Multi-Attribute Search

Jacopo Gambato*

May 12, 2023

Abstract

I study a model of directed search where a consumer inspects products sharing attributes. I embed the resulting optimal search process in an environment in which a multi-product monopolist decides about the bundle of products offered and their prices. Consumers discover their preferences over said attributes rather than over individual products. The emerging search patterns reveal buyers' preferences that the seller can exploit: by setting different prices, the monopolist can encourage specific search paths to arise. In some cases, the monopolist has an incentive to de-list specific products to induce consumption of more expensive products. The results hold in a world in which all products are ex ante identical.

Keywords: targeted search, multi-product monopoly, learning, revealed preferences

JEL Codes: D42, D83, L12, L15

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

1. Introduction

Consider a consumer who wishes to purchase a shirt in a store. She is aware of the selection available but does not know what kind of shirt she wants exactly. Looking for specific kinds of shirts and then trying them on is time-consuming. However, she does not need to try them all: shirts share certain attributes among each other. For example, once the consumer has tried on a red, cotton shirt, she learns whether she wants a shirt that is made of cotton or not, or that is red or not. The next shirt she will try on will change depending on what she has learned. The process will repeat until nothing worth inspecting is left.

In this paper, I study the features of the optimal search process when products are correlated through shared attributes. I consider an environment similar to that introduced by Lancaster (1966) in which consumers value products based on their attributes. This framework allows consumers to select what to search and to adapt their strategy after each realization. The dynamic is as follows: different products sharing an attribute are known to be valued identically with respect to that. Through the search process, consumers learn their preferences for said attributes rather than for individual products. The result of any given inspection makes the consumer update her expectations for the remaining products based on which attributes they share. This, in turn, instructs the next inspection.

I contribute to the existing consumer search literature by allowing inspection of one product to affect the expected return of inspecting a different one. In many circumstances, this represents well consumer search behavior: if a consumer learns that she dislikes a certain attribute in a product, she would rationally try to avoid other products that share that attribute. If attributes can be assumed to be valued independently, this dynamic only arises in a world in which products are defined by multiple attributes. To clarify this distinction, consider the initial example again. Suppose that all shirts were only defined by their color, and that all shirts' color was unique[1]. By trying on a red shirt, the consumer can learn only whether she likes, or not, the color red. If preference over different colors can be assumed to be independent, the consumer cannot use this information to decide what she should search next. When products are defined by multiple attributes the same needs not happen: trying a red, cotton shirt allows the consumer to restrict her attention to shirts that are red, made of cotton, on neither under the same assumption of independence.

Most of the existing consumer search literature studies random search processes. A recent wave of contributions, however, have instead focused on ordered search. In some cases[2], the order in which the consumer searches is fixed exogenously; in others[3] it is determined by the consumer. In this last case, often referred to as "directed" search, consumers determine the optimal search order based on some observable attributes. Recent contributions by Choi et al. (2018) and Haan et al. (2018), for example, assume that prices of all products are known before inspection by the consumer. The latter additionally assumes that products have two attributes, one directly observable and one that needs to be discovered through search.

[1] This is without loss of generality: if color is the only relevant attribute, two shirts sharing the same color are effectively the same product

[2] Arbatskaya (2007), Zhou (2011)

[3] Choi et al. (2018), Haan et al. (2018), Anderson et al. (2020)

The features instructing the order of search in these models are, however, never shared between products. This implies that the order of search cannot be affected by past realizations. A contribution of this paper is to bring together these two features – observable attributes instructing the search process, and shared attributes allowing the consumer to adapt as they search. The resulting optimal search process must accounts for all possible follow-up inspections. I show that in this environment every inspection is optimally selected accounting for what it can imply for the other available products. In particular, inspection of a product informs the consumer of her taste for other available products. Inspecting multiple products not sharing any attribute indirectly informs the consumer of her preference for products that share attributes with them. Each search can be then followed up with a “correction” towards products still uninspected but with known realization. How many available products there are and how they are related through their attributes matter.

Mechanically: with every inspection, expected utility of searching further is updated through the shared attributes. Available products not yet inspected can then be re-ordered based on the expected utility of the search paths their inspection induces. When all products are identical before the search process starts, they carry the same information about other products. They also imply the same search paths conditional on the possible outcomes. This is not the case, for example, when products representing some combinations of available attributes are missing. When products are *ex ante* identical, I show that the optimal search process is equivalent to a myopic one under the assumption of a two point distribution.

