian (1980) and Stahl (1989), in which firms trade off competing for low-cost searchers against exploiting high-cost searchers. More recently, Moraga-González, Sándor, and Wildenbeest (2017) incorporate arbitrary search-cost distributions into a differentiated-products search model with a single product design. In their setting, lower search costs affect prices through two channels. Existing shoppers search more intensively, which pushes prices down, while high-cost consumers who previously stayed out of the market begin to search, which makes average demand less elastic and pushes prices up. When the entry of high-cost shoppers dominates, lower search costs raise the equilibrium price. Our paper extends this insight by endogenizing firms' product designs. Even when the market is fully covered, consumers can switch between generic and niche segments as search costs change. This switching across designs shifts the composition of shoppers in every segment and allows the extensive margin to operate across the full spectrum of equilibrium product designs.

To the best of our knowledge, our paper is the first theoretical model to combine endogenous product design with consumer search-cost heterogeneity. The remainder of the paper is organized as follows. Section 2 introduces the model setup. Section 3 characterizes consumers' and firms' optimal strategies and the equilibrium profile. Section 4 studies the comparative statics of lower search costs. Section 5 concludes, and Section 6 provides parametric examples.

2 Model Setup

There is a continuum of firms and consumers, each with a unit measure. Each consumer has a unit demand for the products. Denote the set of firms by $\mathcal{J}$. Each firm $j \in \mathcal{J}$ chooses its own product design with $s \in \mathcal{S}$, where $\mathcal{S}$ is the set of available designs. The choice of product design determines the distribution of match values that the consumer can derive by consuming the firms' products. Denote by $v_{s,j}$ the match value that an arbitrary consumer can derive by consuming the product of firm j which has adopted design s. It is a random variable with distribution denoted by $F_{s,j}(v_{s,j})$. For each $s \in \mathcal{S}$, let $f_{s,j}$, $[\underline{v}_{s,j}, \bar{v}_{s,j}]$ be the corresponding density and support of product j. We assume the match value is independent across firms, and the distribution $F_{s,j}$ is continuous and twice differentiable. Assume further that the available designs in $\mathcal{S}$ all share the same mean but can be ranked in dispersive order, which is formally defined below.

Definition 2.1. Let X and Y be two random variables with distribution functions F and G, respectively. Let F^{-1} and G^{-1} be the right continuous inverses of F and G respectively,

such that $F^{-1}(\beta) - F^{-1}(\alpha) \leq G^{-1}(\beta) - G^{-1}(\alpha)$ whenever $0 < \alpha \leq \beta < 1$. Then X is said to be smaller than Y in the dispersive order (denoted as $X \leq_{\text{disp}} Y$).

Roughly speaking, a random variable is more dispersive if it has a spread-out density and appeals to consumers with extreme tastes. Dispersive order is commonly used in the random utility framework in the literature. For example, Zhou (2017) also uses this notion to study how valuation dispersion affects price in a frictionless environment, while Choi et al. (2018) show that price increases as the effective value become more dispersed in the dispersive order sense.

Assumption 1. $(\{F_{s,j}|s \in \mathcal{S}\}, \leq_{\text{disp}})$ is a finite totally ordered set with cardinality N, and the corresponding random variables have the same mean μ.

With equal means, dispersive order implies the rotation ordering introduced in Johnson and Myatt (2006), which requires CDF to intersect only once, meaning our design space is more restrictive than Isaac et al.(2012).[2] The above assumption allows the ranking between any pair of product designs to be well-defined so that we can map the design space to the real number. Formally, $s > s'$ if and only if $F_{s,j} \leq_{\text{disp}} F_{s',j}$. With this formulation, a low s corresponds to niches – products that are likely to be either very favorable or unappealing to someone. In contrast, a high s corresponds to broad products, which are inoffensive but not particularly attractive.

Apart from the design, the firm sets a price $P_{s,j}$ to maximize its profit, given consumers' endogenous search behavior and other firms' strategies. Firms announce their designs but not prices to consumers right after the decisions are made. This simplification implies that prices affect consumers' purchasing decisions but not their search behaviors, allowing us to focus on the effect of product designs on directing consumers' searches. Firms have commitment power in the designs; that is, the firm must follow the announced product design.[3] This assumption can be justified by noting that product characteristics are often fixed in the short run.

Consumers' types are characterized by search costs c, which are distributed according to distribution $G(c)$ over the support $[\underline{c}, \bar{c}]$. Let $g(c)$ be the density function of c. The consumer endogenously chooses the optimal search strategy based on her own private search costs c, the product designs committed by the firms, as well as her beliefs in the expected prices. If a consumer with search cost c buys from firm k, her overall payoff is given by

$$v_{s(k),k} - P_{s(k),k} - \sum_{j} c1_{\{\text{ firm j being searched}\}}, \tag{1}$$

where $s(k)$ is the product design chosen by firm k. Contingent on their type c, consumers acquire information sequentially and stop whenever the product obtained so far is sufficiently satisfactory.

The timeline of the model is as follows:

[2]Because $\theta X \leq_{\text{disp}} X$ for any random variable X and $0 < \theta \leq 1$ by Theorem 3.B.4 in Shaked and Shanthikumar (2007), the product design space in our model is richer than Larsen (2013).

[3]We distinguish between announced and observed designs because the firm can manipulate the announcement to attract more consumers while implementing a different one to maximize its profit. Specifically, the announced design determines the type of consumers who will visit, while the actual design decides how likely they will stop and make a purchase. With commitment, this potential inconsistency disappears, and the two can be used interchangeably. We do not discuss this possibility because, in the setting of heterogeneous search costs, marginal consumers span a range rather than a point in the valuation distribution, making it exceedingly difficult to analyze the effect of a design change.

1. Firms simultaneously select their product designs $F_{s,j}$ and price $P_{s,j}$ for all $j \in \mathcal{J}$ and announce their designs to consumers.

2. Based on the announced designs and the beliefs in prices, consumers sequentially incur search costs c to examine each product's price and the realization of its match value. They purchase and exit the market if they identify a satisfactory product.

3. Consumers collect their payoffs according to (1), while firm j's payoff equals the price times the mass of consumers eventually purchasing from it.

We focus on the pure-strategy symmetric perfect bayesian nash equilibrium (SPBE), where all firms selecting the same design adopt the same price. We refer to the firms using design s as s type firms and omit the subscript j whenever the statement is clear enough.

The equilibrium consists of: (i) the consumer's search strategy, (ii) the product design profile $\mathbf{F}^* \equiv \{F^*_{s,j} : \mid j \in \mathcal{J}\}$, (iii) the price profile $\mathbf{P}^* \equiv \{P^*_{s,j} \mid j \in \mathcal{J}\}$, and (iv)consumers' beliefs in prices $\mathbf{P^e} \equiv \{P^e_s \mid s \in \mathcal{S}\}$ for every design including those not used in equilibrium, such that given the correct belief in the equilibrium prices and product designs, consumers choose the optimal search strategy, and firms choose the optimal prices and product designs considering an endogenous shift of consumers' search behavior. As a consequence of the consistency requirement, $P^e_s = P^*_s$ for all $s \in \mathcal{S}^*$. [4] The following assumptions, which are standard in the sequential search literature, are maintained throughout.

Assumptions

- The market is fully covered.[5]
- Consumers must search before they can make purchases.
- Consumers have free recalls.
- Consumers hold passive off-path beliefs. If the consumer observes nonequilibrium price quotes, she maintains the belief that the rest of firms still follow the equilibrium strategies.

To guarantee the existence of the equilibrium, we maintain the following assumption throughout.

Assumption 2. For all $s \in \mathcal{S}$, $f_s(v_s)$ is log-concave.

The assumption implies that F_s, and $1 - F_s$ are log-concave. It also implies that the hazard rate is increasing, so the demand from consumers with higher valuations is more elastic.

[4] The beliefs for prices of off-path designs are more subtle and will be discussed in the section on equilibrium characterization.

[5] Here we abstract away from the participation constraint, which is the main focus of Moraga-González et al to make the key intuition transparent. The assumption is satisfied if the based value of the product is high enough. From a technical standpoint, shifting the valuation distribution parallelly upward does not affect the equilibrium price, provided the elasticity stays the same. However, such a change increases the reservation utility and improves the consumer surplus.

3 Equilibrium characterization

3.1 Consumers' search strategy

According to Weitzman (1979), the optimal search strategy is governed by the reservation utility, which we define as follows. For $s \in \mathcal{S}$, the reservation utility $\hat{x}$ is the unique value satisfying

$$c = \int_{v_s^*}^{\bar{v}_s} (v_s - P_s^e - \hat{x}(c)) d\, F_s(v_s).$$

Alternatively, we can rewrite it as $v_s^*(c) - P_s^e$, where

$$c = \int_{v_s^*}^{\bar{v}_s} (v_s - v_s^*) d\, F_s(v_s) = \int_{v_s^*}^{\bar{v}_s} \left[1 - F_s(v_s)\right], dv_s. \tag{2}$$

Consumers visit stores sequentially in descending order of reservation utility and stop until they find a realized match value exceeding the reservation utility of all the unvisited firms. As there are infinitely many ex-ante symmetric s type firms, the consumer's search problem is stationary.[6] Specifically, a consumer stops searching whenever the match value exceeds the reservation utility. When the consumer does not find an acceptable product and continues searching, she will find herself in the same situation as at the outset.

Based on the observation above, if the consumer chooses to search s type firm in the first place, he will continue visiting other s type firms until he finds a satisfactory product. Therefore, consumers' search behavior is partially directed and partially random. At the very beginning, the consumer decides which type of firms to look into and searches randomly once they have entered a specific market.[7] For ease of exposition, we refer to consumers who constantly search s type firms as s type shoppers and say that they adopt search mode s. Therefore, the consumer's search strategy involves a mapping from her private type c to the search modes.

In determining which designs to look for, a consumer compares the expected consumer surplus of different designs. As shown in Moraga-González et al. (2017), the expected consumer surplus is simply the difference between the consumer's reservation utility and price paid. Interestingly, this is valid whether the consumer engages in active search or not.[8]

Lemma 3.1. *The expected consumer surplus for the s type shopper with search cost c is given by*

$$v_s^*(c) - P_s^e \tag{3}$$

[6] Here, it is possible to just focus on firm types with non-zero-measure because the mass of consumers visiting zero-measure firms must have a zero measure. Otherwise, those zero-measure firms will obtain infinite demand and profits, leading to a contradiction. Therefore, consumers searching for more than one design after exhausting a particular type of firm, and those with consumer surplus differ than lemma 3.1 can only have measure zero.

[7] One caveat is that the above argument relies on the passive off-path belief assumption. Suppose the consumer seeking niche products receives a series of poor valuation realizations at the first few searches and updates her posterior according to the experience; she is likely to revise the consumer surplus downward and redirect her attention to broad products instead.

