// NO.24-033 | 05/2024

DISCUSSION PAPER

// JACOPO GAMBATO

Consumer Search and Firm Strategy With Multi-Attribute Products

Consumer Search and Firm Strategy with Multi-Attribute Products*

Jacopo Gambato†

April 2, 2024

Abstract

I analyze a model of directed search in which a consumer inspects a finite number of products sharing attributes with each others. The consumer discovers her valuation for the attributes of the inspected products and adapts her search strategy based on what she has learned. The consumer anticipates the optimal paths that arise after different realizations; this generates a search rule that accounts for learning systematically. In this search environment, a multiproduct seller commits to a menu of horizontally differentiated products. The seller can exploit the fact that the emerging search paths reveal the consumer's preferences: by setting different prices for *ex ante* identical products, the seller can encourage specific paths to arise and exploit the information that the consumer learned through search. In some cases, the seller optimally limits the set of available products.

Keywords: consumer search, directed search, learning, multiproduct monopoly, pricing, product portfolio

JEL Codes: D42, D83, L12, L15

*I thank the Department of Economics at University of Mannheim and the IOEK department and Junior Research Group for Competition and Innovation at ZEW for guidance and support. In particular, Ithank Martin Peitz, Bernhard Ganglmair, Nicolas Schutz, Volker Nocke, Anton Sobolev, and Atabek Atayev; I also thank Daniel Savelle, Axel Gautier, José Moraga-Gonzales, Justus Preusser, Sander Onderstal, Marco Haan, Julia Reimer, Heiko Karle, Michelangelo Rossi, David Ronayne, Willy Lefez, as well as the participants to the 2023 MaCCI Annual Conference, 11th CRC Retreat, 2023 CLEEN Conference, 12th Workshop on Consumer Search and Switching Costs, 2023 CRESSE Conference, 21st ZEW ICT Conference, 50th EARIE Annual Conference, EEA-ESEM 2023 Conference, XXXVII Jornadas de Economía Industrial, and VfS Annual Conference 2023 for their insightful comments. Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B04) is gratefully acknowledged.

†University of Mannheim, ZEW Mannheim, MaCCI - Jacopo.Gambato@zew.de

1. Introduction

Multiproduct firms are important players in many economic environments. The wide array of strategic choices at their disposal, however, makes them a difficult subject to study. Much has been written about the risk multiproduct firms run to have their products cannibalize each others' demand.[1] Less attention has been devolved to the synergy arising when the products offered by such a firm are correlated and to the effect this has on how consumers interact with the product menu. To study this dimension of the firm's strategic considerations, I develop a framework that allows products to be correlated through shared attributes. In this environment, I study the optimal pricing and menu composition that is chosen by a multiproduct monopolist, and how these choices affect consumer learning when there are search frictions.

The consumer search literature has highlighted the role of search frictions as determinants of market outcomes. The effect of these frictions for intra-firm search, however, has been so far understudied.[2] I contribute to the literature by incorporating correlation across the products offered by a single firm to study how this affects consumers' optimal search and how firms would condition their strategic decisions on it. I define a product as an agglomerate of attributes (Lancaster, 1966) and allow these attributes to be shared across products. These shared attributes are apparent to the consumer beforehand and make two products sharing an attribute perfectly correlated in that dimension. In words, consumers inspect a product knowing how the resulting realization will affect their perception of all other products.

This structure allows for novel insights on consumer search behavior. First, since products share different attributes with different products, the framework allows for the optimal search path to arise endogenously: consumers who learn different things about the same product will choose to inspect different products next. As a consequence, the optimal search process is forward looking – that is, the different implied search paths arising after different realizations matter – so that inspecting any product is more valuable than it would be in isolation because of the existence of correlated alternatives. I further show that when direction of search is endogenously determined, consumers "self-select" towards their preferred option in a way that matches recent evidence of consumers navigating the attribute space strategically as in Hodgson and Lewis (2020).

A multiproduct monopoly firm commits to a menu and posts products' prices anticipating the consumer optimal search process. Prices are posted and contribute to determining the order in which consumer search for their preferred option.[3] Because the outcome of each inspection instructs the next, each inspection reveals the consumer's learned

[1]This consideration is prominently pointed out, for example, in Nocke and Schutz (2018).

[2]Prominent exceptions are Petrikaitė (2018) and Nocke and Rey (2023).

[3]Price advertisement also resolves hold-up problems arising with monopoly pricing in the presence of search frictions as shown in Anderson and Renault (2006); a more in-depth analysis on the matter can be found in Robert and Stahl (1993) and Konishi and Sandfort (2002).

preferences. The firm can price products differently to encourage consumers to self-sort based on the preferences they learn about through the search process, a mechanism reminiscent of that highlighted in Mayzlin and Shin (2011). Unlike in Mayzlin and Shin (2011), however, different prices can emerge in my framework even if products are *ex ante* identical from the consumer's perspective.

Differential prices might induce the consumer to deviate from the seller's preferred order of search. I show that in some cases, when the product menu is relatively small, the seller has an incentive to restrict the supply by removing specific products from the menu and, with them, alternative search paths available to the consumer.[4] The menu restriction induces the firm's preferred order of search to arise, and it is an optimal strategy when the likelihood of a positive realization is high and search is cheap. Whenever this is the case, the seller strictly prefers a uniform pricing strategy over setting different prices for different products. Therefore, both uniform and differential prices can arise in equilibrium.

The results highlight the ability of a multiproduct firm to steer consumers through strategic menu selection. By anticipating how a consumer would react after observing a product, the seller can encourage search towards better suited products, and profit off the consumer's incentive to find good matches. The seller wants the consumer to keep searching whenever possible: what is learned through inspection of a product makes the consumer fine-tune her selection.[5] The seller can increase profits by setting higher prices along paths consistent with positive realizations without discouraging the consumer to search on paths consistent with negative ones. Through this model, I propose strategic menu selection as a new mechanism through which consumers can be steered in their consumption choices.

The model also allows to study product design under a new light when the assumption of all products being *ex ante* identical is relaxed. In particular, it highlights the relative likelihood of sampling general and widely appreciated versus provocative and niche products. In an extension aimed at studying this additional dimension of product positioning, I propose a simple alternative framework that qualifies the view, dominant in the literature, that firms have an incentive to choose polarized designs, that is, either as generic or as daring as possible. Instead, when buyers have agency over the order in which they sample different options, intermediate designs can gain prominence over extremely bland and extremely provocative products.

On the same topic, many markets have seen in recent years floods of new designs tailored for very different tastes and a wider variety and combinations of features than ever before - the so called long tail effect (Bar-Isaac et al., 2012). Endogenous search order naturally explains the rise in popularity of niche designs: with search becoming less and

[4] This strategic choice is fundamentally different from that highlighted by Johnson and Myatt (2002): The authors consider menu pruning in response to entry, while no competitor is ever present in my framework.

[5] The learning component, then, leads to an outcome opposite to that shown in Petrikaitė (2018); the multiproduct monopoly firm studied by the author has an incentive to obfuscate options to increase the probability of selling more expensive alternatives.

less costly less widely appealing but more rewarding products have become less punishing to inspect. Further, as the added variety makes searching more rewarding through the learning process detailed above, more buyers might be induced to start searching and make already popular products reach even more consumers.

The rest of the paper is structured as follows: after reviewing the related literature, I present the framework (Section 2) and characterize the optimal search process with multiple attributes and the learning process they imply (Section 3). Afterwards, I solve the problem of a monopoly seller that selects which products to make available and their prices (Section 4). In Section 5, I extend the model explore the choice of introducing mainstream and niche products in light of this search dynamics. I conclude in Section 6.

Related literature This paper relates to several strands of literature. First, it contributes to the ordered consumer search literature pioneered by Weitzman (1979). Weitzman characterizes the optimal process for a consumer costly searching among n independent boxes. Each box is characterized by a reservation value, a score representing the value that would make the consumer indifferent between opening the box and keeping a certain reward equal to the score. The optimal search order has the consumer opening boxes from the highest to the lowest score. The consumer optimally stops when no unopened box has a score higher than the highest past realization.

The role of search order on market outcomes has been studied extensively in oligopoly settings: Choi et al. (2018) and Haan et al. (2018) study the effect of posted prices on search order. Because sellers want to undercut each other to gain prominence in the search order, pinning down an equilibrium requires consumers to be heterogeneous enough, specifically in the form of different mean expected qualities. Anderson et al. (2020) obtains similar results by introducing heterogeneity through the search cost distribution. The features instructing the order of search in these models are, however, never shared between products. Because the multiproduct monopoly seller I focus on does not have an incentive to undercut himself, moreover, heterogeneity in consumers' characteristics is not necessary in my setting.

Ke and Lin (2022) and Bao et al. (2022) propose a framework similar in concept to the one analyzed here – a discrete number of products sharing one of their two attributes. Ke and Lin (2022), the closest of the two to this paper, studies the relationship between correlation in search and complementarity of the products available, the same topic explored in Anderson et al. (2021). Bao et al. (2022) on the other hand, assumes that the representative consumer inspecting products cannot distinguish the role of each attribute in the utility each product generates, and studies optimal search with Bayesian updating. Conditional search order is also the main subject of Doval (2018)'s famous extension to Weitzman (1979)'s seminal paper on optimal search in which it is shown that allowing the consumer to keep an unopened box without discovering it changes the game to the point that the original result can break down. As my approach more prominently highlights the role of leaning through the search process, however, this paper relates more than the aforementioned ones to this rapidly growing literature (Garcia and Shelegia, 2018, Greminger, 2022, Preuss, 2023).

The paper further contributes to the wide literature on multiproduct firms. Earlier contributions addressed several possible strategies available to this kind of seller. Some, like Mussa and Rosen (1978), focused on price discrimination with vertically differentiated products. Others, like Eaton and Lipsey (1979), discuss market pre-emption through introduction of horizontally differentiated options. Other notable example relate to R&D expenditure (Lin, 2004, Lambertini and Mantovani, 2009) and bundling of products (McAfee et al., 1989).

This paper contributes to the literature on the interaction between menu selection and pricing (Brander and Eaton, 1984, Johnson and Myatt, 2002, Nocke and Schutz, 2018). Novel to the literature is the inclusion of correlation across the products offered by the firm. Correlated products allow consumers to learn their preferences as they inspect options and, therefore, the presence of correlated products affects the value of inspecting each product in isolation. Moreover, correlation affects the value of inspecting products defined by attributes associated with different levels of expected popularity and, therefore, the interaction between product design and search (Bar-Isaac et al., 2012).

Finally, the paper contributes to the literature of pricing in search. The seminal Wolinsky (1986) model, and most of the literature that followed, focuses on competitive settings.[6] Instead, I study within-firm directed search in a monopoly setting as in recent contributions by Petrikaitė (2018) and Nocke and Rey (2023). The latter studies the incentives of a multiproduct seller to "garble" product information to induce consumers to search longer. Because search costs are assumed to be fixed, the firm has no incentive to price discriminate. Petrikaitė (2018), instead, shows that a multiproduct seller can steer consumers towards expensive products by obfuscating cheaper options. In my framework, steering can arise in the form of differential prices being optimally set by the seller without strategic obfuscation of the products made available. The paper, then, is related to the growing steering literature as well (e.g. Ichihashi, 2020, who also considers a monopoly setting).[7]

2. Simplified Framework

The products. I consider an industry with products differentiated with respect to two attributes. A product (i, j) is identified by attributes $A_i \in A$ and $B_j \in B$. I restrict A and B to come in two variants each; Each attribute A_i can be found combined with all

[6]This is true also for most papers studying multiproduct firms; see for example Zhou (2014), Rhodes (2015), Rhodes et al. (2021)

[7]Other notable contributions, although less closely related, can be found in De Corniere and Taylor (2019), Teh and Wright (2022), and Heidhues et al. (2023).