To study the implications of learning in directed search, I embed the search process in an environment in which a multi-product monopolist selects product menu and pricing. Like in Choi et al. (2018) and Haan et al. (2018), I allow for prices to be known to the consumer before search starts. Prices contribute to determine the order in which consumer search for their preferred product. When two products are identical, in expectation, in everything but the price, it is clear that the cheaper one would be inspected first. Since the outcome of each inspection instructs the next, each inspection effectively reveals the consumer’s learned preferences. The monopolist can then price products differently to encourage consumers to self-sort based on their preferences.

Setting different prices, however, might induce the consumer to deviate from the monopolist’s preferred order of search. I show that in some cases the monopolist has an incentive to restrict the supply by removing specific products from the menu. This happens, in particular, when search is cheap and likelihood of a match is high. Searching off the optimal path is relatively less punishing in this case, and it can be more rewarding for the consumer, in expectation, because of the difference in prices. The monopolist, then, has an incentive to remove these alternative paths to induce his preferred order of search. Whenever this is the case,the monopolist strictly prefers a uniform pricing strategy over setting different prices for different products.

The results highlight the ability of multi-product firms to induce consumers to inspect products in a specific order when they can coordinate their product menu. By anticipating how a consumer would react after observing a product, the firm can encourage search towards better suited products, and profit off the consumer’s incentive to find good matches. The monopolist wants the consumer to keep searching whenever possible: what is learned through inspection of a product makes the consumer fine-tune her selection. This fine-tuning implies different search paths for different realizations. The monopolist can increase profits by setting higher prices along these paths without discouraging the consumer to search. Strikingly, different prices can emerge even if products are *ex ante* identical from the consumer's perspective.

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 monopolist that selects which products to make available and their price (Section 4). I conclude in Section 5.

1.1. Related literature

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

Weitzman's result crucially relies on the assumption that boxes are independently distributed. I remove this assumption by introducing attributes shared across the available products. In this environment, Weitzman's result generally does not apply. When products share attributes, each inspection affects more than just the inspected product. As more products are inspected, others that share attributes with them end up being fully revealed as a result. The expected gain from inspecting a different option must account for the probability of correcting towards an uninspected but fully discovered product. This probability depends on the best past realizations attribute by attribute. Therefore, ranking products based on their myopic expected value does not capture the whole value of inspecting them. This implies another difference in the optimal search process: unsampled products are reordered based on the expected utility of inspecting them after each search to account for what the consumer has learned.

Few existing papers incorporate conditional search order. A notable exception is the contribution by Doval (2018): Doval extends Weitzman's search process by allowing the consumer to take an uninspected box. The author shows that this option changes the relative value of the available boxes, since each unopened box represent a different outside option. More closely related to this paper is recent work by Anderson et al. (2021). The authors represent the search process on a tree. The nodes of the tree represent information regarding multiple available products. With this set-up the authors argue that product closer to each others can be considered complements rather than substitutes. When a product gets more valuable in expectation, it re-routs search towards itself. Closer products that would not have been considered before gain prominence as a result. The process is similar to the learning proposed in my framework. While they motivate the process through sequential information acquisition, however, my set-up is better suited to represent search for experience goods.

Most of the literature considers search processes that features unchanging order. This is most obvious when order of search is exogenous, as in Arbatskaya (2007), Zhou (2011), and the prominence literature (Armstrong et al. (2009), Armstrong and Zhou (2011)). These papers focus on the effect search order has on equilibrium prices. When search order is predetermined and products are homogeneous, as in Arbatskaya (2007), equilibrium prices decrease in the search order. Armstrong et al. (2009) and Zhou (2011) show that the opposite result emerges when products are heterogeneous: since consumers who keep searching must have been unhappy with their past realizations, sellers searched later can charge higher prices in equilibrium.

More recently, the opposite relation has been explored: Choi et al. (2018), Haan et al. (2018) and Shen (2015) study the effect of posted prices on search order. They study different competitive settings in which sellers costlessly advertise prices. Prices instruct the order of search of consumers, since consumers want to search for cheaper products first all else equal. Since sellers want to undercut each other to gain prominence in the search order, pinning down an equilibrium requires consumers to be heterogeneous enough. Choi et al. (2018) introduces heterogeneity in the form of different mean expected qualities. Haan et al. (2018) introduces multiple attributes, one of which is known before search starts and, therefore, instructs the optimal order of search. Anderson et al. (2020) obtains similar results by introducing heterogeneity through the search cost distribution. Since the multi-product monopolist I focus on does not have an incentive to undercut himself, this is not necessary in my setting.