[8] The consumer does not engage in active search whenever his search cost is so high that $v_s^*(c)$ goes under $\underline{v_s}$ in the definition (2). In this case, the consumer is still willing to visit a firm if the expected payoff of making a single firm visit is nonnegative. Because these consumers search only once, the expected consumer surplus equals the average match value minus the one-time search cost and the price. It can be easily shown the first two parts add up to $v_s^*(c)$. Intuitively, these consumers purchase because of "the based value of the product."

As niche products are more differentiated and have a thicker tail, they appeal more to consumers with lower search costs, who have their reservation utility located at the right tail of the distribution. Formally, a consumer with search cost c chooses to be a s' type shopper if and only if for all $s \in \mathcal{S}^*$,

$$v^*_{s'}(c) - v^*_s(c) \geq P^e_{s'} - P^e_s. \tag{4}$$

The L.H.S. is the difference of reservation utility between search mode s and s' considering the search costs and match values, while the R.H.S. is the price difference based on equilibrium belief. The lemma below shows that whenever $s' < s$, the L.H.S. of inequality (4) is a non-increasing function of c, so consumers with search costs in the lower range prefer to search niche products and consumers with search costs in the high range prefer to search standardized products.

Lemma 3.2. *For any $s', s \in \mathcal{S}$ with $s' < s$, there exists a unique $\hat{c}$ such that consumers with $c \leq \hat{c}$ weakly prefer to be s' type shoppers and consumers with $c > \hat{c}$ strictly prefer to be s type shoppers. Moreover, the set of consumers indifferent between s' and s is a compact- and convex-valued set, in the form of either $[\underline{c}, c_1]$ or $[\tilde{c}_{s'}, \bar{c}]$.*[9]

Proof. Define $\tilde{c_s} = \int_{\underline{v}_s}^{\bar{v}_s} [1 - F_s(u)]\, du = \mu - \underline{v}_s$ for $s \in \mathcal{S}$. Consumers with $c > \tilde{c_s}$ who specialize in searching type-s firms have $v^*_s(c)$ below $\underline{v}_s$ and will stop and purchase from the first visited firm. It follows from (2) that for $c > \tilde{c_s}$, $v^*_s(c) = \mu - c$. Differentiating (2) with respect to c gives

$$\frac{dv^*_s(c)}{dc} = \frac{-1}{1 - F_s(v^*_s(c))} < 0.$$

We first show that the LHS of inequality (4) is a non-increasing function in c. To this end, it suffices to show that $F_{s'}(v^*_{s'}(c)) \geq F_s(v^*_s(c))$ whenever $s' < s$. As $F_s \leq_{\text{disp}} F_{s'}$ implies that the support of F_s is a strict subset of $F_{s'}$, we have $\underline{v}_s > \underline{v}_{s'}$, so $\tilde{c}_{s'} > \tilde{c_s}$.[10] For $c > \tilde{c_s}$, $F_{s'}(v^*_{s'}(c)) \geq F_s(v^*_s(c))$ clearly holds, so we focus on the case $c \leq \tilde{c_s}$ below.
Equation (2) can be rewritten as

$$c = \int_{F_s(v^*_s(c))}^1 \frac{1-u}{f_s(F_s^{-1}(u))}\, du.$$

Note that $F_s \leq_{\text{disp}} F_{s'}$ implies $f_{s'}(F_{s'}^{-1}(u)) \leq f_s(F_s^{-1}(u))$ for all $u \in [0,1]$. Therefore,

$$c \leq \int_{F_s(v^*_s(c))}^1 \frac{1-u}{f_{s'}(F_{s'}^{-1}(u))}\, du,$$

implying that $F_{s'}(v^*_{s'}(c)) \geq F_s(v^*_s(c))$ as desired.

The set of indifferent consumers is compact and convex-valued comes from that $v^*_{s'}(c) - v^*_s(c)$ is a non-increasing continuous function. Next, we show that if (4) holds with equality

[9] A few caveats are in order. When the statement in the lemma changes to a weak incentive, there could be multiple cutoff points. Second, we assume all consumers play pure strategies except those on the cutoffs, so even though the measure of consumers that are indifferent between two neighboring designs can be positive, the measure of consumers mixing between two search modes is zero. The third caveat is that if all consumers prefer design s over design s' as in the case where $P^e_{s'}$ is much larger than P^e_s, $\hat{c}(s', s) = \underline{c}$. Similarly, $\hat{c}(s', s) = \bar{c}$ if all consumers prefer a more niche design. One should be aware that the cutoff type could strictly prefer one design over another when the above scenarios happen.

[10] This follows from theorem 3.B.15 in Shaked and Shanthikumar (2007).

over a non-degenerate interval of search costs lower than $\tilde{c}_s$, say $[c_0, c_1]$, then it is necessary that $c_0 = \underline{c}$. For (4) to hold with equality, it is necessary that $F_{s'}(v^*_{s'}(c)) = F_s(v^*_s(c))$ for all $c = [c_0, c_1]$. In particular, $F_{s'}(v^*_{s'}(c_0)) = F_s(v^*_s(c_0))$. As

$$c_0 = \int_{F_s(v^*_s(c_0))}^{1} \frac{1-u}{f_s(F_s^{-1}(u))}\, du = \int_{F_s(v^*_s(c_0))}^{1} \frac{1-u}{f_{s'}(F_{s'}^{-1}(u))}\, du,$$

and $f_{s'}(F_{s'}^{-1}(u)) \leq f_s(F_s^{-1}(u))$ for all $u \in [F_s(v^*_s(c_0)), 1]$, it is necessary that $f_s(F_s^{-1}(u)) = f_{s'}(F_{s'}^{-1}(u))$ for all $u \in [F_s(v^*_s(c_0)), 1]$. Consequently, $F_{s'}(v^*_{s'}(c)) = F_s(v^*_s(c))$ for all $c \in [\underline{c}, c_1]$. Finally, as $\frac{dv^*_{s'}(c)}{dc} = \frac{dv^*_s(c)}{dc}$ for $c \in [\tilde{c}_{s'}, \bar{c}]$, if the indifferent set includes any point in $[\tilde{c}_{s'}, \bar{c}]$, it must include the whole interval.

□

Intuitively, consumers with high search costs cannot afford to be picky, so their main objective is to avoid a poor deal. As the variance of generics is less than that of niches, they prefer to be conservative and search for the former. On the other hand, consumers with low search costs prefer niche products because only the right tail of the distribution is relevant for them, and niche product is associated with a thicker right rail. Another implication of the above lemma is that the average number of searches is higher for niche designs because $\frac{1}{1-F_{s'}(v^*_{s'}(c))} \geq \frac{1}{1-F_s(v^*_s(c))}$, where $1 - F_{s'}(v^*_{s'}(c))$ is the probability of finding a satisfactory niche product. Will a higher mixing and matching benefit of exploring specialties outweigh the loss of long-duration searches? The answer is positive. Combining Lemma 3.2 with the fact that $v_{s'}(c) = v_s(c)$ for all $c \geq \tilde{c}_{s'}$, we have $v^*_{s'}(c) \geq v^*_s(c)$ for any $s' < s$, suggesting match benefit net of search costs is higher for niches. This result echos the impact of volatility on option value in the sense that niche design makes the good match better and the bad match worse, but the consumer can avoid the poor match by continuing to search. Hence, high dispersion should improve the consumer's reservation utility weakly, even accounting for the search effort.

Corollary 3.3. , $v^*_{s'}(c) \geq v^*_s(c)$.

Empirically, we use real user-commodities behavior data from season one of Ali Mobile Recommendation Algorithm Competition to test our theory. The dataset comes from Tianchi, a big data platform funded by Alibaba, and includes more than ten million user-item pairs with behavior (click, collect,add-to-cart, and payment) recorded. For our purpose, we approximate consumers' search costs by the average number of clicks and use popularity, how many times a certain item is clicked by all users, as a proxy for the nicheness of a design. Our result shows that consumers who conduct more searches click on less popular items on average, confirming lemma 3.2. The Spearman correlation between the number of clicks of a user and the average popularity of the products she clicks on is -0.1684; this negative correlation is illustrated by figure 1, with each point on the graph representing a single user.

We now generalize the insight to the markets with multiple designs and restrict our discussion to monotone interval partition thereafter. Denote the set of designs announced by firms at the first stage by $\hat{\mathcal{S}}$ and relabel its elements such that $s_1 < s_2 < \cdots < s_n$, where n is the cardinality of set $\hat{\mathcal{S}}$. The lemma/corollary above implies that fixing a price vector $\mathbf{P}^e_{\hat{\mathcal{S}}} = \left\{P^e_{s_1}, P^e_{s_2}, ...P^e_{s_n}\right\}$, one of the optimal search decision takes a monotone interval partition form, given by $\{[a_0, a_1), [a_1, a_2), ..., [a_{n-2}, a_{n-1}), [a_{n-1}, a_n]\}$, where $[a_0, a_1)$ is the interval of search costs that find design s_1 (weakly) optimal, $[a_1, a_2)$ is the interval of search costs that

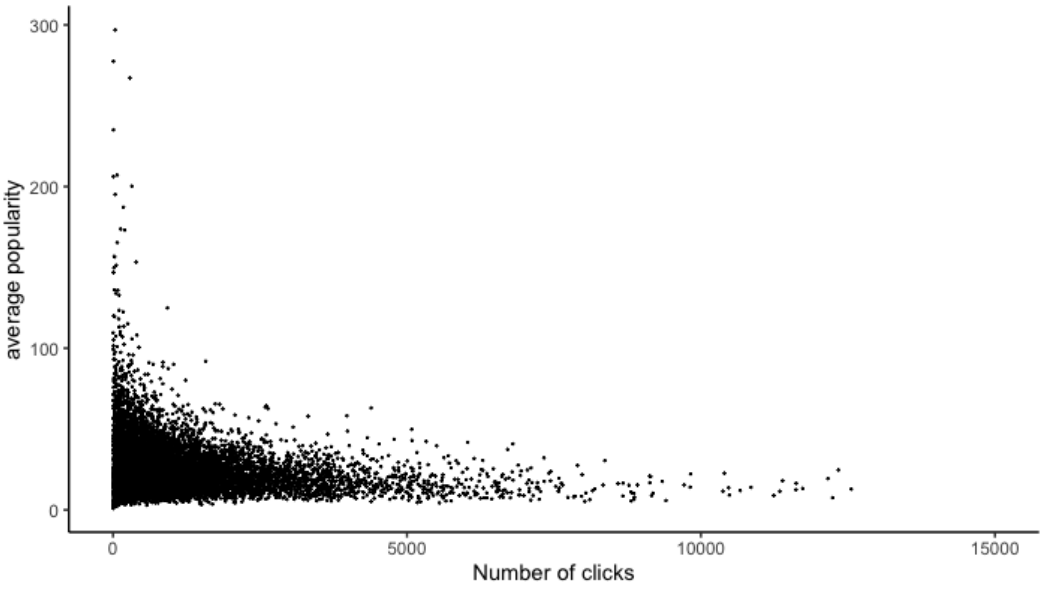


Figure 1: Number of clicks versus average popularity

find design s_2 (weakly) optimal, and so on.[11]Note that if no consumer finds design i optimal, the interval $[a_{i-1}, a_i)$ is empty. Clearly, a monotone interval partition is characterized by a sequence of nondecreasing cutoffs: $\mathbf{a} \equiv \{a_0, a_1, a_2, ..., a_n : a_0 = \underline{c} \text{ and } a_n = \bar{c}\}$.[12]

The number of intervals in the monotone partition may be less than the number of announced designs, as searching for some products can be strictly dominated. If every announced design attracts a positive mass of consumers, then we refer to such a collection of designs as an equilibrium design profile, denoted by $\mathcal{S}^*$, and use $\hat{\mathbf{c}} \equiv \{c_1, c_2, \ldots, c_n\}$ to represent cutoffs associated with this equilibrium partition. We omit $\hat{\mathbf{c}}$'s dependence on $\mathcal{S}^*$ and $P^e_{\mathcal{S}^*}$ whenever it is clear. The following proposition shows that in the equilibrium, consumers' search strategies are captured by a n-interval partition, with the cutoff type being indifferent between adjacent designs.