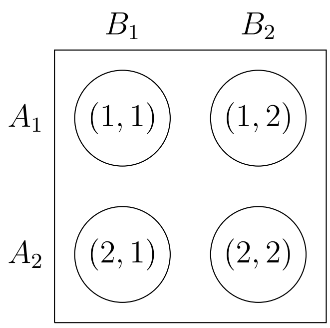

Figure 1: *Products in the same row (resp. column) share attribute* A_i *(resp.* B_j*).*

attributes B_j, $i, j \in \{1, 2\}$, and *vice versa*. One can visualize the products as displayed in a grid, with the rows representing the A attributes, the columns representing the B attributes, and the cells representing products defined by a specific combination of A and B attributes as depicted in Figure 1. Products are only differentiated horizontally through their attribute compositions and are otherwise identical in quality.

The consumer. A representative, risk-neutral consumer (she) has unit demand, is aware of the available products and their attribute composition, and can inspect the products in any order she likes. The consumer has no prior knowledge of her preferences over the available attributes; she learns the realization of each attribute separately by inspecting a product characterized by it. In line with existing models,[8] I assume that *ex post* utility generated by a generic product (i, j) takes the form:

$$u(A_i, B_j) = A_i + B_j = u_{i,j}.$$

I further assume attributes to follow a Binomial distribution: each attribute $y \in A \cup B$ is either a match, generating *ex post* utility one with probability $\alpha \in (0, 1)$, or it is not, generating utility zero instead. The assumption that attributes enter $u_{i,j}$ additively crucially implies that there are no complementarities between attributes: once an attribute is discovered, its realized value affects all products that are defined by it in the same way. The expected utility of an unsampled product (i, j) is then:

$$E[u_{i,j}] = \alpha + \alpha = 2\alpha.$$

Expected utility of a product (i, j) sharing an attribute with a previously sampled product, say A_i, but not the other, is instead:

$$E[u_{i,j}] = A_i + \alpha.$$

In this environment, I study the optimal sequential search process with free recall: a

[8]For example: Choi et al. (2018) and Greminger (2022).

consumer can always go back to a previously inspected product at no additional cost. The cost of inspecting a product is indexed by the constant $s \in (0, 2\alpha)$. The consumer learns the value of each attribute separately after inspecting a product defined by it. Finally, the consumer's outside option is normalized to $u_0 = 0$.

The seller. A multiproduct monopoly seller (he) selects which of the possible products to make available to the representative consumer (that is, he selects $\tilde{N} \subseteq N$), and their respective prices. He is aware of the match probability α and search costs s. The seller can influence the search pattern over available products through prices. Prices are set before the search process starts, cannot be changed, and are observed costlessly by the consumer before she starts searching. All production costs are equal to zero.

Timing and equilibrium concept. The timing of the interaction can be summarized as follows:

1. The seller selects $\tilde{N} \subseteq N$ products to make available and price vector $\mathbf{p}(\tilde{N})$.
2. The consumer observes $\tilde{N}$, $\mathbf{p}(\tilde{N})$, chooses between searching and her outside option, and, if she searches, what to inspect.
3. After each inspection, the consumer chooses between stopping and keeping searching (and what to inspect next) until she either purchases an inspected product or leaves without making a purchase.

I consider Subgame Perfect Equilibria: because the seller commits to menu and prices before the search process starts, and because prices are posted, there is no need to model beliefs explicitly in this environment.

3. A Simple Model of Multi-Attribute Search

Because $A_i \in \{A_1, A_2\}$, $B_j \in \{B_1, B_2\}$, the product space N consists of four products:

$$N = \{(1,1), (1,2), (2,1), (2,2)\}.$$

To illustrate the search dynamic in isolation, I start with the assumption that prices are exogenously set at zero; this assumption will be relaxed in the next section. The consumer can inspect any product in $\tilde{N}$; I start from the case in which $\tilde{N} \equiv N$. At any given point of the search sequence, the set of available products can be partitioned in the set of inspected products, I, and uninspected products, $\tilde{N} \setminus I$.

Updating expected utilities. Suppose that the consumer already inspected one of the products. Because all products are *ex ante* identical, inspecting $(1, 1)$ first is without loss of generality.[9] Whenever the product to inspect can be chosen randomly without loss of generality, I assume that products are inspected in increasing order of their indices. After the first inspection, the consumer has learned realizations A_1 and B_1. Which of the remaining products should be inspected next, if any?

In this simplified framework, it is straightforward to show that the consumer would want to search keeping an attribute she has learned to have positive valuation for (if search costs are low enough), and ignoring one for which she has valuation zero. Formally, given realization $u_{1,1} = A_1 + B_1$, the consumer updates her expectations for the remaining product according to:

$$E(u_{1,2}|I = \{(1, 1)\}) = A_1 + \alpha, \qquad E(u_{2,1}|I = \{(1, 1)\}) = \alpha + B_1,$$
$$E(u_{2,2}|I = \{(1, 1)\}) = 2\alpha.$$

The consumer would next choose to inspect the product with the highest updated expected value as long as:

$$\max_{(i,j)\in N\setminus I} E(u_{i,j}|I) - s > \max_{(i,j)\in I} u_{i,j},$$

which immediately leads to the optimal follow-up search for each possible realization of $(1, 1)$:

- if $A_1 = B_1 = 0$, $(2, 2)$ is searched next; no other search can take place because $A_2 \geq A_1$ and $B_2 \geq B_1$.
- if $A_1 = B_1 = 1$, the consumer stops at $(1, 1)$ because $A_1 \geq A_2$ and $B_1 \geq B_2$.
- if $A_1 > B_1$, $(1, 2)$ is searched next (if $\alpha > s$); no other search can take place because $A_1 \geq A_2$ and B_2 is shared between $(1, 2)$ and $(2, 2)$.
- if $A_1 < B_1$, $(2, 1)$ is searched next (if $\alpha > s$); no other search can take place because $B_1 \geq B_2$ and A_2 is shared between $(2, 1)$ and $(2, 2)$.

Expected utility of searching. Different realizations lead to different search paths being taken every time a new product is inspected. These conditional search paths emerge predictably, and all realizations generate unambiguously an optimal path forward. In turn, this implies that a rational consumer would account for the likelihood of these different paths emerging, and the expected utility they are associated with, when deciding whether

[9] All products share each attribute that characterizes it with another product, and for all products there is one other product that shares no attributes with it. Therefore, all product are *ex ante* identical as long as prices are uniform

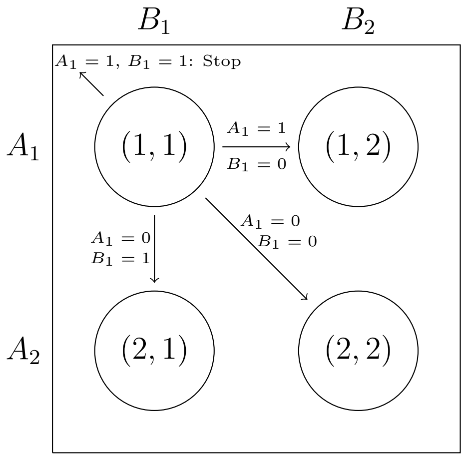

Figure 2: *Optimal search with binomial distribution and all products available, starting from* $(1,1)$.

to start searching or not. From the above, therefore, we obtain the expected utility of searching given the available products and the optimal search paths that can emerge:

$$E(u_{i,j}|I \equiv \emptyset) = 2\alpha^2 + 2\alpha(1-\alpha)\max\{1, 2\alpha + (1-\alpha) - s\} + (1-\alpha)^2(2\alpha - s) - s.$$

The first term refers to (i,j) being the best possible match ($u_{i,j} = 2$, with probability α^2). The second term refers to the eventuality of the consumer liking only one of the two attributes, and incorporates the possible second search that outcome would entail, which only takes place if $s \leq \alpha$. The third refers to the case in which $u_{i,j} = 0$ so that the product sharing no attributes with it would be inspected next. Figure 2 exemplifies the optimal search pattern.

4. Seller's Optimal Strategy

The seller's problem is twofold: he must set up prices to maximize profit, and he must select $\tilde{N}$ to generate trade opportunities. The two decisions are related. The consumer search path depends on the price she observes, and which prices would deter her from searching depend on the available products. In particular, the consumer is willing to search a product priced above its myopic expected value $2\alpha - s$ as long as the expected utility of searching from that point onward is non-negative. As shown above, this can be achieved when products that share attributes with each other are made available. A seller can, in principle, price products above their myopic expected value as long as he made available enough products to justify it.

The two decisions - menu selection and pricing - interact in non-obvious ways. Uniform prices, for example, cannot induce an order of search different from the one characterized above. If these uniform prices are too high, however, some search paths could end prematurely: even if products are identical *ex ante*, the first one searched - $(1,1)$ in the example above - carries more new information than every subsequent search that could arise. It follows that the highest price that would make two products not sharing attributes worth searching is different. If prices are not uniform, however, the consumer could adapt her optimal order of search in response: between a more expensive product for which she has positive information and a cheaper one for which she has no information, that she would inspect the former first is not obvious.

To study these different interactions, I solve the menu and pricing game of the seller considering uniform and differential prices separately. I show that the seller can always manipulate prices to induce a specific ordering of the consumer search. Moreover, I show that the seller has an incentive to strategically restrict the menu of available products to induce his preferred order of search to arise when search is cheap.

4.1. Uniform Prices

Under uniform prices, the seller's trade-off is clear-cut. He wants to raise prices to capitalize on any positive outcome of the consumer search, and he wants to lower prices to incentivize inspections after negative outcomes. The seller is indifferent regarding which product is ultimately purchased, as long as one is. For this reason, I start by assuming that all products are available: $\tilde{N} \equiv N$. I then show the seller's incentive to restrict the menu and the effect this choice has on consumer search.

Consider a generic uniform price level p^u. The seller wants to set the highest level p^u conditional on certain constraints implied by the consumer search process not being violated. Given the optimal search pattern identified in the section above, the expected utility of performing the first inspection is:

$$\begin{aligned}
&E(u_{1,1}|I \equiv \emptyset) = \alpha^2 \max\{2 - p^u, 0\} - s \\
&+ 2\alpha(1-\alpha) \max\{1 - p^u, \alpha \max\{2 - p^u, 0\} + (1-\alpha) \max\{1 - p^u, 0\} - s, 0\} \\
&+ (1-\alpha)^2 \max\{\alpha^2 \max\{2 - p^u, 0\} + 2\alpha(1-\alpha) \max\{1 - p^u, 0\} - s, 0\}.
\end{aligned} \tag{1}$$

That is: the value of inspecting $(1,1)$ is equal to the expected value generated by the search paths that are induced by the possible different realizations. These in turn depend on the relative value of s and α, over which the seller has no control over, and p^u.

At $p^u = 0$, the search problem of the consumer is identical to the one explored in the example above. As prices grow, however, some search paths become inaccessible. The first search path to be prevented by high prices is the one that arises conditional on a bad first match. Indeed, given observation $u_{1,1} = 0$, $(2,2)$ is searched as long as

$$E(u_{2,2}|u_{1,1} = 0) = \alpha^2 \max\{2 - p^u, 0\} + 2\alpha(1-\alpha) \max\{1 - p^u, 0\} - s \geq 0. \tag{2}$$

It is straightforward to show that there exists values p^u such that this condition is not satisfied but $E(u_{1,1}|I \equiv \emptyset)$ is positive: even if the consumer would not search after a bad realization of $(1,1)$, the presence of products sharing attributes with it makes it more likely to find something worth purchasing. As long as p^u is such that $E(u_{1,1}|I \equiv \emptyset)$ is non negative, the consumer can rationally start inspecting products. With this inspection, the consumer can discover that she likes both attributes, after which she always stop searching because she can find no better match. Alternatively, if the consumer likes only one attribute, she is interested in inspecting the other available product that shares it. Suppose $A_1 = 1$, $B_1 = 0$, and $p^u \leq 1$. The consumer would want to perform this additional search if and only if:

$$u_{1,1} = 1 - p^u \leq 1 + \alpha - s - p^u = E(u_{1,2}|I = \{(1,1)\}),$$

which is always satisfied if $s \leq \alpha$, that is, if inspecting a single attribute is worth the necessary search cost. If $s > \alpha$, that is, if s is higher than the expected gain of inspecting one attribute in isolation, the consumer would only ever inspect a product she knows nothing about. In this case, the presence of correlated products is immaterial: because no product can be reached after inspecting a different product with which it shares an attribute, the expected gain of inspecting a product is only ever its expected value. Therefore,

$$p^M = \frac{2\alpha - s}{\alpha(2-\alpha)},$$

(where the superscript M stands for "myopic") is the optimal price when $s > \alpha$.