Correlation in products has been incorporated in other ordered search models: Shen (2015), for example, embeds the search process in a Hotelling framework. Consumers know their location and observe prices, and can choose to search to visit the firm they are farther from. Armstrong and Zhou (2011) also embeds the search process in a Hotelling framework to study the effect of prominence on pricing. In both settings, the available products are perfectly negatively correlated. This implies that there is no uncertainty left regarding consumer preferences after one search. When products are correlated through multiple attributes this is not the case, and a consumer learns about her preferences in a multi-step process.

I further contribute to the growing literature on learning in search. Besides the aforementioned paper by Anderson et al. (2021), a closely related example is the recent contribution by Preuss (2021). Preuss studies a random search process in which a consumer is uncertain about how much she values a certain product. To find out, she must inspect them until she finds a variant she likes. Alternatively, the consumer can randomly learn how much she would value a variant she likes after failing to find it.

The difference lies in how products are defined in these two environments. In Preuss (2021), products are defined by a single attribute, and the consumer can only learn that she should stop searching before finding a product she likes and before exhausting all possible options. This is still true in my setting. Additionally, in my setting, the consumer can learn what she should search next: when products are defined by multiple attributes, not all products need to be affected by one inspection in the same way.

Another notable example comes from Garcia and Shelegia (2018). The authors study search when consumers can learn about the distribution of match value of differentiated products by observing other consumers' purchases. Unlike in my setting, this interpretation of learning affects search outside of the search process itself. The same is true for Greminger (2022), who reinterprets learning as information acquisition as a separate action available to the consumer.

Finally, the paper contributes to the literature of pricing in search. The standard Wolinsky (1986) model, and most of the literature that followed, focus on competitive settings. Instead, I study pricing as implemented by a multi-product monopolist who can coordinate prices of differentiated products. In most papers with the same premise[4], consumers visiting a multi-product sellers learn pricing and valuation for all the products he offers at the same time. Instead, I study within-firm directed search. From this point of view, the paper can be positioned in between recent contributions by Petrikaitė (2018) and Nocke and Rey (2022).

Petrikaitė (2018) shows that a multi-product monopolist can manipulate search costs to induce consumers to stop at the most profitable products. By inducing a certain search order, the monopolist makes consumers self-sort to maximize profit. In my model, the same happens. The result is, however, the opposite: different prices emerge when the monopolist encourages, rather than hinders, search. The difference lies in the learning process. When a consumer learns what he likes, he has an incentive to search further to find a more suitable product. With every inspection, the consumer restricts her attention to more suitable options. The monopolist can profit off this pattern by setting different prices for different products. In particular, the monopolist can set higher prices on a path that would only be reached by a consumer whose preference align with it.

Nocke and Rey (2022), instead, studies the incentives of a multi-product seller to "garble" product information to induce consumers to search longer. Since search costs are assumed to be fixed, the firm has no incentive to price discriminate. This is not the case in my setting: different prices emerge because different search patterns reveal different preferences learned through inspection. This could not be the case in Nocke and Rey (2022) because, like Preuss (2021), they focus on single-attribute products: garbling information affects incentives to keep searching, not what should be optimally searched next.

2. The Model

I consider an industry with products differentiated with respect to two attributes[5]. A product (a, b) is identified by attributes $a \in \{a_1, a_2\}$, $b \in \{b_1, b_2\}$. Thus, the product space N consists of 4 products:

$$N = \{(a_1, b_1), (a_1, b_2), (a_2, b_1), (a_2, b_2)\}$$

The monopolist selects which of these products to make available; the set of available products is denoted by $A \subseteq N$. Products are only differentiated through their attributes and are otherwise identical in quality.[6]

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 for the available attributes; she learns the realization of each attribute separately by inspecting a product characterized by it. In line with existing models[7] I assume that *ex post* utility generated by a generic product (a_i, b_j) takes the

[4]Zhou (2014); Rhodes (2015); Rhodes et al. (2021)

[5]The framework is adapted from Smolin (2020)

[6]As an example: shirts differ in their color (attribute a) and their fabric (attribute b). A shirt is either red (a_1) or blue (a_2), and either made of cotton (b_1) or polyester (b_2).