Proposition 3.4. *For any $\mathcal{S}^*$, the partition $\hat{\mathbf{c}}$ is chosen in a way that boundary types are indifferent between two neighboring designs.*

$$\hat{c}_i = \hat{c}(s_i, s_{i+1}), \qquad \forall i \in \{1, 2, \ldots, n-1\}, \quad \textit{and } s_i \in \mathcal{S}^*. \tag{5}$$

That is,

$$v^*_{s^i}(\hat{c}(s_i, s_{i+1})) - v^*_{s^{i+1}}(\hat{c}(s_i, s_{i+1})) = P^e_{s_i} - P^e_{s_{i+1}}. \tag{6}$$

Proof. We show that consumers in the ith interval chooses search mode i if and only if $\hat{c}_i = \hat{c}(s_i, s_{i+1})$. For the only if part, we prove by contrapositive. Suppose $\hat{c}(s_i, s_{i+1}) < \hat{c}_i$,

[11]A non-degenerated interval of search costs that allows indifference between multiple designs can lead to non-monotone partition.

[12] Define $a_i = \max(a_{i-1}, \min\{\hat{c}_{i,j}, j > i\})$ for $i = 1, \ldots, n-1$, where $\hat{c}_{i,j}$ is the cutoff between design s_i and s_j. When $a_i = a_{i-1}$, either no consumer find design i (weakly) optimal or precisely consumers with $c = a_i$ are indifferent between multiple designs, e.g. $i-1$,i and $i+1$. We assume none of them search design i in the latter case, so the set of consumers with a specific search mode is either empty or has a positive measure.

consumers with c slightly below $\hat{c}_i$ prefer design s_{i+1} over s_i, contradicting to the fact that consumers in the i th interval favor design s_i over the others. Similarly, if $\hat{c}(s_i, s_{i+1}) > \hat{c}_i$, consumers with search costs slightly above $\hat{c}_i$ lies in the $i+1$ interval but prefer design s_i over s_{i+1}, contradiction.

For the if part, we first observe consumers in the kth interval lies on the right-hand side of $\hat{c}(s_i, s_{i+1})$ for $i \in \{1, 2, \ldots, k-1\}$. Applying the transitivity of consumers' preference, consumers in the kth interval prefer design s_k to all designs more niche than it. Similarly, since consumers in the kth interval lies on the left hand side of $\hat{c}(s_i, s_{i+1})$ for $i \in \{k, \ldots, n\}$, they prefer design s^k to all broader designs. □

For each $\mathcal{S}^*$ and $P^e_{\mathcal{S}^*}$, there could be multiple $\hat{\mathbf{c}}$ satisfying equation (5) and (6). As one will see in the later discussion, we will exploit this multiplicity to establish the existence of equilibria. The partition structure not only yields an intuitive property on consumers' search behavior but significantly reduces the difficulty in deriving the demand and prices for different designs, which is the objective of the next section.

3.2 Firm's strategy

In this subsection, we focus on a specific set of equilibrium designs $\mathcal{S}^*$ and abbreviate design s_k as design k. Define $\mathbf{q} \equiv (q_1, q_2, \ldots, q_n)$, where q_k is the measure of type k firms in the market. We can decompose firms' decision-making into two stages. In the first stage, they commit to a specific design. In the second stage, firms select the optimal prices according to consumers' search behavior solicited by designs chosen in the first stage. Assuming firm j chooses design k in the first stage, we now focus on the second-stage problem. By symmetry, $g(c)dc/q_k$ unit of type k shoppers with search costs c come to firm j on their first visit and come to firm j on their second visit whenever they could not find satisfactory products at the first firm, which happens with probability $F_k(v^*_k(c))$. Since the type-k shoppers purchasing at firm j can visit no, one, two, ... up to an infinite number of firms before, the measure of type-k shoppers with search costs c visiting firm j is $\frac{g(c)dc}{q_k\left(1-F_k(v^*_k(c))\right)}$. Conditional on the visit, the type k shopper stops and purchases at the firm j whenever her actual surplus is higher than the expected surplus ($v_{k,j} - P_{k,j} \geq v^*_k(c) - P^e_k$). The demand from type-$k$ shoppers with search costs c is thus

$$D_k(P_{k,j}|\mathbf{P^e}, \mathbf{F}^*, \mathbf{q}, c)dc \equiv \frac{g(c)\left(1 - F_k(v^*_k(c) - P^e_k + P_{k,j})\right)}{q_k\left(1 - F_k(v^*_k(c))\right)}dc. \tag{7}$$

[13] Notice that if firm j charges a price $P_{k,j} \geq \bar{v}_k - v^*_k(c) + P^e_k$, no consumer with search cost lower than $c'(P_{k,j})$ will stop at firm j, where $c'(P_{k,j})$ solves $\bar{v}_k - v^*_k(c') + P^e_k - P_{k,j} = 0$. Specifically, any price $P_{k,j} > \bar{P}_k(P^e_k) \equiv \bar{v}_k - v^*_k(\bar{c}) + P^e_k$ prohibits all types of consumers from visiting and is strictly dominated. It is, therefore, without loss to confine the choice set to be $X \coloneqq \{P_{1,j}, P_{2,j}, \ldots, P_{n,j}|P_{k,j} \in [0, \min(\bar{v}_k, P_k(P^e_k))]\}$. Since corner solutions yield zero profits, given the belief in equilibrium prices, the optimal pricing strategy for firm j must lie in the interior of this choice set.

From proposition 3.4, the impact of heterogeneous search costs on consumers' search modes can be summarized by equation (5). Since the target customers of type-k firms are those with search costs ranging from $\hat{c}_{k-1}$ to $\hat{c}_k$, the profit function of type-k firm is

[13] The demand from type k shoppers with $c > \tilde{c_k}$ is $\frac{g(c)\left(1-F_k(v^*_k(c)-P^e_k+P_{k,j})\right)}{q_k}$. This is because all type k shoppers stop and purchase on their first visit in the equilibrium($v_{k,j} \geq v^*_k(c)$ for all $c > \tilde{c_k}$). Therefore, when firm j encounters type k shoppers with $c > \tilde{c_k}$, it knows with certainty that it must be their first visit.

$$\pi_k(P_{k,j}|\mathbf{P^e},\mathbf{F}^*,\mathbf{q}) \equiv P_{k,j}\int_{\hat{c}_{k-1}}^{\hat{c}_k} D_k(P_{k,j}|\mathbf{P^e},\mathbf{F}^*,\mathbf{q},c)\,dc. \tag{8}$$

If the profit maximization admits an interior solution, $P_k^* = P_k^e$ together with FOC implies

$$\int_{\hat{c}_{k-1}}^{\hat{c}_k}\left[1 - P_k^*\frac{f_k(v_k^*(c))}{1-F_k(v_k^*(c))}\right]dG(c) = 0. \tag{9}$$

For all $k \in \mathcal{S}^*$, denote the hazard rate corresponding to distribution F_k by $H_k(v) \equiv \frac{f_k(v)}{1-F_k(v)}$, and the conditional distribution of search costs on the kth interval by $\tilde{G}(c;k) \equiv \frac{G(c)-G(\hat{c}_{k-1})}{G(\hat{c}_k)-G(\hat{c}_{k-1})}$ for all $c \in [\hat{c}_{k-1},\hat{c}_k]$, with $\tilde{g}(c;k) = \tilde{G}'(c;k)$. A necessary condition for $\mathcal{S}^*$ to be sustained as an equilibrium design profile is

$$P_k^*(\mathcal{S}^*) = \frac{1}{\int_{\hat{c}_{k-1}}^{\hat{c}_k} H_k(v_k^*(c))\,d\tilde{G}(c;k)}. \tag{10}$$

The equality indicates that the equilibrium price negatively correlates with the conditional expectation of the hazard rate. To explain the intuition, let's restrict our attention to type-k shoppers with search cost c. Because all such shoppers stopping at the firm j obtain expected utility $v_k^*(c) - P_k^*$, a marginal price increase prevents those with valuations equal to the reservation utility from purchasing. Therefore, the marginal consumers accounts for $\frac{f_k(v_k^*(c))}{1-F_k(v_k^*(c))}$ portion of the population eventually buy from firm j. Since the target of type-k firms consists of consumers with search costs ranging from $\hat{c}_{k-1}$ to $\hat{c}_k$, $P_k^* = \frac{1}{\mathcal{H}_k(\hat{c}_{k-1},\hat{c}_k|\mathcal{S}^*)}$ is simply the inverse of the conditional expectation of the hazard rate over the range, where $\mathcal{H}_k(\hat{c}_{k-1},\hat{c}_k|\mathcal{S}^*) \equiv \int_{\hat{c}_{k-1}}^{\hat{c}_k} H_k(v_k^*(c))\,d\tilde{G}(c;k)$. Differentiating the conditional hazard rate with respect to the upper and lower bound, we get

$$\frac{\partial \mathcal{H}_k}{\partial \hat{c}_k} = \tilde{g}(\hat{c}_k;k)\left[H_k(v_k^*(\hat{c}_k)) - \mathcal{H}_k\right] < 0,$$

and

$$\frac{\partial \mathcal{H}_k}{\partial \hat{c}_{k-1}} = \tilde{g}(\hat{c}_{k-1};k)\left[\mathcal{H}_k - H_k(v_k^*(\hat{c}_{k-1}))\right] < 0.$$

An upward shift of the lower or the upper bound induces consumers to search less intensively due to higher average search costs; less choosy consumers correspond to a strictly less elastic demand and a higher price.