Suppose now that $s \leq \alpha$. The seller can select one of two pricing profiles: on one hand, he can elect to price products in a way that encourages a follow-up search after a first bad realization. These prices must make a product just myopically worth searching, or, they must solve equation (2) with equality:

$$\mathbf{p^E} = \begin{cases} p_L^E = p^M & \text{if } \alpha^2 \leq s \leq \alpha, \\ p_H^E = \frac{2\alpha^2 - s}{\alpha^2} & \text{if } 0 < s < \alpha^2, \end{cases}$$

where E stands for "encourage", L stands for "low", and H stands for "high".

Alternatively, the seller can select higher prices that discourage search after a bad first realization. These prices must be strictly higher than the encouraging ones and lead to a lower probability of trade, but a higher return conditional on the consumer finding something to purchase. These prices are such that $E[u_{1,1}]|_{I\equiv\emptyset} = 0$, because for any higher price the consumer would not start searching:

$$\mathbf{p^D} = \begin{cases} p_L^D = \frac{2\alpha(1+(1-\alpha)(\alpha-s))-s}{\alpha(2-\alpha)} & \text{if } \frac{3\alpha^2-2\alpha^3}{1+2\alpha-2\alpha^2} \leq s \leq \alpha, \\ p_H^D = \frac{2\alpha(\alpha(3-2s)-(1-\alpha)s)-s}{\alpha^2(3-2\alpha)} & \text{if } 0 < s < \frac{3\alpha^2-2\alpha^3}{1+2\alpha-2\alpha^2}, \end{cases}$$

where D stands for "discourage".

Uniform price selection. Lower prices can always be selected whenever higher ones do not prevent the consumer from searching. The seller is, however, not interested in his products being inspected, but in his products being purchased. Trade is maximized for $p \leq 1$: any higher price requires the consumer to like both attributes in a product to purchase it. Notice that it holds:

$$p_L^E > 1 \iff 0 < s < \alpha^2.$$

Therefore, the price that maximizes search and trade can be identified as the minimum between p_L^E and 1. To simplify the notation, I define:

$$p_T = \min(p_L^E, 1),$$

where T stands for "trade", as p_T is the price that maximizes the probability of trade. Overall, when selecting p^{u*} among the candidate equilibrium prices displayed above, the seller chooses between maximizing search efforts, maximizing per-sale revenue, and maximizing probability of trade. Higher prices discourage search and reduce probability of trade for a given search pattern; lower prices encourage search but lead to lower revenue conditional on trade taking place.

By plugging in the various (feasible) prices for the various combinations of α and s and following the search path different prices induce according to Equation (1), one can obtain the expected profit of the seller. These profits can then be directly compared and lead to a unique equilibrium price for all possible combinations of α and s. In particular, when $s > \alpha$, the only candidate price and relative expected profit is:

$$p^M < 1 \quad \rightarrow \quad \pi^M = p_M \left(1 - (1-\alpha)^4\right).$$

Instead, when $s \leq \alpha$, the candidate prices obtained above lead to expected profits:

$$p_T \leq 1 \quad \rightarrow \quad \pi_L^E = p_T \left(1 - (1-\alpha)^4\right),$$

which maximizes probability of trade and is always valid,

$$p_L^D < 1 \quad \rightarrow \quad \pi_I^D = p_L^D \left(1 - (1-\alpha)^2\right),$$

which prevents any further inspection after a bad first realization if $0 < s < \frac{3\alpha^2 - 2\alpha^3}{1+2\alpha-2\alpha^2}$, but generates trade if any one inspected attribute is appreciated,

$$p_H^E > 1 \quad \rightarrow \quad \pi_H^E = p_H^E \left[\alpha^2(1 + 2(1-\alpha) + (1-\alpha)^2)\right],$$

which always allows for a second inspection if $0 < s < \alpha^2$, but requires the consumer to find a product to like in both attributes to lead to a purchase, and

$$p_H^D > 1 \quad \rightarrow \quad \pi_H^D = p_H^D \left[\alpha^2(1 + 2(1 - \alpha))\right],$$

which does not allow for another search after a bad first realization. In all cases, expected profit is calculated as price times the probability of trade generated because production costs are assumed to be equal to zero. The candidate prices reflect the relative importance of encouraging search and extracting rent conditional on search taking place. In particular, the seller trades off higher probability of trade by encouraging search and revenue conditional on trade taking place by discouraging it.

Intuitively, higher prices are preferable, for the seller, for low search costs and high probability of a match α. For such parameters the consumer is easily encouraged to start searching. If $s > \alpha$, there is only one candidate price, p^M, the lowest of the candidate prices. For $s \leq \alpha$, instead, which of the four candidate prices is selected depends on the relative value of s: the lower s is, the higher prices can be set without impeding search.

Implications for the optimal menu selection. Because prices can encourage or discourage search, they also determine the optimal product menu selection. When all products are available, any product can be rationally selected to be the first to inspect by the consumer. For high enough prices, however, not all products can be inspected after fixing a starting point. From the above discussion it emerges that if the seller optimally selects $p^{u*} = \mathbf{p^D}$, conditional on the consumer starting from $(1, 1)$, inspection of $(2, 2)$ could not rationally take place. Indeed, $(2, 2)$ would only be inspected after a bad first realization, but $p^{u*} = \mathbf{p^D}$ prevents this search altogether. When the seller selects a price that prevents search after a bad first realization, introducing three or four products is equivalent from the seller's perspective. Because this equivalence is a byproduct of the unrealistic assumption of zero fixed costs associated with each product, it is sensible to assume that, in this case, only three products would be introduced.

Notice that this does not affect the expected utility of search if inspection starts from the right product. If $(2, 2)$ were to be removed, search starting from $(1, 1)$ would be unaffected. Starting from any other product, however, would generate negative expected utility of search. Suppose for example that $(2, 2)$ was removed and that the consumer started from $(1, 2)$ or $(2, 1)$. Then, not only she would not rationally inspect the unrelated product, but she would not be able to inspect $(2, 2)$ after learning something positive about it. This cannot be optimal.

By removing a product, the seller effectively "locks" the consumer into a specific search path. The values α, s and p^u determine which search paths can be taken; given these search paths, products are introduced. For example: if it $s > \alpha$, inspection of a single attribute is never rational. Then, the only feasible search paths affect products that share no attributes. It follows that, in this case, only products that share no attribute would be introduced. The discussion motivates the following result:

Proposition 1. *Consider a multiproduct seller selecting optimal menu $\tilde{N} \subseteq N$ and uniform pricing p^u of multi-attribute products. In equilibrium:*

- *If $s > \alpha$: $p^{u*} = p^M$, $|\tilde{N}| = 2$, and the consumer can start searching from any available product.*
- *If $s < \alpha$ and $p^{u*} = \mathbf{p^E}$: $\tilde{N} \equiv N$, and the consumer can start searching from any available product.*
- *If $s < \alpha$ and $p^{u*} = \mathbf{p^D}$: $|\tilde{N}| = 3$, and the consumer is steered toward a specific search path.*

Proof. All calculations and precise cut-offs for α and s can be found in Appendix A. ■

Discussion. The seller values higher probabilities of trade taking place: because prices are uniform, the seller is not concerned with which product is purchased as long as one is. Selecting prices that do not hinder the probability of trade is often optimal. Raising prices is only worth it if the loss of a potential trade is compensated when trade does take place. In particular, α must be high enough that the chances of not liking the first product inspected are low, and s must be low enough that search is not discouraged. Whenever this is the case, the seller can raise price and not introduce all possible variants; as a consequence, there is a loss in trade efficiency. When the supply is restricted, moreover, the seller effectively induces a specific order of search. strategic menu selection can give rise to endogenous prominence based on the relative position of the products.

At uniform prices the consumer retains some positive expected value from search when the seller has an incentive to maximize trade by keeping prices low. Whenever this is the case, moreover, the consumer is free to start from any of the available products. As I will show in the next section, however, the seller generally has a profitable deviation if he is allowed to set different prices for these products and soften the trade-off between encouraging search after bad realizations and profiting whenever fine-tuning after a good, but not great, match is possible.

4.2. Differential Prices

When prices are assumed to be uniform, the choice of the seller is between keeping prices low to maximize search, and raising them to capitalize on good realizations. Ideally, the seller wants both: low prices to make the consumer keep searching after bad realizations, and high prices to profit off the consumer learning what she likes. This can be achieved if the seller can price products differently.

The trade-off of the seller under uniform prices refers to different search paths. Low prices encourage further search whenever the consumer finds nothing to like with her first inspection. High prices generate higher profits when the consumer partially likes at least the first option inspected. By pricing along these paths differently, the seller can achieve both higher probability of trade compared to the high uniform price case, and higher expected profit compared to the low uniform price case.

To see why, consider again the uniform price p_T that generates the maximum probability of trade but low rent extraction. When this price is optimally selected, it allows the consumer to keep searching after a bad first realization, and trade is likely to take place. In particular, what is needed is that the first product inspected, say $(1,1)$, and the product that would be searched next conditional on $A_1 = B_1 = 0$, $(2,2)$, to be priced at p_T. On this path, if the other products were priced higher than p_T, nothing would change because $(1,2)$ and $(2,1)$ would not be considered even at uniform prices, as long as the consumer can rationally start searching.

If the consumer, instead, learns that she likes an attribute inspected in the first search, she would like to search next along that attribute. This is clearly true if prices are uniform. Suppose, however, that $(1,2)$ and $(2,1)$ were priced slightly higher than $(1,1)$. If the consumer has learned that she likes A_1 (resp. B_1), and if the price difference is not too high, she would still want to search the more expensive product. Going backwards: the consumer would start her search from the cheaper option given that products are *ex ante* identical. As long as the price differential is not too high, the consumer has no incentives to stop searching early, nor to deviate towards a different search path. By pricing $(1,1)$ and $(2,2)$ at $p = p_T$, and the remaining products at a higher price the seller can then achieve both higher prices and higher probability of trade. In doing so, the seller erodes at the consumer expected utility without preventing search. When considering the equilibrium strategy of the seller, the following result emerges:

Proposition 2. *Consider a multiproduct seller selecting optimal menu $\tilde{N} \subseteq N$ and pricing $\mathbf{p}(\tilde{N})$ of multi-attribute products. There exist values $\underline{\alpha} \in (0,1)$ and $\underline{s} \in (0,\alpha)$ such that, in equilibrium:*

- *For $\alpha \in (0, \underline{\alpha}]$:*
  - *all products are introduced at different prices for $s \in (0,\alpha]$, and*
  - *two uncorrelated products are introduced and priced at $p = p^M$ for $s \in (\alpha, 2\alpha)$.*
- *For $\alpha \in (\underline{\alpha}, 1)$:*
  - *three products are introduced and priced at $p \in p^D$ for $s \in (0, \underline{s}]$,*
  - *all products are introduced at different prices for $s \in (\underline{s}, \alpha]$, and*
  - *two uncorrelated products are introduced and priced at $p = p^M$ for $s \in (\alpha, 2\alpha)$.*

Proof. All calculations and precise cut-off values for $\underline{\alpha}$ and $\underline{s}$ can be found in Appendix A. ■

Determining the optimal pricing vector with differential prices is challenging in this environment. In particular, the difference in prices can induce the consumer to adapt their search strategy to avoid the more expensive product and retain some expected utility. We are interested in finding out the optimal price spread from the seller's point of view, in which cases this spread does not affect the optimal search order, and, when it does, what is the seller optimal "reply". Henceforth, I assume that $(1,1)$ and $(2,2)$ have lower prices and therefore act as possible starting points; furthermore, I keep the assumption of products over which the consumer is indifferent to be searched in increasing order of their indices.