When will FOC be sufficient for the optimality of price? In the following, we provide two conditions that guarantee the quasi-concavity of the profit function and assume they are satisfied thereafter. One can easily verify that the second-order derivative of the profit function is non-positive if condition (i) holds. Condition (ii) follows directly from Theorem 1 in Moraga-González et al.(2017). The basic idea is to observe that the demand from a particular consumer is log-concave in the price and the reservation utility under condition (ii). After that, they show that log-concavity is preserved under aggregation by the Prékopa Theorem(1973). Since the profit function is a product of the price and the aggregate demand, the profit function is log-concave and, therefore, quasi-concave.

Assumption 3. For all k either (i) $f_k'(v_k) \geq 0$ or (ii) $g(v_k^{*-1}(t))$ is log-concave.

3.3 equilibrium

We first point out any equilibrium profile can be fully characterized by $\mathcal{S}^*$, $\hat{\mathbf{c}}$, $\mathbf{q}^*$, and $\mathbf{P^e}$. For each equilibrium design profile $\mathcal{S}^*$, consumers search according to the partition structure given in (5), and firms set on-path prices according to equation (10), both of which are solely determined by $\hat{\mathbf{c}}$. The product design profile $\mathbf{F}^*$, by ex-ante symmetry of firms, can be equivalently translated into the composition of firms $\mathbf{q}^*$, which is endogenously determined by the profit indifference condition (11). For each $k, k' \in \mathcal{S}^*$,

$$\pi_k(P_k^*|\mathbf{P^e}, \mathbf{F}^*, \mathbf{q}^*) = \pi_{k'}(P_{k'}^*|\mathbf{P^e}, \mathbf{F}^*, \mathbf{q}^*) = \pi^*(\mathcal{S}^*). \tag{11}$$

Finally, beliefs in the prices of on-path designs should be consistent with the actual prices, and beliefs in the prices of off-path designs should prevent any consumer from exploring it in case appearing at the first stage. As consumers have all the freedom to assign beliefs out of the equilibrium path, setting expected prices of unseen designs to be exceedingly high ensures no consumer is willing to look into them; such an equilibrium exists as long as $\hat{\mathbf{c}}$ and $\mathbf{P}^*$ satisfying equations (6) and (10).[14] An important inquiry concerns the rationality of consumer beliefs in off-path designs, specifically whether the firm will indeed adhere to the anticipated price if it adopts an unseen design. Given consumer belief, every price yields zero payoffs, so the firm should be indifferent between them. But is some prices more sensible than others? Firms should certainly exclude some prices from their consideration, such as any price above $\frac{1}{H_k(v_k^*(\bar{c}))}$ or below $\frac{1}{H_k(v_k^*(\underline{c}))}$ because any price within this range cannot be supported by any customer base $\mathcal{A} \subseteq [\underline{c}, \bar{c}]$. Moreover, if a design can attract any consumer, the set of customers must be convex. Given these observations, we can define rational belief formally as follows.

Definition 3.1. A belief system $\mathbf{P^e}$ is rational if for all k, there exists a convex set $\mathcal{A}_k \subseteq [\underline{c}, \bar{c}]$ such that P_k^e solves

$$\max_{[0,\bar{v_k}]} p_k \int_{\mathcal{A}_k} \frac{1 - F_k(v_k^*(c) - P_k^e + p_k)}{1 - F_k(v_k^*(c))}\, dG(c).$$

Next, we show that an equilibrium exists under a rational belief system and briefly describe the construction, with detail relegated to the appendices.

Proposition 3.5. *There exists an equilibrium* $(\mathcal{S}^*, \hat{\mathbf{c}}, \mathbf{q}^*, \mathbf{P^e})$ *supported by a rational belief.*

The idea is to construct a self-map from the product space of a weakly increasing sequence of cutoffs $\mathbf{a}$ and expected prices $\mathbf{P^e}$ to itself. For any belief in prices, the optimal search decision is captured by a correspondence of weakly increasing cutoffs summarizing the associated monotone interval partition in a fashion similar to footnote 12. Firms choose optimal prices according to a sequence of cutoffs and expected prices, assuming consumers between the $k-1$th and the kth cutoff search for design k. If this interval degenerates, no consumer prefers to search design k, and the firm forms its profit maximization problem assuming only consumers with search costs equal to the cutoff pay a visit. Then we show these two correspondences are upper-hemicontinuous, compact-and convex-valued and apply the Kakutani fixed-point theorem to prove the existence. Our next objective is to

[14] In particular, any market with a single design k can be sustained as an equilibrium by setting the expected price of an off-path design j to satisfy $P_j^e > v_j^*(\underline{c}) - v_k^*(\underline{c}) - P_k^*(\{s_k\})$ for all $j < i$ and $P_j^e > v_j^*(\bar{c}) - v_k^*(\bar{c}) - P_k^*(\{s_k\})$ for all $j < i$. The question is whether this belief is rational.

examine the condition in which multiple designs coexist in the market, known as a separating equilibrium, as well as the situation where a single design dominates, referred to as a pooling equilibrium. To facilitate analysis, we limit our focus to equilibria with a maximum of two designs and outline some key properties.

3.4 Separating equilibrium ($|\mathcal{S}^*| = 2$)

Lemma 3.6. *There exists a separating equilibrium only if $\hat{c}_1 < \tilde{c}_2$.*

Proof. Suppose towards a contradiction; there exists a separating equilibrium in which all broad shoppers have search costs $c \geq \hat{c}_1 \geq \tilde{c}_2$. In the equilibrium, almost all broad shoppers except the one with $c = \tilde{c}_2$ will stop and purchase at their first visit, which makes the demand infinitely inelastic locally. Increasing the price of the broad products by ε is equivalent to shifting consumers' match value downward by ε. The demand from consumers with search costs higher than $c^*(\varepsilon)$ remains intact since these consumers will purchase immediately like before, where $c^*(\varepsilon)$ solves $v_2^*(c^*) + \varepsilon = \underline{v}_2$. On the other hand, the probability of purchasing for consumers with $\tilde{c}_2 \leq c < c^*(\varepsilon)$ will decline by $F_2(v_2^*(c) + \varepsilon)$ since now broad shoppers continue to search if $v_2 - \varepsilon < v_2^*(c)$. Define $\varepsilon_1 \equiv c^*(\varepsilon) - \tilde{c}_2$. The marginal impact of raising P_2^* by ε is, therefore, at least $[1 - G(c^*(\varepsilon))]\varepsilon - P_2^* \int_{\tilde{c}_2}^{\tilde{c}_2+\varepsilon_1} F_2(v_2^*(c) + \varepsilon)\, dG(c)$. The second term is negligible compared to the first term because $F_2(v_2^*(c) + \varepsilon) < F_2(\underline{v}_2 + \varepsilon) < \varepsilon_2$ for $\varepsilon_2 > 0$, implying raising price by ε is profitable. □

The equilibrium prices are determined by three factors: the dispersion effect, the location effect, and the compositional effect. For any fixed v, and $s' < s$, $H_{s'}(v) \leq H_s(v)$. This result comes from theorem 2.1 of Khaledi (2002), stating that under increasing failure rate(which is ensured by assumption 2), dispersed order implies hazard rate order. Since the hazard rate corresponds to the portion of marginal consumers susceptible to an increase in price, a higher hazard rate is associated with more elastic marginal demand and a lower price. The first force, named the dispersion effect, makes standardized products cheaper. Second, the position of the marginal consumers in the distribution varies for different search modes. Suppose all consumers have search costs c; niche shoppers have higher reservation utilities due to greater mixing and matching benefits. Given the same product design, the demand from niche shoppers is closer to the upper tail of the distribution and hence is more elastic. We term this the location effect. The compositional effect captures the difference in search costs between niche and broad shoppers, that is, the domain of search costs which the hazard rate function integrates over. As broad shoppers are composed of consumers with high search costs, firms have a higher local monopoly power over standardized products. In other words, the search-discouraging impact is stronger for broad shoppers assuming the same valuation distribution. Mathematically, $H_k(v_k^*(c))$ is a non-increasing function of c, suggesting the demand from high-cost shoppers is less elastic. In summary, the dispersion effect works opposite to the compositional and the location effect in the direction of raising (reducing) the elasticity of the broad (niche) products.

The following lemma states that the dispersion effect should dominate the combination of the location effect and composition effect, leading to higher prices for niche products. Essentially, a firm can charge a higher price for the niche product because it caters to consumers' tastes better than the generic product, captured by $v_1^*(c) \geq v_2^*(c)$ for all c.

Lemma 3.7. *A niche product is no cheaper than a standardized product, and the number of consumers per firm is higher for standardized products.*

Proof. From the consumer's indifference condition and corollary 3.3, we know the price of a niche product must be no lower than a standardized product. The second part of the statement follows directly from the profit-indifference condition. □

Corollary 3.8. *For any $\mathcal{S}^*$, $H_1(v_1^*(\hat{c}_1)) < H_2(v_2^*(\hat{c}_1))$.*

Proof. From the previous lemma, we understand the dispersion effect works against and dominates the other two effects. Observing the dispersion effect outweighs the location effect in the absence of the composition effect completes the proof. Mathematically,

$$\frac{1}{H_1(v_1^*(\hat{c}_1))} > P_1^* = \frac{1}{\mathcal{H}_1(\underline{c}, \hat{c}_1|\mathcal{S}^*)} > P_2^* = \frac{1}{\mathcal{H}_2(\hat{c}_1, \bar{c}|\mathcal{S}^*)} > \frac{1}{H_2(v_2^*(\hat{c}_1))}.$$

□

The corollary positions that the dispersion effect dominates the location effect for cutoff-type consumers. In general, it's hard to check whether the condition is satisfied due to difficulty in calculating the reservation utility. However, it's easy to verify the condition holds for all c if the valuation density is uniform.[15] Now we are ready to provide a sufficient condition for the existence of $\hat{c}_1$. Define $h(\hat{c}_1|S^*) \equiv v_1^*(\hat{c}_1) - v_2^*(\hat{c}_1) - P_1^*(\hat{c}_1) + P_2^*(\hat{c}_1)$ as the change in consumer surplus when the cutoff type switch from searching the broad design to the niche design. Since the difference between the reservation utility is non-increasing in search costs, $h(\hat{c}_1|\mathcal{S}^*)$ decreases with $\hat{c}_1$ if $P_1^* - P_2^*$ is increasing. The slope of the price difference is given by

$$\frac{d}{d\hat{c}_1}(P_1^* - P_2^*) = -\frac{g(\hat{c}_1)[H_1(v_1^*(\hat{c}_1)) - \mathcal{H}_1]}{\mathcal{H}_1^2 G(\hat{c}_1)} + \frac{g(\hat{c}_1)[\mathcal{H}_2 - H_2(v_2^*(\hat{c}_1))]}{\mathcal{H}_2^2[1 - G(\hat{c}_1)]}$$

Therefore, $h(\hat{c}_1|\mathcal{S}^*)$ is decreasing if

$$|\frac{dP_1^*}{d\hat{c}_1} / \frac{dP_2^*}{d\hat{c}_1}| = \frac{\mathcal{H}_2^2[H_1(v_1^*(\hat{c}_1)) - \mathcal{H}_1]}{\mathcal{H}_1^2[\mathcal{H}_2 - H_2(v_2^*(\hat{c}_1))]} \times \frac{1 - G(\hat{c}_1)}{G(\hat{c}_1)} > 1.$$

Since the first term is positive, the condition holds if there is not much density accumulating in the low-cost region or $g(c)$ is an increasing and sufficiently convex function. In what follows, we show a unique separating equilibrium exists if $G(\tilde{c}_2)$ is small enough. Intuitively speaking, an upward shift of the cutoff would induce median-cost shoppers to switch from searching for niche products to generic products. When the density of low-cost consumers is insignificant compared to that of high-cost consumers, this change in extensive margin introduces a greater proportion of median-cost consumers to enter the market for niches than to exit the market for generics. That is to say, the search-discouraging effect intensifies more for the niche product, causing $P_1^* - P_2^*$ to decrease arbitrarily fast with the cutoff. The condition $\underline{c} = 0$ ensures that the price of the niche product is sufficiently low, so consumers with the lowest search costs prefer to be niche shoppers. The proof holds through as long as the search costs of the lowest type satisfy $\lim_{\hat{c}_1 \to \underline{c}} h(\hat{c}_1|\mathcal{S}^*) > 0$.

Proposition 3.9. *Suppose $\underline{c} = 0$, there exists a k^*, such that there is a unique separating equilibrium if $G(\tilde{c}_2) < k^*$.*

[15] Consider uniform valuation distribution $F(v;\gamma)$ with mean $\frac{1}{2}$ over the support $[\frac{1-\gamma}{2}, \frac{1+\gamma}{2}]$, where γ measures the level of dispersion. Let X_γ be the random variable associated with distribution $F(v;\gamma)$. One can easily verify $X_\gamma \leq_{\text{disp}} X_{\gamma'}$ if $\gamma' > \gamma$. Define the hazard rate of X_γ by H_γ and the reservation utility by v_γ^*. $H_\gamma(v_\gamma^*(c)) = \frac{1}{\sqrt{2\gamma c}}$, which is a decreasing function of γ.

Proof. We first focus on the case where $\tilde{c_2} < \bar{c}$. Since $G(c)$ has no atom, $\mathcal{H}_i^2$, and $H_i(v_i^*(c)) - \mathcal{H}_i$ are both bounded above and below. There exists $\varepsilon > 0$, such that

$$\frac{\mathcal{H}_2^2[H_1(v_1^*(\hat{c}_1)) - \mathcal{H}_1]}{\mathcal{H}_1^2[\mathcal{H}_2 - H_2(v_2^*(\hat{c}_1))]} > \varepsilon, \text{ and } -\frac{g(\hat{c}_1)[H_1(v_1^*(\hat{c}_1)) - \mathcal{H}_1]}{\mathcal{H}_1^2} > \varepsilon.$$

Setting $G(\tilde{c_2}) < k^* = \min(\frac{\varepsilon}{2+\varepsilon}, \frac{\varepsilon \tilde{c}_2}{2[\bar{v}_1 - \bar{v}_2 + P_2^*(0)]})$ ensures $|\frac{dP_1^*}{d\hat{c}_1} / \frac{dP_2^*}{d\hat{c}_1}| > 2$ and therefore implies for all $\hat{c}_1 < \tilde{c}_2$,

$$\frac{d}{d\hat{c}_1}(P_1^* - P_2^*) > \frac{1}{2}\frac{d}{d\hat{c}_1}P_1^* = -\frac{g(\hat{c}_1)[H_1(v_1^*(\hat{c}_1)) - \mathcal{H}_1]}{2\mathcal{H}_1^2 G(\hat{c}_1)} > \frac{\varepsilon}{2G(\hat{c}_1)} > \frac{\bar{v}_1 - \bar{v}_2 + P_2^*(0)}{\tilde{c}_2}.$$

If $G(\hat{c}_2) < k^*$, we have

$$\lim_{\hat{c}_1 \to \tilde{c}_2} h(\hat{c}_1|\mathcal{S}^*) = h(0|\mathcal{S}^*) + \int_0^{\tilde{c}_2} \frac{dh(\hat{c}_1|\mathcal{S}^*)}{d\hat{c}_1} d\hat{c}_1 \leq \bar{v}_1 - \bar{v}_2 + P_2^*(0) - \int_0^{\tilde{c}_2} \frac{dP_1^* - P_2^*}{d\hat{c}_1} d\hat{c}_1 < 0,$$

and

$$\lim_{\hat{c}_1 \to 0} h(\hat{c}_1|\mathcal{S}^*) = \bar{v}_1 - \bar{v}_2 + P_2^*(0) > 0.$$

Here $\bar{v}_1 - \bar{v}_2 > 0$ follows from the property that the support of v_2 is a strict subset of the support of v_1. By the intermediate value theorem, there exists a $\hat{c}_1^*$ solving $h(\hat{c}_1|\mathcal{S}^*) = 0$. The uniqueness follows from the fact that $h(\hat{c}_1|\mathcal{S}^*)$ is decreasing on $[0, \tilde{c}_2]$. Proof for the case of $\bar{c} < \tilde{c}_2$ is similar by observing the slope of $h(\hat{c}_1|\mathcal{S}^*)$ can be made arbitrarily negative. □

One may wonder whether the assumption that the proportion of low-cost consumers is small meets real-life situations. As indicated by graph 1, the majority of users accumulate at the left part; and as the number of clicks increase(corresponding to lower search costs), the density becomes more and more sparse. The data set shows that most of the users click no more than 5000 times, and only 256 out of 10000 users sampled more than that, suggesting condition stated in the hypothesis of proposition 3.9 is likely to be met in this particular example.

The next result shows that pooling equilibrium with the most niche design generates the highest profit.

Lemma 3.10. *Let $\mathcal{S}^*$ be any design profile of a separating equilibrium, $\pi^*(\{s_1\}) > \pi^*(\mathcal{S}^*)$, where s_1 is the most niche design in $\mathcal{S}^*$.*

Proof.

$$\pi^*(\{s_1\}) = P_1(\{s_1\}) > P_1(\mathcal{S}^*) \geq \sum_{i=1}^{n} P_i^*(\mathcal{S}^*)[G(\hat{c}_i) - G(\hat{c}_{i-1})] = \pi^*(\mathcal{S}^*).$$

The first inequality follows from the fact that only consumers with the lowest search costs choose the most niche design in a separating equilibrium, in contrast with the pooling equilibrium, where firms target the whole spectrum of buyers. □

4 Comparative statics ($|\mathcal{S}| = 2$)

In this section, we parameterize the distribution of search costs by a positive number β. Specifically, denoting the parametrized CDF by $G(c;\beta)$ and denote the corresponding pdf by $g(c,\beta)$ defined over the support $[\underline{c},\bar{c}]$. Here we assume the support does not depend on β so as to isolate the effect of an endogenous shift of the cutoff. Furthermore, we denote the conditional distribution of c truncated above at $\hat{c}_1$ by $\bar{G}(c;\beta,\hat{c}_1)$, the counterpart truncated below at $\hat{c}_1$ by $\underline{G}(c;\beta,\hat{c}_1)$, with density $\bar{g}(c;\beta,\hat{c}_1)$ and $\underline{g}(c;\beta,\hat{c}_1)$ respectively. $G(c;\beta)$ is continuously differentiable with β. With this definition, the conditional hazard rates of niche shoppers and broad shoppers are given by $\mathcal{H}_1(\underline{c},\hat{c}_1;\beta) = \int_{\underline{c}}^{\hat{c}_1} H_1(v_1^*(c))\,d\bar{G}(c,\hat{c}_1;\beta)$ and $\mathcal{H}_2(\hat{c}_1,\bar{c};\beta) = \int_{\hat{c}_1}^{\bar{c}} H_2(v_2^*(c))\,d\underline{G}(c,\hat{c}_1;\beta)$ respectively. The objective is to study how changes in the distribution of search costs affect the cutoff and equilibrium prices.

Definition 4.1. Search costs decrease with β in the FOSD sense if $G(c;\beta) \le G(c;\beta')$ for all c and $\beta' \le \beta$. i.e. $G_\beta(c;\beta) \le 0$ for all c with strict inequality for some c.

An equivalent definition of FOSD is that for every non-decreasing function $u : \mathbb{R} \to \mathbb{R}$, $\int u(x)\,dG(c;\beta) \ge \int u(x)\,dG(c;\beta')$. This alternative definition helps proving the next result.

Lemma 4.1. *In the pooling equilibrium, a FOSD decrease in search costs reduces the equilibrium price. i.e. $P_s(\beta') \le P_s(\beta)$ for all $\beta' < \beta$.*

Proof. Since $H_s(v_s^*(c))$ is non-increasing in c, we have for all $\beta > \beta'$,

$$\int_{\underline{c}}^{\bar{c}} H_s(v_s^*(c))\,dG(c;\beta) \le \int_{\underline{c}}^{\bar{c}} H_s(v_s^*(c))\,dG(c;\beta').$$

Taking the reciprocal of the inequality completes the proof. □

In a separating equilibrium, the cutoff $\hat{c}_1^*(\beta)$ changes with the distribution of search costs, and solves

$$L(\hat{c}_1,\beta) \equiv v_1^*(\hat{c}_1) - v_2^*(\hat{c}_1) - \frac{1}{\mathcal{H}_1(\underline{c},\hat{c}_1;\beta)} + \frac{1}{\mathcal{H}_2(\hat{c}_1,\bar{c};\beta)} = 0.$$

Provided $\frac{\partial f}{\partial \hat{c}_1} \neq 0$ at the equilibrium cutoff $\hat{c}_1^*(\beta)$, the implicit function theorem suggests the effect of a marginal decrease in β on the cutoff is

$$\frac{d\hat{c}_1^*(\beta)}{d\beta} = -\frac{\frac{\partial L(\hat{c}_1^*(\beta),\beta)}{\partial \beta}}{\frac{\partial L(\hat{c}_1^*(\beta),\beta)}{\partial \hat{c}_1}}.$$