First, consider the optimal price spread. The search rules determine two separate constraints. Prices must be such that search can start. Moreover, prices must be consistent with the search process as it unfolds. The price increase being profitable relies on the consumer learning about which attribute she likes: a higher price can arise only on a path dictated by the consumer finding an attribute to keep. Suppose the consumer inspects $(1,1)$ and observes $A_1 = 1, B_1 = 0$. Suppose moreover that the optimal base price selected by the seller is $p_T \leq 1$. Conditional on inspecting one attribute being worth the cost of inspection ($s < \alpha$), the consumer would want to search $(1,2)$ if:

$$u_{1,1} = 1 - p_{1,1} \leq 1 + \alpha - s - p_{1,2} = E(u_{1,2}|I = \{(1,1)\}),$$

which implies $p_{1,2} \leq p_{1,1} + \alpha - s$, where $p_{1,1}$ and $p_{1,2}$ are the observed prices for $(1,1)$ and $(1,2)$, respectively. The higher price $p_{1,2}$ effectively captures the expected gain of searching that product after learning positive information about it by inspecting a different product. Because the seller is interested in the highest price that does not dissuade the search, the following candidate prices profile arises:

$$p_{1,1} = p_{2,2} = p^* = p_T \quad p_{1,2} = p_{1,2} = p^{**} = p_T + \alpha - s = p_T + \delta_L,$$

if $\alpha^2 < s < \alpha$, and:

$$p_{1,1} = p_{2,2} = p^* = p_H^E > 1 \quad p_{1,2} = p_{1,2} = p^{**} = 2 - \frac{s}{\alpha},$$

if $0 < s < \alpha^2$. The latter can be found following the same steps as the former, accounting for the fact that at these prices only a product that the consumer likes in both its attributes can be purchased.

Given search as characterized above, these pricing structure lead to the same probability of trade as their uniform counterparts. Compared to them, however, they lead to higher expected profit because the more expensive products are purchased with positive probability. Notice that this deviation preserves the internal consistency of the search process because the consumer would always inspect the cheapest product first if she has no information on any of the available products.

Consumer adaptation and firm response. Differential prices can distort the optimal search order of the consumer after the first realization. In particular, the consumer could find it optimal to ignore the more expensive product even if she learns that she likes something about it. In this case, the consumer would search $(2,2)$ hoping to find a good realization instead, and would only inspect the more expensive product if she knows she likes both of its attributes and nothing else. Consider again the candidate prices profile $p_{1,1} = p_{2,2} = p_T$, $p_{1,2} = p_{1,2} = p_T + \delta_L$. After realization $A_1 = 1$, $B_1 = 0$:

$$u_{1,1} = 1 - p_{1,1}, \quad E(u_{1,2}|I = \{(1,1)\}) = 1 + \alpha - s - p_{1,2},$$

$$E(u_{2,2}|I = \{(1,1)\}) = \alpha^2(2 - p_{2,2}) + (1-\alpha)(1 - p_{2,2}) + \alpha(1-\alpha)(2 - s - p_{1,2}) - s$$

When prices are uniform, a consumer would always want to inspect $(1,2)$ after learning $A_1 = 1$, $B_1 = 0$. This is not necessarily the case. It is possible that the consumer, observing the different prices, decides to change the order in which to inspect the remaining products. In particular, she could elect to inspect $(2,2)$ first and learn her realizations for all attributes. Then, the consumer could discover that $u_{2,2} = 2$, which she would not be able to by inspecting $(1,2)$. If she were to learn that $A_2 = 0$ and $B_2 = 1$, instead, then and only then would she inspect $(1,2)$ and purchase it.

For α high enough and s low enough, inspecting $(2,2)$ before $(1,2)$ is a rational deviation: search in this case is cheap, and the likelihood of liking both attributes A_2 and B_2 is relatively high. This deviation is at the detriment of the seller: the more expensive products now are reached with lower probability. The seller can optimally reply in three ways:

- the seller can let the consumer search $(2,2)$ first, and further increase $p_{1,2}$ and $p_{2,1}$ to $(p_{1,1} + 1 - s)$, or
- the seller can reduce prices $p_{1,2}$ and $p_{2,1}$ to encourage his preferred order of search to arise, or
- the seller can remove $(2,2)$ to induce his preferred order of search and keep the same prices for all other products.

The first reply further highlights the ability of the seller to condition prices on search behavior. If the consumer has an incentive to search $(2,2)$ after $(1,1)$ conditional on $A_1 + B_1 = 1$, the seller knows that the other two products would only be reached if they are the only product generating utility equal to 2. The probability of this happening, however, is lower than in the optimal price profile. Alternatively, the seller can make $(1,2)$ and $(2,1)$ cheaper. Because the consumer is interested in first searching $(2,2)$ because the alternative is too expensive, this deviation re-establishes the most profitable search order. Because the prices need to be lower, however, these paths are now less profitable than without the deviation. Finally, removing $(2,2)$ forces the consumer to take the path that the seller wants her to. This, however, reduces the probability of trade. These deviation are only necessary as long as $(\alpha, s) \in (0,1) \times (0, \alpha^2)$: when $s > \alpha^2$, search costs are too high for the consumer to be interested in searching $(2,2)$ when the seller would want her to inspect $(1,2)$ or $(2,1)$.

Each of the above strategies generates different expected profits for the seller. Given $p^* = p_T$, $p^{**} = p_T + \delta_L$:

- if the seller allows the consumer to deviate and raises p^{**} to $\overline{p} = p^* + 1 - s$,

$$\overline{\pi} = (1 - (1-a)^4)p^* + 2\alpha^2(1-\alpha)^2(\overline{p} - p^*);$$

- if the seller reduces p^{**} to $\underline{p}$ to induce seller preferred order,

$$\underline{\pi} = (1 - (1-a)^4)p^* + 2\alpha^2(1-\alpha)(\underline{p} - p^*);$$

- if the seller removes $(2,2)$ to prevent the deviation deviation,

$$\widehat{\pi} = (1 - (1-a)^2)p^* + 2\alpha^2(1-\alpha)(p^{**} - p^*).$$

All three options are optimal for some combinations of α and s. The same exercise can be applied to the alternative pricing profile $p_{1,1} = p_{2,2} = p_H^E > 1$, $p_{1,2} = p_{1,2} = 2 - \frac{s}{\alpha}$: in this case, deviation by the consumer is always feasible, and so the seller must react accordingly as well. In particular, for this alternative profile, removing $(2,2)$ always dominates the other two strategies.

Comparison. The feasible expected profits under differential prices must be compared to the highest expected profit under uniform prices obtained in the previous subsection. Two results emerge. First, whenever the seller has an incentive to select uniform prices that encourage search, he has an incentive to differentiate prices. This is intuitive: the lowest prices when products are priced differently are the same as the trade-maximizing uniform price. Because prices are set up to generate strictly higher profits while maintaining the same probability of trade, it is clearly an improvement to set differential prices.

Second, whenever it is optimal to remove a product to prevent deviation by the consumer, the high uniform prices generate higher profits. This, too, is straightforward: when the seller's best option is to give up on an inspection in case of a bad first match, uniform prices generate higher expected profits because, when prices are different, the consumer always starts from the cheaper one.

The leftmost graph in Figure 3 summarizes the equilibrium menu and pricing selection for all feasible combinations of α and s. The two decisions are intertwined. The seller has an incentive to make all products available only if they can all be reached, and purchased, with positive probability.

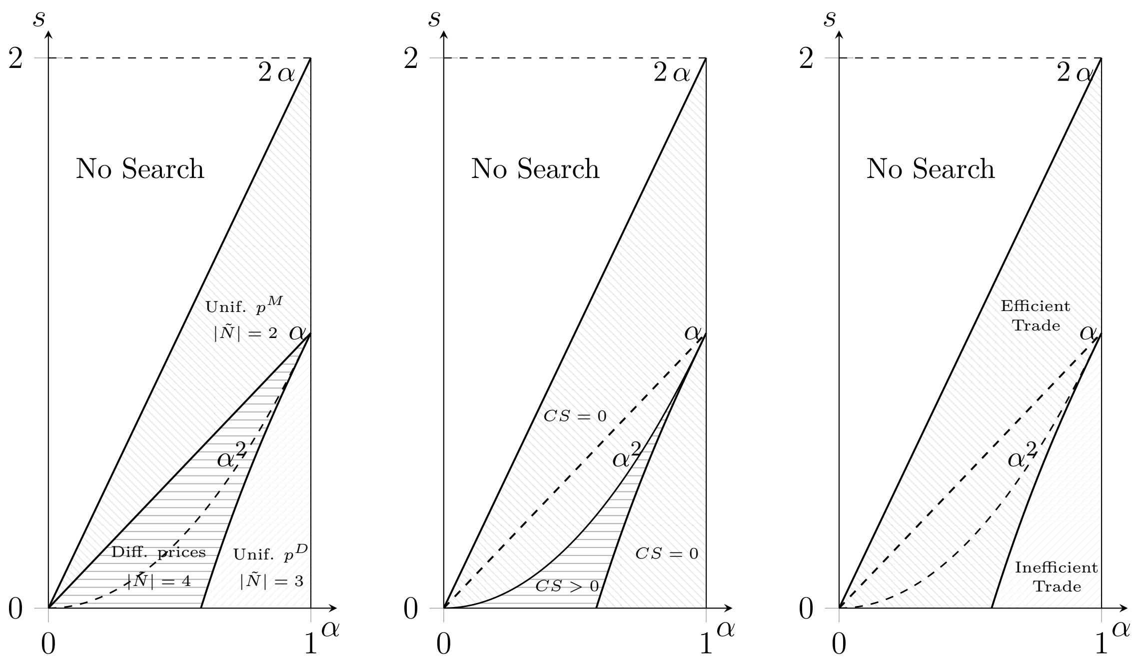

Figure 3: *Equilibrium monopoly menu selection and pricing (leftmost), expected consumer surplus (center), and trade efficiency (rightmost) for all feasible combinations of* $\alpha \in (0, 1)$ *and search costs* $s \in (0, 2\alpha)$

When search costs are very high, only a bad first realization induces the consumer to keep searching: introducing more than two products allows the consumer to randomize her starting point but at no benefit to the seller. On the other hand, when search costs are very low, the seller prefers to set prices that prevent some search paths to arise if probability of a match is relatively high. Lower search costs do not necessarily translate to more product variety, nor to efficient trade: when uniform high prices are selected, probability of trade is not as high as it could feasibly be because the menu is strategically restricted as well.

Finally, whenever all products are introduced, they are never priced uniformly. This pricing structure allows the monopolist to more efficiently extract rent while maintaining the highest probability of trade. Only when the consumer can deviate and force a reaction in the monopolist optimal pricing (that is, for $s < \alpha^2$), the consumer preserves some positive expected utility as long as the menu is not optimally restricted by the monopolist. Otherwise, the monopolist is able to capture it all through strategic pricing and menu selection.

4.3. Discussion of the Results

The results of this section highlight the incentives of a multiproduct seller to strategically determine the menu of available products to extract rent efficiently. To do so, he leads consumers towards specific search paths consistent with different outcomes of past inspections. With differential prices the seller is able to profit off the learning component of search in this environment.

This finding is at odds with the standard prediction of search models with multiproduct firms. In environments in which inspection of a product does not inform consumers of their taste for alternatives, strategic obfuscation of alternatives is the general outcome. Petrikaitė (2018), for example, shows that a multiproduct seller, like the one studied here, has an incentive to increase search cost of inspecting one product to induce consumers to inspect the easier-to-find, and more expensive, alternative first. Strikingly, the prediction goes in the opposite direction in the framework presented here. The learning component induces the seller to display some products more prominently, at a lower price, to let consumers learn about their tastes. Encouraging search, rather than discouraging it, allows the seller to sell more expensive products.

A possible application of this framework relates to the practice of businesses to offer free samples of new products to attract interest. In particular, by making some products prominent and easy to assess, a firm can encourage potential buyers to learn about their taste for novelties and alternatives that they might not have considered otherwise. In doing so, the firm can use the positive experience associated with the sampling to increase the willingness to pay of consumers unaware of their preferences for said products. Together with the strategic ordering shown above, this points at the importance of menu selection and positioning of options in environments with search frictions. It also creates a peculiar parallel with the prominence literature (Armstrong et al., 2009, Zhou, 2011): the products inspected early are cheaper not to retain consumers as in the aforementioned literature, but to induce additional inspections of more expensive alternatives.

The model carries implications for digital markets, particularly in relation to recommendation systems and price discrimination based on consumers' search history. Recommendation systems have been objects of great interest and scrutiny in the past few years because of their crucial role in the digital economy. A good recommendation system reduces frictions and, therefore, increases efficiency of trade. It is clear, however, that such systems can be objects of manipulation. The results of the model imply that the learning component relevant when searching products sharing attributes creates incentives to bias recommendations. The seller modelled here does not want to make prominent the best match possible. Rather, he wants the consumer to start from a subpar match and then self-select towards a more expensive product after learning her preferences because she might be discouraged from inspecting an expensive product without any information about it.

Consumers self-selecting based on taste also creates the incentive to condition pricing on their search history. Algorithmic pricing, the practice of pricing items automatically to adapt to the state of the market, are more and more commonly used in the digital world.[10] The model highlights the role that a product's position in the attribute space

[10] Airline companies and, more recently, e-commerce retailers are prime examples of this practice being

plays in their pricing. Consumers are willing to search more expensive products only if they have already learned something positive about them by inspecting a different option. Equivalently, one can imagine a reactive pricing system that adapts as search unfolds. If two products sharing attributes are inspected in sequence, the ordering signals that the consumer has learned something positive about those attributes. The price of the second, then, can safely be raised by an algorithm trained to recognize these patterns. On the other hand, if two products not sharing any feature are inspected in sequence, both should be priced low to maintain the consumer engaged with the search.

5. Extension - Mainstream and Niche attributes

Not all products in a given category are equal: some are generic and tend to have a broad appeal while others are more polarizing and tend to be either loved or hated. This distinction, and product positioning more in general, has been studied extensively. Recent work by Johnson and Myatt (2006) has provided a useful framework to study the incentives of firms to strategically select between different degrees of "niche" design; the framework has then been adapted to accommodate consumer search by Bar-Isaac et al. (2012). The main predictions on the topic can be summarized as follows: first, firms have an incentive to select a "polarized" design, either as generic or as controversial as possible. Second, a reduction is search costs can explain both a superstar effect (the most popular designs get more popular with falling search costs) and a long-tail effect (the least popular designs get more popular with falling search costs).

I now adapt the baseline model presented above in an application that focus on the relative likelihood of sampling mainstream and niche products. The resulting dynamics challenge some of the literature's results: a targeted search process reduces the incentives to chase the most niche design possible, while the learning process makes mixed designs profitable. Further, I show that while the long-tail effect is a first-order consequence of reduced search costs, the superstar effect fails to manifest in the absence of an increase in the number of consumers searching in the market.

Adapting the framework Suppose that attributes still follow a binomial distribution, but are not identically distributed. In particular, suppose that A_1 and B_1 generate utility V_1 with probability α_1, and that A_2 and B_2 generate normalized utility V_2 with probability α_2. To model niche and mainstream attributes, I impose:

$$\alpha_1 = \alpha \qquad \alpha_2 = (1-\varepsilon)\alpha \qquad \varepsilon \in (0,1)$$

$$V_1 = 1 \qquad V_2 = 1+\varepsilon \qquad \varepsilon \in (0,1)$$

widely in use.

So that attributes A_1 and B_1 have a higher expected value, $\alpha > (1-\varepsilon^2)\alpha$, but conditional on a match attributes A_2 and B_2 lead to a higher utility.[11]

Of the products defined by these attributes, then, $(1,1)$ can be considered generic or "mainstream", with a broad appeal and lower risk of dissatisfaction; $(2,2)$ can be consider provocative or "niche", with a lower expected appeal but a higher value for fans; $x_{1,2}$ and $x_{2,1}$ can be considered mixed. How does this additional differentiation affect optimal search order?

We can distinguish between different threshold levels of s that imply which kind of search can be rationally conducted:[12]

- 2α, above which search cannot take place,
- $(1+(1-\varepsilon^2))\alpha$, above which only $(1,1)$ can be rationally inspected,
- $2(1-\varepsilon^2)\alpha$, above which any fully unsampled product except for $(2,2)$ can be searched
- α, above which searching for an individual mainstream attribute leads to a negative expected gain from search
- $(1-\varepsilon^2)\alpha$, above which searching for an individual niche attribute leads to a negative expected gain from search

which are ordered differently depending on ε:

$$2(1-\varepsilon^2)\alpha > \alpha \qquad \Longleftrightarrow \qquad \varepsilon < \frac{\sqrt{2}}{2}$$

Depending on the relative values of the three variables above different search decisions can incur and thus the analysis must distinguish the various possible cases. Besides 2α and 0, the thresholds between which search takes place at all, the only thresholds the position of which is fixed in the order are the highest, $(1+(1-\varepsilon^2))\alpha$, and the lowest, $(1-\varepsilon^2)\alpha$. If $(1+(1-\varepsilon^2))\alpha < s \leq 2\alpha$, only $(1,1)$, the mainstream product, can be rationally searched and is therefore trivially the optimal starting point in this segment.

Suppose instead that $s \leq (1-\varepsilon^2)\alpha$ so that every product can be searched and every individual attribute is worth exploring after a partial match. In this case it is generally optimal to start from $(2,2)$, the niche product, but there are combinations of α and s such that $(1,1)$ is still the better starting point.

This can be shown by direct comparison of the expected utility associated with different

[11] For simplicity, and to highlight only the effect of these changes on the search environment, I ignore the pricing dimension in this application.

[12] Notice that, for $\varepsilon = 0$, the thresholds collapse to the two considered in the baseline framework.

starting points. Starting from $(1, 1)$ leads to:

$$\begin{aligned}E[u_{1,1}|I \equiv \emptyset] = 2\alpha^2 - s \\ &+ 2\alpha(1-\alpha)[(1-\varepsilon)\alpha(1+(1+\varepsilon)) + (1-(1-\varepsilon)\alpha) - s] \\ &+ (1-\alpha)^2[2(1-\varepsilon)^2\alpha^2(1+\varepsilon) + 2(1-\varepsilon)\alpha(1-(1-\varepsilon)\alpha)(1+\varepsilon) - s]\end{aligned}$$

Since sampling again after $(1, 1)$ must target one or two niche attributes. Starting from $(2, 2)$ and $(1, 2)$ (or $(2, 1)$), instead, has associated expected utility:

$$\begin{aligned}E[u_{2,2}|I \equiv \emptyset] = 2(1-\varepsilon)^2\alpha^2(1+\varepsilon) - s \\ &+ 2(1-\varepsilon)\alpha(1-(1-\varepsilon)\alpha)[\alpha(1+(1+\varepsilon)) + (1-\alpha)(1+\varepsilon) - s] \\ &+ (1-\varepsilon)^2(1-\alpha)^2[2\alpha^2 + 2\alpha(1-\alpha) - s]\end{aligned}$$

and:

$$\begin{aligned}E[u_{1,2}|I \equiv \emptyset] = \alpha(1-\varepsilon)\alpha(1+(1+\varepsilon)) - s \\ &+ \alpha(1-(1-\varepsilon)\alpha)[2\alpha + (1-\alpha) - s] \\ &+ (1-\varepsilon)\alpha(1-\alpha)[2(1-\varepsilon)\alpha(1+\varepsilon) + (1-(1-\varepsilon)\alpha)(1+\varepsilon) - s] \\ &+ (1-\alpha)(1-(1-\varepsilon)\alpha)[\alpha(1-\varepsilon)\alpha(1+(1+\varepsilon)) \\ &\quad + \alpha(1-(1-\varepsilon)\alpha) + (1-\varepsilon)\alpha(1-\alpha)(1+\varepsilon) - s]\end{aligned}$$

The direct comparison, fully contained in Appendix B, leads to the following observations:

- $(2, 2)$ is a strictly better starting point than $(1, 2)$, $\forall s, \alpha, \varepsilon$,
- $(1, 1)$ is strictly better than $(2, 2)$ for $s \in \left(\frac{2(1-\varepsilon)}{2-\varepsilon}, (1-\varepsilon^2)\alpha\right)$.

Intuitively, when search costs are relatively low it is generally worth inspecting first the riskier option, $(2, 2)$, because adjustment is always possible. Only for relatively high values of both α, which makes the mainstream option almost sure to be appreciated, and s it would be better for the buyer to try to avoid to search twice. Notably, $(1, 2)$ and $(2, 1)$ are never optimal starting points, but their presence affects the value of inspecting the others, and benefit the niche product relatively more since they imply a correction towards a "safer" attribute if either $A_2 = 0$ or $B_2 = 0$.

For intermediate segments of search costs, $(1-\varepsilon^2)\alpha < s \leq (2-\varepsilon^2)\alpha$, the magnitude of ε determines the relative position of the remaining thresholds. For ε low enough ($\varepsilon < \frac{\sqrt{2}}{2}$), $2(1-\varepsilon^2) > \alpha$, so that for some parameters it is feasible to search $(2, 2)$ but not follow up on any partial match. *Vice versa*, for ε high enough it becomes viable to sample an undiscovered mainstream attribute but not $(2, 2)$.

Despite this multiplicity, the comparison of expected utilities generated starting from any product worth searching leads to clear-cut results:

Proposition 3. *There exist a unique value $\tilde{s}(\alpha, \varepsilon)$ such that, $\forall(\alpha, \varepsilon) \in (0,1) \times (0,1)$:*

- *It is optimal for the consumer to start searching from niche product $(2,2)$ if $s \leq \tilde{s}(\alpha, \varepsilon)$,*
- *It is optimal for the consumer to start searching from mainstream product $(1,1)$ if $s > \tilde{s}(\alpha, \varepsilon)$,*
- *it holds:*

$$\frac{\partial \tilde{s}(\alpha, \varepsilon)}{\partial \alpha} \geq 0, \quad \frac{\partial \tilde{s}(\alpha, \varepsilon)}{\partial \varepsilon} < 0.$$

Proof. All calculations can be found in Appendix B. ■

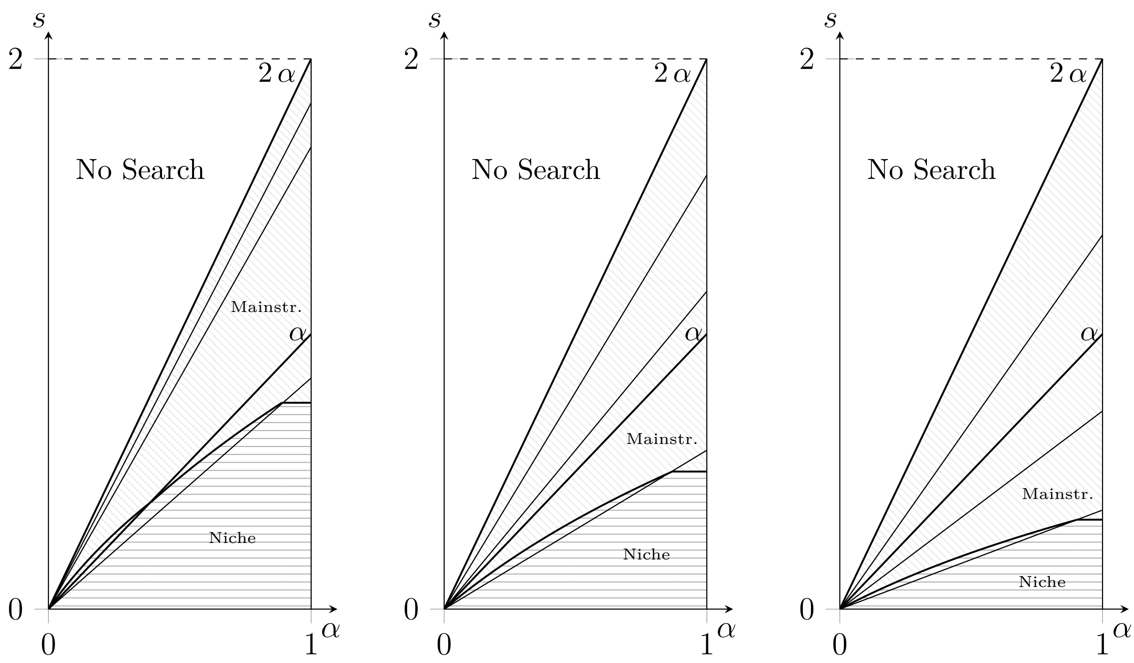

Figure 4: *Optimal starting point in the complete two-by-two grid for $\varepsilon = 0.4, 0.65, 0.8$ respectively; The moving thresholds, from highest to lowest, are: $s = (2-\varepsilon^2)\alpha$, $s = 2(1-\varepsilon^2)\alpha$, and $s = (1-\varepsilon^2)\alpha$.*

The niche product is associated with lower expected utility and, therefore, it is only viable to inspect it for relatively low search costs. Sampling the mainstream product, however, implies a lower chance to follow up on a partial match. Only for very low s starting from $(2,2)$ can ever be a rational strategy. As shown in Figure 4, which summarizes the results for different values of ε from lower (leftmost) to higher (rightmost) differentiation, this notion is reinforced by the fact that the higher ε is, the lower $\tilde{s}(\alpha, \varepsilon)$ is as well.