By proposition 3.9, the denominator is negative given there are sufficiently many high-cost consumers. The sign of $\frac{d\hat{c}_1^*(\beta)}{d\beta}$ therefore depends on the sign of $\frac{\partial L(\hat{c}_1),\beta)}{\partial \beta}$ at the equilibrium cutoff, which is

$$\frac{1}{(\mathcal{H}_1)^2}\frac{\partial}{\partial \beta}\mathcal{H}_1(\underline{c},\hat{c}_1;\beta) - \frac{1}{(\mathcal{H}_2)^2}\frac{\partial}{\partial \beta}\mathcal{H}_2(\hat{c}_1,\bar{c};\beta)\Big|_{\hat{c}_1=\hat{c}_1^*}. \tag{12}$$

Once the sign of $\frac{d\hat{c}_1^*(\beta)}{d\beta}$ is obtained, we can use it to derive the comparative statics of equilibrium prices. As the equilibrium price is inversely proportional to the conditional hazard rate, we focus on how search costs affect the latter. A marginal decrease in search

costs affects the equilibrium prices in two ways, which can be seen clearly from the derivative of the conditional hazard rate given by

$$\frac{d}{d\beta}\mathcal{H}_1 = \bar{g}(\hat{c}_1^*;\hat{c}_1^*,\beta)\left[H_1(v_1^*(\hat{c}_1^*(\beta))) - \mathcal{H}_1\right]\frac{d\hat{c}_1^*(\beta)}{d\beta} + \frac{\partial \mathcal{H}_1}{\partial \beta}. \tag{13}$$

The first term is the compositional effect (extensive margin), an endogenous shift of the cutoff induces consumers to join or drop out of the market of niches. In the case that $\frac{d\hat{c}_1^*(\beta)}{d\beta} < 0$, a decrease in the search cost would incentivize marginal consumers originally purchasing standardized products to switch to niches. As these newcomers have higher search costs, the compositional effect tends to reduce the average elasticity and raise the price of niche products. The second term captures the search-discouraging effect (intensive margin), which is equivalent to change in the conditional distribution of search costs in the absence of a cutoff change, indicating the average search intensity had the cutoff remain unchanged. Similarly, the price decomposition for standardized products is

$$\frac{d}{d\beta}\mathcal{H}_2 = \underline{g}(\hat{c}_1^*;\hat{c}_1^*,\beta)\left[\mathcal{H}_2 - H_2(v_2^*(\hat{c}_1^*(\beta)))\right]\frac{d\hat{c}_1}{d\beta} + \frac{\partial \mathcal{H}_2}{\partial \beta}. \tag{14}$$

In general, it is hard to determine the sign of the search-discouraging effect, equivalently, whether a FOSD decrease of the distribution of search costs would result in a FOSD decrease of the distribution truncated from above (below). Proposition 3 and 4 of Moraga-González et al.(2017) show that the statement is true if the density has the increasing likelihood ratio property (ILRP) and the reserve result— a FOSD decrease leads to a FOSD increase of the distribution truncated above— is valid under decreasing likelihood ratio property(DLRP). Nonetheless, how cutoff changes is determined by the relative magnitude of the search-discouraging effect between two markets. Therefore, knowing the direction of the change of the stochastic ordering is insufficient to pin down the comparative statics of either the cutoff or prices. As a simplification, below we focus on two special cases where the search discouraging effect is absent for one market ($\frac{\partial \mathcal{H}_k}{\partial \beta} = 0$ for either $k = 1$ or $k = 2$). These special settings allow us to derive unambiguous comparative statics.

Proposition 4.2. *Let $\hat{c}_1^*(\beta)$ be the unique separating equilibrium cutoff corresponding to density $g(c;\beta)$. Provided a unique equilibrium exists for a small perturbation around β, we have the following.*

(i) If a FOSD decrease (increase) in search costs does not affect distribution above $\hat{c}_1^*(\beta)$, the change results in a higher (lower) cutoff and higher (lower) prices for standardized products. The price response of niches is ambiguous.

(ii) If a FOSD decrease (increase) in search costs does not affect distribution below $\hat{c}_1^*(\beta)$, the change results in a lower (higher) cutoff, lower (higher) prices for both niches and standardized products.

Proof. Under condition (i), we have $\frac{\partial \mathcal{H}_1(\underline{c},\hat{c}_1^*;\beta)}{\partial \beta} \leq 0$ and $\frac{\partial \mathcal{H}_2(\hat{c}_1^*,\bar{c};\beta)}{\partial \beta} = 0$ by the definition of FOSD of the distribution truncated above. Equation (12) reveals the cutoff rises with a marginal decrease in $\beta(\frac{d\hat{c}_1^*(\beta)}{d\beta} \leq 0)$. Since $\mathcal{H}_2(v_2^*(c))$ is non-increasing, $\mathcal{H}_2 - \mathcal{H}_2(v_2^*(\hat{c}_1^*(\beta))) \leq 0$. Equation (14) then implies that in the absence of the search discouraging effect, the compositional effect lowers the conditional hazard rate and drives up the prices of standardized products. The impact on niche products is ambiguous since the compositional effect works

against the search-encouraging effect. If the mass of consumers switching their search behavior is substantial, the newly join high-cost niche shoppers could more than counteract the reduction of search costs in the low-cost region, leading to higher prices of niche products. Under condition (ii), we have $\frac{\partial \mathcal{H}_1(\underline{c},\hat{c}_1^*;\beta)}{\partial \beta} = 0$, $\frac{\partial \mathcal{H}_2(\hat{c}_1^*,\bar{c};\beta)}{\partial \beta} \leq 0$, and $\frac{d\hat{c}_1^*(\beta)}{d\beta} \geq 0$.When search costs decline in the upper part of the distribution, standardized products become cheaper, which induces more consumers to search for blockbusters, therefore shifting the cutoff downwards. Only the compositional effect is in play for niches; as those medium-cost consumers change to purchasing standardized products, the average search costs of niche shoppers decrease, leading to lower prices. For the standardized product, the result is unambiguous since the two effects work in the same direction. A decline in search costs for those with high search costs intensifies the price competition for standardized products and induces more median-cost consumers to join the generic market. As the demand from these consumers is relatively elastic compared to those originally in the market, the compositional effect reinforces the search discouraging effect, making the competition more fierce. □

The above proposition implies the group of people enjoying an improvement in search technology matters for the price response. Suppose heterogeneity in search costs is due to accessibility to the internet. For instance, the old generation may go to a supermarket physically if they want to do grocery shopping, whereas the young generation can submit their order through an e-commerce platform. If the improvement occurs mainly for consumers with low search costs, such as the introduction of a recommendation engine. In that case, the effect is likely anti-competitive(especially when the extensive margin plays a crucial role). In contrast, a more efficient retail display or better traffic control is pro-competitive. With this observation, we contend that companies may use more obfuscation offline than online. The implication is particularly important as the parallel operation on both markets gains increasing popularity. Two examples are Amazon's acquisition of Whole Foods Markets and Walmart's purchase of Jet.com.

5 Conclusion

In this paper, we study how search costs influence product designs through the demand channel and demonstrate the negative relationship between search costs and designs carried over to an individual level. Our first contribution is to provide a theoretical framework to model how search friction affects consumers' preferences toward different designs. The assumption of dispersive ordering provides a tractable partition structure that allows us to simplify the demand and equilibrium characterization. Empirically, a tendency toward niche products for consumers with low-search costs is verified from the e-commence data provided by Taobao. Our second contribution is to offer another justification for the unconventional price response in addition to the previous literature (Chen and Zhang 2011, Zhou 2014, MorMoraga-González et al. 2017, Choi et al. 2018). When search costs decrease, medium-type consumers could endogenously switch from seeking generics to niches, raising the average search intensity and prices in both markets. Specifically, we highlight that reduction in search costs can have either pro-competitive or anti-competitive effects depending on the distribution's shape, implying that obfuscation can harm firms in some cases.

Finally, our model can be applied to the study of bundling. Given that a bundle is more dispersed than the item, our first result indicates consumers with low-search costs incline towards purchasing individual items separately. This is in line with Harris et al.(2006), in

which they show experimentally that subjects with less motivation to process information have a stronger preference toward bundles.

6 Example

6.1 Uniform valuation with a ploynomial cost function

In the following, we assume $F_1(v_1) = v_1$ for $v_1 \in [0,1]$, and $F_2(v_2) = 2(v_2 - \frac{1}{2}) + \frac{1}{2}$ for $v_2 \in [\frac{1}{4}, \frac{3}{4}]$. Let $G(c) = (\frac{c}{\bar{c}})^\alpha$ for $c \in [0, \bar{c}]$, where $\bar{c} = \tilde{c_2} = \frac{1}{4}$. The reservation utilities of a niche and a standardized product are $v_1^*(c) = 1 - \sqrt{2c}$ and $v_2^*(c) = \frac{3}{4} - \sqrt{c}$, respectively. FOC is sufficient to pin down the optimal pricing strategy because $f_i'(v) = 0$ for $i = 1, 2$.

By proposition 3.9, there exists a unique separating equilibrium if α is sufficiently large. Specifically, there exists a equilibrium if there is a $\hat{c} \in [0, 1/4]$ such that $v_1^*(\hat{c}) - v_2^*(\hat{c}) - P_1^*(\hat{c}) + P_2^*(\hat{c}) = 0$. Plugging the equilibrium price in (10) into the expression gives

$$h(\hat{c}; \alpha) \equiv \frac{\sqrt{\hat{c}}(2\alpha - 1)((4\hat{c})^\alpha - 1)}{2\alpha\left((4\hat{c})^\alpha - 2\sqrt{\hat{c}}\right)}.$$

Evaluating the limit at $\hat{c} = 1/4$ gives

$$\frac{1}{2} - \frac{(2\alpha - 1)}{2\sqrt{2}a} + (3/4 - \sqrt{\frac{1}{2}}).$$

Using the same logic in the proof of proposition 3.9, a separating equilibrium exists for $\alpha > \frac{\sqrt{2}}{4\sqrt{2}-5} \approx 2.15$. The uniqueness of the equilibrium cannot be proved formally due to analytical intractability. However, we can numerically show $h_1(\hat{c}, \alpha) < 0$ for all $0 \leq \hat{c} \leq \frac{1}{4}$ and $2.5 < \alpha < 6$ by the MATLAB, suggesting for each α, there is a unique cutoff $\hat{c}$.

Figure 2 shows that a FOSD increase in the search cost lowers the prices of both niche and broad products for $\alpha \in [2.5, 3.15]$, whereas the change leads to an increase in the niche price and a decrease in broad price for $\alpha \in [3.15, 6]$. To disentangle the reason behind the unconventional price response in the former case, we define $\hat{c} = 0.2$ as the benchmark and compare the price changes with and without the compositional effect. If $\alpha = 3.45$, the actual equilibrium cutoff coincides with the benchmark cutoff, and the equilibrium prices equal the hypothetical prices without an endogenous shift of consumers' search behavior. Since $g(c; \alpha)$ satisfies the increasing likelihood ratio property, an increase in α implies a FOSD increase in the conditional distribution truncated above and truncated below by theorem 1.C.6 in Shaked and Shanthikumar (2007). Therefore, in the absence of change in the cutoff, the search-discouraging effect prevails, resulting in higher prices for both designs, as shown by the dotted line.