5.1. Discussion

Several observations follow from this exercise. First and foremost, it brings into question the incentives of firms to select polarized designs for their products. Johnson and Myatt (2006) and the rotation of demand studied therein states firm profits to be "U-shaped" in the design choice so that firms will optimally selects either as bland and generic or as controversial designs as possible. Bar-Isaac et al. (2012) confirms this result after developing a model of random search. The targeted nature of the process considered here, however, complicates the issue: niche products might be more profitable, generating a higher value conditional on a good match, but are also riskier and, as such, consumers sample them consistently only when it is very cheap to do so.

It can be readily shown that introducing a product located "in-between" $(1,1)$ and $(2,2)$ in terms of riskiness can be optimal: consider again Figure 4 and suppose $s = \tilde{s}(\alpha, \varepsilon)$ for $\varepsilon = 0.8$ (or, anywhere on the thick line in the rightmost figure). Then, if attributes A_2, B_2 were associated with $\varepsilon = 0.6$ instead, they would become the prominent one in the search process since, in this scenario, it is a strictly better starting point than $(1,1)$ (as shown in the central figure). Since prominent products have a higher probability of being purchased, and since niche attributes generate more value conditional on search, it follows that these "intermediate" attributes would be more profitable to introduce than more extreme ones.

The notion can also be challenged considering the mixed design products $(1,2)$ and $(2,1)$, both of which are clearly purchased with positive probability if search costs are low enough to allow consumers to make use of the information learned with the first sampling. Their presence also indirectly makes niche products like $(2,2)$ more appealing since it is always easier to follow up on a partial match sampling a mainstream attribute (with expected value α) than a niche one (with expected value $(1-\varepsilon^2)\alpha$). This observation is relevant for the same multiproduct firm selecting a menu to offer studied in the baseline search model: by introducing products with different combinations of niche and mainstream attributes and guiding search, such a firm can induce buyers to find the match for which they have the highest willingness to pay. Consumer learning and the low search costs associated with the digital economy could also make such a positioning feasible for a single-product firm trying to enter an established market.

The model also speak to the long-tail and superstar effects mentioned in the beginning of this section. With search costs decreasing over the last years thanks to the ever-growing prevalence of digital markets, both niche and already well-known products have seen a rise in popularity. The model provides an intuitive explanation of the first phenomenon and hints at the main drivers of the second.

It is clear that niche products have a more prominent role the lower the search costs are: when sampling is cheap, buyers prefer to sample first products associated with stronger but less likely appeal and allow for ever more daring designs to be introduced. Consider the space defined by $0 < s \leq 2(1-\varepsilon^2)\alpha$: both mainstream and niche designs can be rationally sampled at an expected gain. For relatively high values of s, however, sampling a niche product is conditional on a safe one resulting in a poor match; as s decreases, however, the risk associated with the more daring design becomes less and less relevant to the point that for low enough search costs, niche designs acquire prominence and are searched first by buyers and are, therefore, more likely to be purchased. It follows that endogenous search order can explain the reported long-tail effect arising as search costs drop.

On the other hand, a rise in prominence in niche designs has a negative effects on the chance of mainstream products to be sampled and purchased. The superstar effect cannot be directly explained by the changes introduced by targeted search and learning through shopping. Notice however that, as search costs drop, the expected utility of any search process increases: the lower s is, the more likely it is for a buyer to find a good match since more and more options become rationally available for sampling. The model presented in this paper assumes a representative consumer who starts searching with probability 1 and, therefore, assumes away any extensive margin that would manifest if heterogeneous consumers had to choose to start searching at the beginning of the game. The wider variety of products that lower search costs make profitable to introduce, and the consequent increase in the expected utility generated through search, however, would encourages more consumers to start searching. This effect on the extensive margin of a lower search costs, then, can explain the rise in popularity of modern superstar products.

6. Conclusion

In this paper, I study the implications of product correlation through shared attributes for directed search and the associated incentives of a seller to introduce different products and prices to capitalize on consumer learning. The framework highlights a novel interaction between pricing and optimal order of inspection in directed search: consumers have an incentive to find better matches in their search process as they learn what they like. This dictates their strategy predictably in a way that highlights how a multiproduct seller is able to profit off the learning process by setting differential prices to let consumers self-select based on their preferences.

The framework's predicted search patterns align well with recent evidence of spatial learning in search: Hodgson and Lewis (2020) reports evidence of search for digital cameras to be characterized by a learning process consistent with the one in this framework. Consumers are shown to inspect a broader set of attributes early only to close in on their preferred alternatives in later stages, getting closer and closer to the product they ultimately choose to purchase. This pattern cannot be easily reconciled with standard search models, but is well in line with the prediction of this framework. Further, the model presented here can more easily rationalize the pervasive tendency of consumers to retrace their steps while searching for products.

The implications of this model for recommendation systems and algorithmic pricing schemes have been addressed in an earlier section. It is worth stressing out, however, that these implications go beyond the specific market structure studied here. Coordination of menu and pricing allows a multiproduct seller to induce specific search paths to arise. Equivalently, one can imagine e-commerce platforms to do the same through manipulation of the options presented to captured consumers and the information therein. This is especially true in a world in which data on consumers' decisions, consumption and search patterns is abundant, and algorithmic pricing and recommendation systems are ever more effective at predicting human behavior. In line with recent work on consumption steering and self-preferencing, then, the model's results suggest the need for meticulous regulatory oversight over the algorithms determining what consumers shopping online are shown, and when.

References

Anderson, Simon P. and Régis Renault, "Advertising Content," *American Economic Review*, 2006, *96* (1), 93–113.

Anderson, Simon P., Maxim Engers, and Daniel Savelle, "Ordered Search: Equilibrium and Optimum," *Unpublished manuscript*, 2020.

Anderson, Simon P., Maxim Engers, and Daniel Savelle, "Consumer Journeys with Nested Weitzman Scores and Informational Complementarities," 2021. Unpublished manuscript.

Armstrong, Mark, John Vickers, and Jidong Zhou, "Prominence and Consumer Search," *RAND Journal of Economics*, 2009, *40* (2), 209–233.

Bao, Renjie, Sanxi Li, and Jun Yu, "Search with Interdependent Values," *Unpublished manuscript*, 2022.

Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat, "Search, Design, and Market Structure," *American Economic Review*, 2012, *102* (2), 1140–60.

Brander, James A. and Jonathan Eaton, "Product Line Rivalry," *American Economic Review*, 1984, *74* (3), 323–334.

Choi, Michael, Anovia Yifan Dai, and Kyungmin Kim, "Consumer Search and Price Competition," *Econometrica*, 2018, *86* (4), 1257–1281.

De Corniere, Alexandre and Greg Taylor, "A Model of Biased Intermediation," *RAND Journal of Economics*, 2019, *50* (4), 854–882.

Doval, Laura, "Whether or Not to Open Pandora's Box," *Journal of Economic Theory*, 2018, *175*, 127–158.

Eaton, B. Curtis and Richard G. Lipsey, "The Theory of Market Pre-Emption: the Persistence of Excess Capacity and Monopoly in Growing Spatial Markets," *Economica*, 1979, *46* (182), 149–158.

Garcia, Daniel and Sandro Shelegia, "Consumer Search with Observational Learning," *RAND Journal of Economics*, 2018, *49* (1), 224–253.

Greminger, Rafael P., "Optimal Search and Discovery," *Management Science*, 2022, *68* (5), 3904–3924.

Haan, Marco A., José L. Moraga-González, and Vaiva Petrikaitė, "A Model of Directed Consumer Search," *International Journal of Industrial Organization*, 2018, *61*, 223–255.

Heidhues, Paul, Mats Köster, and Botond Kőszegi, "Steering Fallible Consumers," *Economic Journal*, 2023, *133* (652), 1430–1465.

Hodgson, Charles and Gregory Lewis, "You Can Lead a Horse to Water: Spatial Learning and Path Dependence in Consumer Search," *Unpublished manuscript*, 2020.

Ichihashi, Shota, "Online Privacy and Information Disclosure by Consumers," *American Economic Review*, 2020, *110* (2), 569–595.

Johnson, Justin P. and David P. Myatt, "Multiproduct Quality Competition: Fighting Brands and Product Line Pruning.," *Unpublished manuscript*, 2002.

Johnson, Justin P. and David P. Myatt, "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 2006, *96* (3), 756–784.

Ke, T. Tony and Song Lin, "Informational Complementarity," *Management Science*, 2022, *66* (8), 3699–3716.

Konishi, Hideo and Michael T. Sandfort, "Expanding Demand Through Price Advertisement," *International Journal of Industrial Organization*, 2002, *20* (7), 965–994.

Lambertini, Luca and Andrea Mantovani, "Process and Product Innovation by a Multiproduct Monopolist: a Dynamic Approach," *International Journal of Industrial Organization*, 2009, *27* (4), 508–518.

Lancaster, Kelvin J., "A New Approach to Consumer Theory," *Journal of Political Economy*, 1966, *74* (2), 132–157.

Lin, Ping, "Process and Product R&D by a Multiproduct Monopolist," *Oxford Economic Papers*, 2004, *56* (4), 735–743.

Mayzlin, Dina and Jiwoong Shin, "Uninformative Advertising as an Invitation to Search," *Marketing science*, 2011, *30* (4), 666–685.

McAfee, R. Preston, John McMillan, and Michael D. Whinston, "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values," *Quarterly Journal of Economics*, 1989, *104* (2), 371–383.

Mussa, Michael and Sherwin Rosen, "Monopoly and Product Quality," *Journal of Economic theory*, 1978, *18* (2), 301–317.

Nocke, Volker and Nicolas Schutz, "Multiproduct-Firm Oligopoly: An Aggregative Games Approach," *Econometrica*, 2018, *86* (2), 523–557.

Nocke, Volker and Patrick Rey, "Consumer Search and Choice Overload," *Forthcoming (Journal of Political Economy)*, 2023.

Petrikaitė, Vaiva, "Consumer Obfuscation by a Multiproduct Firm," *RAND Journal of Economics*, 2018, *49* (1), 206–223.

Preuss, Marcel, "Search, Learning and Tracking," *RAND Journal of Economics*, 2023, *54* (1), 54–82.

Rhodes, Andrew, "Multiproduct Retailing," *Review of Economic Studies*, 2015, *82* (1), 360–390.

Rhodes, Andrew, Makoto Watanabe, and Jidong Zhou, "Multiproduct Intermediaries," *Journal of Political Economy*, 2021, *129* (2), 421–464.