In contrast, when the cutoff is endogenous, an increase in α causes the price of a niche product to rise more than a standardized product, rendering generic designs more attractive. From figure 3, the proportion of niche shoppers decreases with α over the whole domain. When this countervailing force from the compositional change dominates the search discouraging effect, equilibrium prices could fall with search costs as opposed to conventional wisdom.

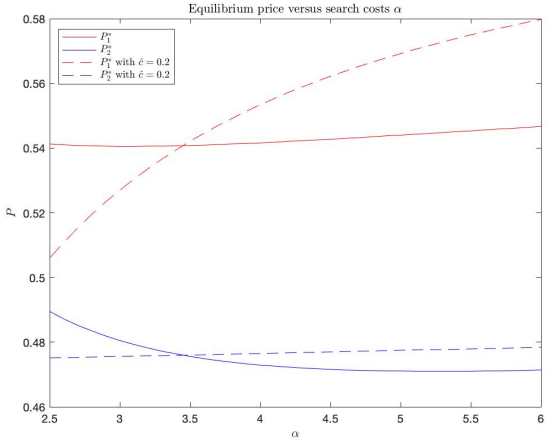


Figure 2: Equilibrium prices with and without the compositional effect

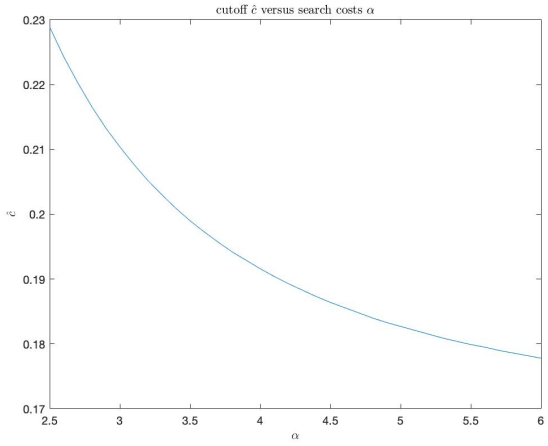


Figure 3: Cutoff versus search costs

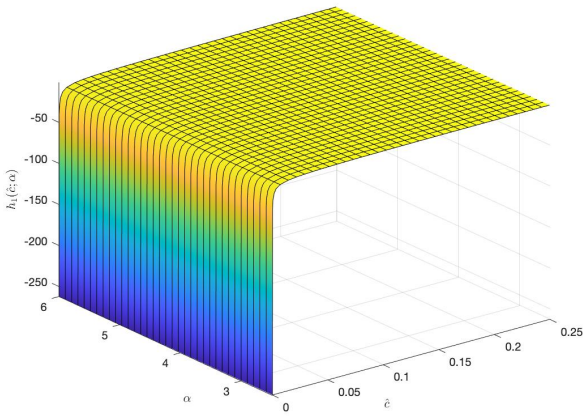


Figure 4: First order derivative with respect to cutoff

A Appendix

Definition A.1. A correspondence $\Gamma : X \rightrightarrows Y$ is upper-hemicontinuous at x if and only if for every sequence $x_\bullet = (x_n)_{n=1}^{\infty} \in X$, $y_n \in \Gamma(x_n)$ with $x_\bullet \to x$ and $y_\bullet = (y_n)_{n=1}^{\infty} \to y \in Y$, we have $y \in \Gamma(x)$.

Denote the set of all sequences associated with a monotone interval partition by $K \equiv \left\{\{a_1..,a_{N-1}\} \in [\underline{c},\bar{c}]^{N-1} : a_1 \leq a_2 \leq ... \leq a_{N-1}\right\}$. We first construct the consumer's best response in terms of search decisions. Let $P = \prod_{k=1}^{N}[0,\bar{v}_k]$ be the product price space. Define a correspondence $\mathbf{C} : P \rightrightarrows [\underline{c},\bar{c}]^{\frac{1}{2}N(N-1)}$ as follows. For each $\mathbf{P}^e \in P$, one can construct sequences of cutoffs

$$\hat{\mathbf{c}} = \{\hat{c}_{ij} \in [\underline{c},\bar{c}] : i < j \text{ and } i,j = 1,2,...,N\} \in [\underline{c},\bar{c}]^{\frac{1}{2}N(N-1)}$$

satisfying the following condition.

- If $v_i^*(c) - P_i^e = v_j^*(c) - P_j^e$ for some $c \in [\underline{c},\bar{c}]$, set $\hat{c}_{ij}$ to be any one of the solution.
- If $v_i^*(c) - P_i^e > v_j^*(c) - P_j^e$ for all $c \in [\underline{c},\bar{c}]$, set $\hat{c}_{ij} = \bar{c}$.
- If $v_i^*(c) - P_i^e < v_j^*(c) - P_j^e$ for all $c \in [\underline{c},\bar{c}]$, set $\hat{c}_{ij} = \underline{c}$.

Denote by $\mathbf{C}(\mathbf{P}^e)$ the set of cutoff sequences consistent with the conditions above. That is, for each pair $i,j = 1,2,...,N$ with $i < j$,

$$\mathbf{C}_{ij}(\mathbf{P}^e) = \begin{cases} \{\bar{c}\} & \text{if } v_i^*(c) - P_i^e > v_j^*(c) - P_j^e \text{ for all } c \in [\underline{c},\bar{c}] \\ \{\underline{c}\} & \text{if } v_i^*(c) - P_i^e < v_j^*(c) - P_j^e \text{ for all } c \in [\underline{c},\bar{c}] \\ \text{solutions to } v_i^*(c) - P_i^e = v_j^*(c) - P_j^e & \text{otherwise} \end{cases}$$

Note that if $\mathbf{C}_{ij}(\mathbf{P}^e)$ takes multiple values, i.e., $v_i^*(c) - P_i^e = v_j^*(c) - P_j^e$ admits mutliple solutions, we know from the proof of the lemma 3.2 that this can happen only over an interval of search costs. It is clear that $\mathbf{C}(\mathbf{P}^e)$ is compact- and convex-valued.

Lemma A.1. $\mathbf{C}(\mathbf{P}^e)$ *is an upper-hemicontinuous correspondence.*

Proof. Take arbitrary sequences $\left\{\mathbf{P}^{e,k}\right\}$ and $\left\{\hat{\mathbf{c}}^k\right\}$ such that $\mathbf{P}^{e,k} \to \mathbf{P}^e$, $\hat{\mathbf{c}}^k \in \mathbf{C}\left(\mathbf{P}^{e,k}\right)$ and $\hat{\mathbf{c}}^k \to \hat{\mathbf{c}}$. Take a pair of designs i,j with $i < j$. If $\hat{c}_{ij} \notin \{\underline{c},\bar{c}\}$, then it is necessary that $\hat{c}_{ij}^k$ satisfies

$$v_i^*\left(\hat{c}_{ij}^k\right) - P_i^{e,k} = v_j^*\left(\hat{c}_{ij}^k\right) - P_j^{e,k}, \text{ for } k \text{ sufficiently large.}$$

Using the continuity of v_i^* and v_j^*, in the limit, $v_i^*(\hat{c}_{ij}) - P_i^e = v_j^*(\hat{c}_{ij}) - P_j^e$.

Next suppose $\hat{c}_{ij} = \underline{c}$. At least one of the following subsequences exist: $\left\{c_{ij}^{k_n}\right\}$ with $c_{ij}^{k_n} = \underline{c}$, or $\left\{c_{ij}^{k_n}\right\}$ with $c_{ij}^{k_n} > \underline{c}$. In the former case, $v_i^*\left(c_{ij}^{k_n}\right) - P_i^{e,k_n} \leq v_j^*\left(c_{ij}^{k_n}\right) - P_j^{e,k_n}$ for all n. By the continuity of v_i^* and v_j^*, in the limit, $v_i^*(\underline{c}) - P_i^e \leq v_j^*(\underline{c}) - P_j^e$. It follows from corollary 3.2 that $v_i^*(c) - P_i^e \leq v_j^*(c) - P_j^e$ for all $c \in [\underline{c},\bar{c}]$. In the latter case, $v_i^*\left(c_{ij}^{k_n}\right) - P_i^{e,k_n} = v_j^*\left(c_{ij}^{k_n}\right) - P_j^{e,k_n}$ for all n. By the continuity of v_i^* and v_j^*, in the limit, $v_i^*(\underline{c}) - P_i^e = v_j^*(\underline{c}) - P_j^e$. The case of $\hat{c}_{ij} = \bar{c}$ is symmetric to that above and omitted. We have thus established that $\hat{\mathbf{c}} \in \mathbf{C}(\mathbf{P}^e)$. □

For each $\hat{\mathbf{c}} \in \mathbf{C}(\mathbf{P}^e)$, one can construct a nondecreasing cutoff sequence by $a_i(\hat{\mathbf{c}}) = \max\{a_{i-1}, \min\{\hat{c}_{ij} : j > i\}\}$, $i = 1, 2, .., N-1$. The monotone interval partition is characterized by cutoffs $\mathbf{a} = \{a_1, a_2, ..., a_{N-1}\} \in K$ describes an optimal search decision of the consumer in response to price vector $\mathbf{P}^e$. (It should be apparent that $\mathbf{a} \in \mathbf{A}(\mathbf{P}^e)$ and $\{\hat{c}_{ij} \in [\underline{c}, \bar{c}] : i < j \text{ and } i, j = 1, 2, ..., N\}$ are two equivalent ways of describing the optimal search.) Denote by $\mathbf{A}(\mathbf{P}^e)$ the set of all nondecreasing cutoff sequences that describe the consumer's optimal search given price vector $\mathbf{P}^e$. Specifically,

$$\mathbf{A}(\mathbf{P}^e) \equiv \left\{ \begin{array}{c} (a_1, a_2, ..., a_N) : a_1 = \min\{\hat{c}_{1j} : j > 1\}, a_2 = \max\{a_1, \min\{\hat{c}_{2j} : j > 2\}\}, ..., \\ a_{N-1} = \max\{a_{N-2}, \min\{\hat{c}_{N-1,j} : j > N-1\}\}, \hat{\mathbf{c}} \in \mathbf{C}(\mathbf{P}^e) \end{array} \right\}.$$

It follows from the continuity of each a_i in $\{\hat{\mathbf{c}}\}$ that $\mathbf{A}(\mathbf{P}^e)$ is also upper hemicontinuous and compact-valued.

Lemma A.2. $\mathbf{A}(\mathbf{P}^e)$ *is upper-hemicontinuous.*

Proof. Take arbitrary sequences $\{\mathbf{P}^{e,k}\}$ and $\{\mathbf{a}^k\}$ such that $\mathbf{P}^{e,k} \to \mathbf{P}^e$, $\mathbf{a}^k \in \mathbf{A}(\mathbf{P}^{e,k})$ and $\mathbf{a}^k \to \mathbf{a}$. As $\mathbf{a}^k \in \mathbf{A}(\mathbf{P}^{e,k})$ for each k, there exists $\hat{\mathbf{c}}^k \in \mathbf{C}(\mathbf{P}^{e,k})$ such that $a_i^k = a_i(\hat{\mathbf{c}}^k) = \max\{a_{i-1}^k, \min\{\hat{c}_{ij}^k : j > i\}\}$ for all $i = 1, 2., ..., N-1$. By Bolzano–Weierstrass, there exists a subsequence $\{\hat{\mathbf{c}}^{k_n}\}$ such that $\hat{\mathbf{c}}^{k_n} \to \hat{\mathbf{c}}$ for some $\hat{\mathbf{c}} \in [\underline{c}, \bar{c}]^{\frac{1}{2}N(N-1)}$. We know from the previous lemma (uhc of $\mathbf{C}$) that $\hat{\mathbf{c}} \in \mathbf{C}(\mathbf{P}^e)$. By the continuity of $a_i(\cdot)$, $a_i(\hat{\mathbf{c}}) = a_i$ for each i, and so $\mathbf{a} \in \mathbf{A}(\mathbf{P}^e)$. □

Lemma A.3. $\mathbf{A}(\mathbf{P}^e)$ *is convex-valued.*

Proof. Fix an arbitrary $\mathbf{P}^e$. An alternative way to express $\mathbf{A}(\mathbf{P}^e)$ is as follows. Let $A_i(\mathbf{P}^e)$ be all possible values at the i-th entry of the set $\mathbf{A}(\mathbf{P}^e)$ — this is the set of all possible cutoffs for design s_i. Let $\hat{c}_i(\mathbf{P}^e) \equiv \max\{\min\{\hat{c}_{ij} : j > i\} : \hat{\mathbf{c}} \in \mathbf{C}(\mathbf{P}^e)\}$ be the highest search cost consistent with (weakly) preferring design s_i to higher designs. Below we suppress the dependence on $\mathbf{P}^e$ for notational simplicity.

For design s_1, A_1 is either a singleton $\{\hat{c}_1\}$, an interval $[\underline{c}, \hat{c}_1]$ or an interval $[\tilde{c}_1, \hat{c}_1]$.

1. In the first case, set $\hat{A}_1 = \{\hat{c}_1\}$. Continue the procedure for search costs in the interval $[\hat{c}_1, \bar{c}]$. Namely, A_2 is either a singleton $\{\hat{c}_2\}$, an interval $[\hat{c}_1, \hat{c}_2]$ or an interval $[\tilde{c}_2, \hat{c}_2]$.... Define
$$\hat{A}_2 = \{(a_1, a_2) : a_1 = \hat{c}_1, a_2 \in A_2\}$$

2. In the second case, it is necessary that there exists some designs such that $v_1^*(c) - P_1^e = v_j^*(c) - P_j^e$ over the interval $[\underline{c}, \hat{c}_1]$; let $j \geq 2$ be the smallest such index. For any $j' < j$ (if exists), it is necessary that $v_j^*(c) - P_j^e > v_{j'}^*(c) - P_{j'}^e$ for all c, so $\hat{c}_{j'} = \underline{c}$. Moreover, $\hat{c}_j \geq \hat{c}_1$. There are two possibilities about A_j: it is either a singleton $\{\hat{c}_j\}$ or is an interval $[\underline{c}, \hat{c}_j]$. If $A_j = \{\hat{c}_j\}$, set
$$\hat{A}_j = \{(a_1, ..., a_j) : a_1 \in [\underline{c}, \hat{c}_1], a_2 = ... = a_{j-1} = a_1, a_j = \hat{c}_j\},$$
which is clearly convex. If A_j is an interval $[\underline{c}, \hat{c}_j]$, set
$$\hat{A}_j = \{(a_1, ..., a_j) : a_1 \in [\underline{c}, \hat{c}_1], a_2 = ... = a_{j-1} = a_1, a_j \in [a_1, \hat{c}_j]\},$$
which is clearly convex. Here, it is necessary that there is some other design such that $v_j^*(c) - P_j^e = v_k^*(c) - P_k^e$ over the interval $[\underline{c}, \hat{c}_j]$; let $k \geq j$ be the smallest such index.

For any $k' < k$ (if exists), it is necessary that $v_j^*(c) - P_j^e > v_{k'}^*(c) - P_{k'}^e$ for all c, so $\hat{c}_{k'} = \underline{c}$. Moreover, $\hat{c}_k \geq \hat{c}_j$.There are three possibilities about A_k: it is either a singleton $\{\hat{c}_k\}$, an interval $[\underline{c}, \hat{c}_k]$, or $[\tilde{c}_k, \hat{c}_k]$....

3. If

In the former case, set $\hat{A}_1 = \{\hat{c}_1\}$. Continue the procedure for search costs in the interval $[\hat{c}_1, \bar{c}]$. Namely, A_2 is either a singleton $\{\hat{c}_2\}$ or an interval $[\hat{c}_1, \hat{c}_2]$....

In the latter case, it is necessary that there exists some designs such that $v_1^*(c) - P_1^e = v_j^*(c) - P_j^e$ over the interval $[\underline{c}, \hat{c}_1]$; let $j \geq 2$ be the smallest such index.Continue the procedure for search costs in the interval $[\hat{c}_j, \bar{c}]$. Namely, A_{j+1} is either a singleton $\{\hat{c}_{j+1}\}$ or an interval $[\hat{c}_j, \hat{c}_{j+1}]$....

If A_j is an interval $[\underline{c}, \hat{c}_j]$, set

$$\hat{A}_j = \{(a_1, ..., a_j) : a_1 \in [\underline{c}, \hat{c}_1], a_2 = ... = a_{j-1} = a_1, a_j \in [a_1, \hat{c}_j]\},$$

which is clearly convex. Here, it is necessary that there is some other design such that $v_j^*(c) - P_j^e = v_k^*(c) - P_k^e$ over the interval $[\underline{c}, \hat{c}_j]$; let $k \geq j$ be the smallest such index. For any $k' < k$ (if exists), it is necessary that $v_j^*(c) - P_j^e > v_{k'}^*(c) - P_{k'}^e$ for all c, so $\hat{c}_{k'} = \underline{c}$. Moreover, $\hat{c}_k \geq \hat{c}_j$. There are two possibilities about A_k: it is either a singleton $\{\hat{c}_k\}$ or is an interval $[\underline{c}, \hat{c}_k]$....

It is clear that every $\hat{A}_j, j = 1, 2, ..., N-1$, is convex. Proceed inductively until reaching $\hat{A}_{N-1} = \mathbf{A}(\mathbf{P}^e)$.

□

For each nondecreasing cutoff sequences $\mathbf{a} = \{a_1, a_2, ..., a_N\}$, let

$$\pi_i(\mathbf{a}, \mathbf{P}^e) \equiv \arg \max_{p_i \in [0, \bar{v}_i]} p_i \times \int_{[a_{i-1}, a_i]} \frac{[1 - F_i(v_i^*(c) + p_i - P_i^e)]}{1 - F_i(v_i^*(c))} dG(c)$$

be the set of optimal prices for firms adopting design i, assuming that only consumers with search cost in the interval $[a_{i-1}, a_i]$ search for this design. By assumption, the objective function above is continuous, and quasiconcave in p_i. The argmax correspondence is thus upper-hemicontinuous, and compact- and convex-valued.

Define a correspondence $\Gamma : P \times K \rightrightarrows P \times K$ by $\Gamma(\mathbf{P}^e, \mathbf{a}) \equiv (\Pi_{i=1}^n \pi(\mathbf{a}, \mathbf{P}^e), \mathbf{A}(\mathbf{P}^e))$. As each of π_i and $\mathbf{A}$ are convex-valued, compact-valued, and upper-hemicontinuous, it has a fixed point by the Kakutani fixed-point theorem. At this fixed point, the firms and consumers are playing a mutual best response.

References

Anderson, Simon P, and Regis Renault. 1999. "Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model." *The RAND Journal of Economics*, 719–735.

Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat. 2012. "Search, design, and market structure." *American Economic Review*, 102(2): 1140–1160.

Brynjolfsson, Erik, Yu Hu, and Duncan Simester. 2011. "Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales." *Management Science*, 57(8): 1373–1386.

Brynjolfsson, Erik, Yu Jeffrey Hu, and Michael D Smith. 2006. "From niches to riches: Anatomy of the long tail." *Sloan management review*, 47(4): 67–71.

Cachon, Gérard P, Christian Terwiesch, and Yi Xu. 2008. "On the effects of consumer search and firm entry in a multiproduct competitive market." *Marketing Science*, 27(3): 461–473.

Chen, Yongmin, and Tianle Zhang. 2011. "Equilibrium price dispersion with heterogeneous searchers." *International Journal of Industrial Organization*, 29(6): 645–654.

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

Harris, Judy, and Edward A Blair. 2006. "Consumer preference for product bundles: The role of reduced search costs." *Journal of the Academy of Marketing Science*, 34(4): 506–513.

Johnson, Justin P, and David P Myatt. 2006. "On the simple economics of advertising, marketing, and product design." *American Economic Review*, 96(3): 756–784.

Khaledi, Baha-Eldin, and Subhash Kochar. 2002. "Stochastic orderings among order statistics and sample spacings." In *Uncertainty and optimality: Probability, statistics and operations research*. 167–203. World Scientific.

Kuksov, Dmitri. 2004. "Buyer search costs and endogenous product design." *Marketing Science*, 23(4): 490–499.

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

Prékopa, András. 1973. "On logarithmic concave measures and functions." *Acta Scientiarum Mathematicarum*, 34: 335–343.

Shaked, Moshe, and J George Shanthikumar. 2007. *Stochastic orders.* Springer.

Stahl, Dale O. 1989. "Oligopolistic pricing with sequential consumer search." *The American Economic Review*, 700–712.

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

Weitzman, Martin L. 1979. "Optimal search for the best alternative." *Econometrica: Journal of the Econometric Society*, 641–654.

Wolinsky, Asher. 1986. "True Monopolistic Competition as a Result of Imperfect Information." *The Quarterly journal of economics*, 101(3): 493–511.

Zhou, Jidong. 2014. "Multiproduct search and the joint search effect." *American Economic Review*, 104(9): 2918–39.