Robert, Jacques and Dale O. Stahl, "Informative Price Advertising in a Sequential Search Model," *Econometrica*, 1993, pp. 657–686.

Teh, Tat-How and Julian Wright, "Intermediation and Steering: Competition in Prices and Commissions," *American Economic Journal: Microeconomics*, 2022, *14* (2), 281–321.

Weitzman, Martin L., "Optimal Search for the Best Alternative," *Econometrica*, 1979, *47* (3), 641–654.

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

Zhou, Jidong, "Ordered Search in Differentiated Markets," *International Journal of Industrial Organization*, 2011, *29* (2), 253–262.

Zhou, Jidong, "Multiproduct Search and the Joint Search Effect," *American Economic Review*, 2014, *104* (9), 2918–39.

Appendix

A. Simplified framework: monopoly pricing

Uniform prices As in the main text, I start by assuming $\tilde{N} \equiv N$ and obtain equilibrium prices for different combinations of α, s. Then, I show the optimal restriction of $\tilde{N}$ conditional on the optimal prices.

The seller is interested in finding prices that maximize probability of trade times price. Given expected utility of search as per Equation 1:

$$\begin{aligned}
&E[u_{1,1}]|_{I\equiv\emptyset} = \alpha^2 \max(2-p^u,0) - s \\
&+ 2\alpha(1-\alpha)\max(1-p^u,(1-\alpha)\max(1-p^u,0) + \alpha\max(2-p^u,0) - s, 0) \\
&+ (1-\alpha)^2 \max(\alpha^2 \max(2-p^u,0) + 2\alpha(1-\alpha)\max(1-p^u,0) - s, 0)
\end{aligned}$$

the highest prices that make consumers start search can be computed as prices that make the expression reach a value of zero:

$$\mathbf{p^D} = \begin{cases} p^M = \frac{2\alpha-s}{\alpha(2-\alpha)} & \text{if } \alpha \leq s < 2\alpha \\ p_L^D = \frac{2\alpha(1+(1-\alpha)(\alpha-s))-s}{\alpha(2-\alpha)} & \text{if } \frac{3\alpha^2-2\alpha^3}{1+2\alpha-2\alpha^2} \leq s < \alpha \\ p_H^D = \frac{2\alpha(\alpha(3-2\alpha)-(1-\alpha)s)-s}{\alpha^2(3-2\alpha)} & \text{if } 0 < s < \frac{3\alpha^2-2\alpha^3}{1+2\alpha-2\alpha^2} \end{cases}$$

The highest prices that allows for inspection after a bad first realization, instead, are:

$$\mathbf{p^E} = \begin{cases} p_L^E = \frac{2\alpha-s}{\alpha(2-\alpha)} & \text{if } \alpha^2 \leq s < 2\alpha \\ p_H^E = \frac{2\alpha^2-s}{\alpha^2} & \text{if } 0 < s < \alpha^2 \end{cases}$$

In each segment identified among the two sets of prices above, lower prices are always feasible, as they generate positive expected utility of search. Lower prices can induce more extensive search and higher probability of trade. Therefore, we look for profitable price reductions for each segment in consideration.

If $\alpha \leq s < 2\alpha$, only $p^M = p_L^E$ is feasible among the candidates above. Furthermore, it can be shown that:

$$\alpha \leq s < 2\alpha \;\rightarrow\; p^M < 1$$

By plugging in p^M in equation 1, one sees that at this prices the consumer stops and purchase if $u(a,b) \neq 0$, and is willing to search again if $u_{1,1} = 0$. It is clear that no deviation from p^M can be profitable: if prices are any higher, expected utility of search would be negative and search would not start; if prices were any lower, no additional probability of trade would be generated. Therefore, in this segment, $p^{u*} = p^M$.

If $\frac{3\alpha^2-2\alpha^3}{1+2\alpha-2\alpha^2} \leq s < \alpha$, both p_L^D and p_L^E are feasible. Moreover, it holds $p_M = p_L^E < p_L^D$

for the whole segment. Therefore, it is sufficient to compare expected profits under p_L^E and p_L^D. Notice that p_L^D is such that searching again after a bad first realization is not possible. In this segment:

$$\alpha^2(2 - p_L^D) + 2\alpha(1-\alpha)(1-p_L^D) - s < 0$$

Therefore, the seller compares:

$$\pi_L^E = (1-(1-\alpha)^4)p_L^E$$

$$\pi_I^D = \alpha^2(1+2(1-\alpha))p_L^D$$

Direct comparison indicates that p_L^D is selected for some combination of high α and relatively low s:

$$\pi_I^D > \pi_L^E \iff \frac{4\alpha^2 - 2\alpha}{3\alpha - 1} < s < \alpha$$

p_L^E is selected otherwise.

If $0 < s < \frac{3\alpha^2 - 2\alpha^3}{1+2\alpha-2\alpha^2}$, several distinctions must be made. First, $p_L^E < 1 \iff \alpha^2 < s < \alpha$. Therefore, for $s < \alpha^2$, $p_T = 1$ becomes a feasible deviation as it is the price that maximizes probability of trade. Further, p_H^D is now a feasible price to select: it only leads to a purchase if an inspected product is liked in both attributes, and allow for a second search after finding one liked attribute but not after a bad first realization. p_H^E also requires two attributes to be liked by the consumer, but always allow for a follow up search. p_H^E, which is always true in this segment, only allows for a follow-up search if $0 < s < \alpha^2$. This final segment must be split in two sub-segments.

If $\alpha^2 < s < \frac{3\alpha^2 - 2\alpha^3}{1+2\alpha-2\alpha^2}$, $p_L^E < 1$ is always the best choice:

$$\pi_L^E > \pi_H^D = p_H^D(\alpha^2(1+2(1-\alpha))$$

If $0 < s < \alpha^2$, $p_H^E > p_T$; the choice is between:

$$\pi_T = (1-(1-\alpha)^4)p_T$$

$$\pi_H^E = (\alpha^2(1+2(1-\alpha)+(1-\alpha)^2)p_H^E$$

$$\pi_H^D = (\alpha^2(1+2(1-\alpha))p_H^D$$

Direct comparison indicates that all three pricing levels can be optimal: π_T is optimal for:

$$\min\left(\frac{3\alpha^3 - 6\alpha^2 + 2\alpha}{\alpha - 2}, \frac{-\alpha^4 + 8\alpha^3 - 12\alpha^2 + 4\alpha}{2\alpha^2 - 2\alpha - 1}\right) < s < \alpha^2$$

π_H^E is optimal for:

$$0 < s < \min\left(\frac{3\alpha^3 - 6\alpha^2 + 2\alpha}{\alpha - 2}, \frac{2\alpha^2}{3}\right)$$

and α high enough. Otherwise, π_H^D is optimal.

All feasible combinations of $\alpha \in (0,1)$ and $s \in (0, 2\alpha)$ are then accounted for when restricting the seller to a uniform pricing strategy.

Differential prices It must be shown that the price deviations shown in the main text lead to a higher expected profit. Consider $p^u = p_L^E$. As long as at this price level consumers have a strictly positive expected utility of search, the seller can introduce differential prices profitably. In particular, consider pricing such that:

$$p_{1,1} = p_L^E < 1 \quad p_{2,2} = p_L^E < 1 \quad p_{1,2} = p_L^E + \alpha - s \quad p_{2,1} = p_L^E + \alpha - s$$

Which is valid for $p_L^E < 1$ or, $\alpha^2 < s$. As shown in the main text, for $s > \alpha$ the consumer has no reason to search again after finding something she likes, and indeed would lead to a lower, rather than higher, price level for $p_{1,2}$ and $p_{2,1}$. In this segment ($\alpha^2 < s < \alpha$), such prices lead to strictly higher expected profits. Indeed, when the consumer starts from $(1,1)$ (equivalently, $(2,2)$), she only searches the more expensive product if she already knows that she likes it in some attribute. The consumer cannot start from any other product: if she starts from the more expensive product, her expected utility of search in this segment is negative.

Finally, the difference in prices do not induce changes in the optimal search path. To see why, consider the optimal deviation available to the consumer on the path in which she would want to inspect $(1,2)$: inspecting $(2,2)$ leads to utility equal to two with probability α^2, and allows to correct to $(1,2)$ if she learns that she likes B_2 but not A_2, which happens with probability $\alpha(1-\alpha)$. The expected utility along this alternate path is equal to:

$$(\alpha^2(2 - p_L^E) + (\alpha(1-\alpha) + (1-\alpha)^2)(1 - p_L^E) + \alpha(1-\alpha)(2 - s - (p_L^E + \alpha - s)) - s$$

which is lower than the expected utility of searching $(1,2)$ directly if $s > \alpha^2$. Therefore, no deviation is possible in this segment.

If $s < \alpha^2$, two changes must be accounted for. First, p_T is the preferred option, because $p_L^E > 1$ does not lead to trade taking place. In turns, this implies that because base prices are lower than the myopic expected value of inspecting a product, consumer surplus is above zero if $s < \alpha^2$ under differentiated prices. Further, the consumer would want to search the cheaper $(2,2)$ first, because search costs are low. The seller can react by:

- letting the consumer do so, increase the price of $(1,2)$ to $p_T + 1 - s$
- reducing the price $(1,2)$ to induce his preferred order of search
- removing $(2,2)$.

The first reaction re-establishes the equilibrium: the consumer now inspects the more expensive product only if he knows it is the only product that leads to utility equal to two. Because this is the case, its price can be increased, because the search process took away all uncertainty about it. This product is purchased with probability $\alpha^2(1-\alpha)^2$ and leads to expected profit:

$$\overline{\pi} = (1-(1-\alpha)^4)p_T + 2\alpha^2(1-\alpha)^2(1-s)$$

which is still a strictly higher expected profit than the respective uniform price strategy.

The second reaction also re-establishes the equilibrium: by setting a lower price for $(1,2)$, the seller makes sure that the consumer has no incentive to deviate. Because $s<\alpha^2$, the baseline price is $p=p_T$ and the level p that prevents the deviation solves:

$$\alpha(2-p)-s = \alpha^2 + (1-\alpha)\alpha(-p-s+2)-s \iff p = 1+s\left(\frac{1-\alpha}{\alpha}\right)$$

which leads to expected profits:

$$\underline{\pi} = (1-(1-\alpha)^4)p_T + 2\alpha^2(1-\alpha)\left(s\left(\frac{1-\alpha}{\alpha}\right)\right)$$

Finally, removing $(2,2)$ prevents the deviation from taking place at all. Because no follow-up search in case of a bad first realization is possible without $(2,2)$, however, overall probability of trade decreases. Expected profits in this case are:

$$\widehat{\pi} = (1-(1-\alpha)^2)p_T + 2\alpha^2(1-\alpha)(\alpha-s)$$

By direct comparison, one finds that all three can be optimal for different values of α, s. In particular, $\widehat{\pi}$ is optimal for α high enough, that is, for:

$$0 < s < \min\left(\frac{3\alpha^2+\alpha-2}{2\alpha^2}, \frac{1}{2}\left(\alpha^2+3\alpha-2\right)\right)$$

$\underline{\pi}$ is optimal for:

$$\max\left(\frac{\alpha}{\alpha+1}, \frac{1}{2}\left(\alpha^2+3\alpha-2\right)\right) < s < \alpha^2$$

while $\overline{\pi}$ is optimal otherwise.

The same argument can be applied to the trade-off between p_H^E and p_H^D when $0<s<\alpha^2$. In this segment, p_H^E is such that trade only happens if the consumer learns that she likes both attributes about a product, but the parameters encourage the consumer to search again after a bad first realization. Here, too, the seller can choose an intermediate strategy between uniform prices at p_H^E and uniform prices at p_H^D. Suppose the consumer inspected $(1,1)$ and learned $A_1=1$, $B_1=0$. Then, she would want to inspect $(1,2)$. She does so as long as:

$$\alpha(2-p_{1,2})-s \geq 0 > 1-p_H^E$$

which implies:

$$p_{1,2} = 2 - \frac{s}{\alpha}$$

It can be shown that the consumer always reacts to this price level by inspecting $(2,2)$ instead of $(1,2)$. Indeed, if $0 < s < \frac{\alpha^2}{1+\alpha}$, it holds:

$$\alpha^2(2-p_I) + (1-\alpha)\alpha\left(-\left(2-\frac{s}{\alpha}\right) - s + 2\right) - s > \alpha\left(2 - \left(2 - \frac{s}{\alpha}\right)\right) - s = 0$$

Once again, the seller can react by allowing the deviation and further increasing $p_{1,2}$ to $2-s$, reducing $p_{1,2}$ to $\frac{2\alpha^3-2\alpha^2-2\alpha-3\alpha^2 s+5\alpha s-s}{(\alpha-2)\alpha}$ to make the consumer search according to his preferred order, or remove $(2,2)$.

Unlike in the previous case, the latter option is always optimal. When the sellers selects differentiated prices, then, for α high and s low the consumer has an incentive to adapt in a way that makes the seller restrict the menu of available products.

Comparison Comparison between the optimal uniform price strategy and the deviation shown above is straightforward. First, it is trivial that whenever $p^{u*} = p_T$, all deviations are strictly preferable: indeed, the strategy with differentiated prices preserves the total probability of trade but generates higher profits for some positive probability. To compare the above strategy with the other uniform prices the seller can optimally select, direct comparison of the profit is sufficient. The same applies to the case in which $p^{u*} = p_L^D$ and $0 < s < \alpha^2$.

Two results emerge: when selecting p_T as base product and the consumer does not adapt their search strategy, this is always optimal. Second, when there is adaptation by consumer and seller, those profits must be compared with the relevant uniform price in the segment, that is, p_H^D.

Direct comparison indicates that p_H^D dominates different prices whenever the optimal reply of the seller to the consumer adapting his search strategy is to restrict the supply. This follows from the fact that, with different prices, consumers always search the cheapest one first. Therefore, the only comparisons left are between π_H^D and the best between $\overline{\pi}$ and $\underline{\pi}$ when $p^* = p_T$. It holds:

$$\underline{\pi} > \pi_H^D \iff \frac{\alpha^4 - 8\alpha^3 + 12\alpha^2 - 4\alpha}{2\alpha^3 - 6\alpha^2 + 4\alpha + 1} < s < \alpha^2$$

$$\overline{\pi} > \pi_H^D \iff \frac{\alpha^4 + 4\alpha^3 - 10\alpha^2 + 4\alpha}{2\alpha^4 - 4\alpha^3 + 4\alpha^2 - 2\alpha - 1} < s < \alpha^2$$

Which delimit the lower right area in Figure 3 in the main text.

B. Extension - mainstream and niche attributes

Suppose that $s \leq (1-\varepsilon^2)\alpha$ so that every product can be searched and every individual attribute is worth exploring after a partial match.

Expected utility associated with starting from $(1,1)$ is:

$$\begin{aligned}E[u_{1,1}|I \equiv \emptyset] = 2\alpha^2 - s \\ &+ 2\alpha(1-\alpha)[(1-\varepsilon)\alpha(1+(1+\varepsilon)) + (1-(1-\varepsilon)\alpha) - s] \\ &+ (1-\alpha)^2[2(1-\varepsilon)^2\alpha^2(1+\varepsilon) + 2(1-\varepsilon)\alpha(1-(1-\varepsilon)\alpha)(1+\varepsilon) - s]\end{aligned}$$

Expected utility associated with starting from $(2,2)$ is:

$$\begin{aligned}E[u_{2,2}|I \equiv \emptyset] = 2(1-\varepsilon)^2\alpha^2(1+\varepsilon) - s \\ &+ 2(1-\varepsilon)\alpha(1-(1-\varepsilon)\alpha)[\alpha(1+(1+\varepsilon)) + (1-\alpha)(1+\varepsilon) - s] \\ &+ (1-\varepsilon)^2(1-\alpha)^2[2\alpha^2 + 2\alpha(1-\alpha) - s]\end{aligned}$$

Expected utility associated with starting from $(1,2)$(or $(2,1)$) is:

$$\begin{aligned}E[u_{1,2}|I \equiv \emptyset] = \alpha(1-\varepsilon)\alpha(1+(1+\varepsilon)) - s \\ &+ \alpha(1-(1-\varepsilon)\alpha)[2\alpha + (1-\alpha) - s] \\ &+ (1-\varepsilon)\alpha(1-\alpha)[2(1-\varepsilon)\alpha(1+\varepsilon) + (1-(1-\varepsilon)\alpha)(1+\varepsilon) - s] \\ &+ (1-\alpha)(1-(1-\varepsilon)\alpha)[\alpha(1-\varepsilon)\alpha(1+(1+\varepsilon)) \\ &\quad + \alpha(1-(1-\varepsilon)\alpha) + (1-\varepsilon)\alpha(1-\alpha)(1+\varepsilon) - s]\end{aligned}$$

The three can be directly compared:

- $E[u_{2,2}|I \equiv \emptyset] > E[u_{1,2}|I \equiv \emptyset] \iff s < 1$, which is larger than the threshold delimiting the segment: $(1-\varepsilon^2)\alpha$,
- $E[u_{2,2}|I \equiv \emptyset] > E[u_{1,2}|I \equiv \emptyset] \iff s < \frac{2(1-\varepsilon)}{2-\varepsilon}$, which is smaller than $(1-\varepsilon^2)\alpha$ for α large enough.

The order of the remaining thresholds depends on the magnitude of ε. Suppose first $\varepsilon < \frac{\sqrt{2}}{2}$ so that $(2-\varepsilon^2) > 2(1-\varepsilon^2) > 1$. If $s > 2(1-\varepsilon^2)\alpha$, $(2,2)$ can never be searched and no partial match can be followed up:

$$E[u_{1,1}|I \equiv \emptyset] = 2\alpha - s$$

$$
\begin{aligned}
E[u_{1,2}|I \equiv \emptyset] = {} & \alpha(1-\varepsilon)\alpha(1+(1+\varepsilon)) - s \\
& + \alpha(1-(1-\varepsilon)\alpha) \\
& + (1-\varepsilon)\alpha(1-\alpha)(1+\varepsilon) \\
& + (1-\alpha)(1-(1-\varepsilon)\alpha)[\alpha(1-\varepsilon)\alpha(V+(1+\varepsilon)) \\
& \qquad + \alpha(1-(1-\varepsilon)\alpha)V + (1-\varepsilon)\alpha(1-\alpha)(1+\varepsilon) - s] \\
= {} & \left[(2-\varepsilon^2)\alpha - s\right](2-\alpha(2-\alpha(1-\varepsilon)-\varepsilon))
\end{aligned}
$$

Clearly, as $s \to (2-\varepsilon^2)\alpha$ the latter shrinks to zero. Suppose $s \to 2(1-\varepsilon^2)\alpha$. Then:

$$E[u_{1,1}|I \equiv \emptyset] = 2\alpha\varepsilon^2$$

$$E[u_{1,2}|I \equiv \emptyset] = 2\alpha\varepsilon^2 - \alpha^2\varepsilon^2(2-\alpha-\varepsilon+\alpha\varepsilon)$$

Since for the lowest possible s in the interval $(1,1)$ is the optimal starting point and since $(2-\alpha(2-\alpha(1-\varepsilon)-\varepsilon)) > 1$, which implies that expected utility generated starting from $(1,2)$ decreases faster than the one generated starting from $(1,1)$, it follows that the latter is always the optimal starting point in this interval.

Suppose now $\varepsilon > \frac{\sqrt{2}}{2}$ so that $1 > 2(1-\varepsilon^2)$: now, following up on a partial match towards a mainstream attribute is possible, but it is not possible to inspect a niche attribute in isolation. $E[u_{1,1}|I \equiv \emptyset]$ is unchanged; $E[u_{1,2}|I \equiv \emptyset]$ now has an additional follow-up search possible:

$$
\begin{aligned}
E[u_{1,2}|I \equiv \emptyset] = {} & \alpha(1-\varepsilon)\alpha(1+(1+\varepsilon)) - s \\
& + \alpha(1-(1-\varepsilon)\alpha)(\mathbf{2}\alpha + (\mathbf{1}-\alpha) - \mathbf{s}) \\
& + (1-\varepsilon)\alpha(1-\alpha)(1+\varepsilon) \\
& + (1-\alpha)(1-(1-\varepsilon)\alpha)[\alpha(1-\varepsilon)\alpha(1+(1+\varepsilon)) \\
& \qquad + \alpha(1-(1-\varepsilon)\alpha) + (1-\varepsilon)\alpha(1-\alpha)(1+\varepsilon) - s]
\end{aligned}
$$

Suppose $s \to 2(1-\varepsilon^2)\alpha$. Then:

$$E[u_{1,2}|I \equiv \emptyset] = 2\alpha\varepsilon^2 - \alpha[\alpha(1-\varepsilon^3) - \alpha^2(1-\varepsilon)^2] < 2\alpha\varepsilon^2$$

The same argument applies: $(1,1)$ is the optimal starting point whenever $s > 2(1-\varepsilon^2)\alpha$.

Consider now the segment $(1-\varepsilon^2)\alpha < s \leq 2(1-\varepsilon^2)\alpha$: $(2,2)$ can be searched but individual niche attributes cannot. If $\varepsilon > \frac{\sqrt{2}}{2}$, $2(1-\varepsilon^2) < 1$ and so searching one mainstream attribute is feasible. Then:

$$E[u_{1,1}|I \equiv \emptyset] = (2\alpha - s)(2-\alpha(2-\alpha)) - 2\alpha\varepsilon^2(1-\alpha)^2$$

$$E[u_{2,2}|I \equiv \emptyset] = 2\alpha\left(2-\alpha(1-\varepsilon)-\varepsilon^2\right) - s\left(2-\alpha^2(1-\varepsilon)^2\right)$$

$$E[u_{1,2}|I \equiv \emptyset] = \alpha\left(4-\alpha(3-(1-\varepsilon)\varepsilon)(1-\varepsilon)+\alpha^2(1+\varepsilon)(1-\varepsilon)^2-2\varepsilon^2\right) - s(2-\alpha(1-\varepsilon))$$

Comparing the two latter for the admissible values of ε leads to the following condition:

$$E[u|x^1 = x_{2,2}] > E[u|x^1 = x_{1,2}] \iff s < \frac{1-\alpha+\varepsilon(1-\varepsilon+\alpha\varepsilon)}{1-\alpha+\alpha\varepsilon}\alpha$$

which is clearly higher than α. Comparing the first and the second instead leads to:

$$E[u|x^1 = x_{2,2}] > E[u|x^1 = x_{1,1}] \iff s < \frac{1-\alpha+\varepsilon(1-2\varepsilon+2\alpha\varepsilon)}{2-2\alpha+2\alpha\varepsilon-\alpha\varepsilon^2}2\alpha$$

which can be shown to be between the two relevant thresholds up until the value α that splits the final threshold in too. This fully characterize the optimal starting point when $\varepsilon > \frac{\sqrt{2}}{2}$.

If $\varepsilon < \frac{\sqrt{2}}{2}$, the same result applies as for the $s \leq \alpha V$ segment. The only segment left is $2(1-\varepsilon)\alpha \geq s > \alpha V$. Now, no partial match can be optimally followed up: after a bad match, however, the product that shares no attribute with the first sampled can be searched.

The usual procedure leads to $(2,2)$ being the optimal starting point for $s < \frac{4(1-\alpha)(1-\varepsilon)}{2-\alpha(2-\varepsilon)}\alpha$, which is compatible with the thresholds as long as $\varepsilon < \frac{1}{2}\left(< \frac{\sqrt{2}}{2}\right)$, and $(1,1)$ being the optimal starting point otherwise. The results are displayed in Figure 4.

// IMPRINT

ZEW – Leibniz-Zentrum für Europäische Wirtschaftsforschung GmbH Mannheim

ZEW – Leibniz Centre for European Economic Research

L 7,1 · 68161 Mannheim · Germany
Phone +49 621 1235-01
info@zew.de · zew.de