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

form:

$$u(a_i, b_j) = a_i + b_j = u_{i,j}$$

where $u_{i,j}$ represents the realized ex post utility. I assume each attribute to be i.i.d random variables distributed according to a cumulative distribution function F. Defining generic attributes $a \in \{a_1, a_2\}$, $b \in \{b_1, b_2\}$, and $y \in \{a_1, a_2, b_1, b_2\}$, I further assume $F(y)$ be twice-differentiable everywhere. Expected utility of an unsampled product (a, b) is then:

$$E[u(a, b)] = \int \int (a + b) dF(a) dF(b)$$

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

$$E[u(a, b)] = a + \int b \, dF(b)$$

The sequential search process is with free recall: a 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. The consumer knows all distributions and s.

A multi-product monopolist (he) selects which of these products to make available to the representative consumer (that is, he selects $A \subseteq N$), and their respective prices. He is also aware of distribution F and search costs s. The monopolist can influences 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. I assume all production costs to be constant and normalized at zero.

To solve the monopolist pricing problem I assume attributes to follow a Binomial distribution: $F(y) \equiv B(1, \alpha)$, where $\alpha \in (0, 1)$ refers to the probability that the consumer likes any inspected attribute y[8]. Expected utility of an uninspected and of a partially known product in this case are:

$$E[u(a, b)] = 2\alpha^2 + 2\alpha(1 - \alpha) = 2\alpha$$

$$E[u(a, b)] = a + \alpha$$

respectively. Awareness of the distribution in this simplified environment requires α to be known by all agents.

For the pricing game, I consider Sub-game Perfect Equilibria; sequential rationality in this context refers to prices and expected profits to be internally consistent with the optimal search pattern of the consumer, and viceversa. When the consumer is indifferent between stopping and searching again, or when she is indifferent between two products to purchase, indifference is always resolved in favor of the monopolist – that is, the most profitable outcome is selected.

[8]In this specification, one can think of the search process as an effort to find an attribute that satisfies a need of the consumer.

3. Search

The consumer can inspect any product in $A \subseteq N$. Search is sequential: after any inspection, the consumer chooses to stop or search a different product among the uninspected available ones. If she chooses to keep searching, she must also choose which one. 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, $A \setminus I$. As is shown below, the set of uninspected products can be further partitioned based on the amount of information that the consumer has acquired about them through the search process.

To fix ideas, consider the case in which $A \equiv N$ (that is, all products are available). To focus on the search process itself, I assume for now that the price of all available products are normalized to zero. Suppose the consumer already inspected one of the products. Since all products are available, inspecting (a_1, b_1) first is without loss of generality[9]. More in general: whenever the first product to inspect can be chosen without loss of generality, I assume that products are inspected in increasing order based on their indices. After the first inspection, the consumer then has learned realizations a_1 and b_1. We are interested in finding which of the remaining products should be inspected next, if any.

3.1. Discovering two attributes

Suppose the consumer has already inspected (a_1, b_1). Further, suppose that the consumer decided to search the product sharing no attributes with it – that is, (a_2, b_2). This product generates value in two ways. Myopically, it can lead to a realization $u(a_2, b_2) > u(a_1, b_1)$. Additionally, inspecting (a_2, b_2) reveals the remaining unknown information regarding the other two products. Once (a_1, b_1) and (a_2, b_2) are inspected, all products' *ex post* utility is revealed. Thus, inspecting (a_2, b_2) allows the consumer to correct towards the remaining products if doing so is worth the cost of the additional search[10].

The expected utility of inspecting (a_2, b_2) needs to distinguish four possible outcomes. First, (a_1, b_1) can still be the best product to keep after all information is acquired. Alternatively, if both $a_2 > a_1$ and $b_2 > b_1$, (a_2, b_2) would be selected. Finally, if either $a_1 - s > a_2$ and $b_2 - s > b_1$, or $b_1 - s > b_2$ and $a_2 - s > a_1$, the consumer would perform an additional search to reach (a_1, b_2) or (a_2, b_1) respectively. The expected utility of each of these scenarios can be directly computed from the realization of a_1 and b_1. Importantly, the probability of a correction taking place (either towards (a_1, b_2) or (a_2, b_1)) depends both on the realization of $a_1 + b_1$ as a whole and the individual realizations of a_1 and b_1. This is why the reservation price approach as proposed by Weitzman (1979) cannot be generally applied. In the standard Pandora search process, each product can be represented by a single index determining the value of inspecting it. In this environment, the same value depends non-trivially on the past realizations. The expected value of the inspection given these past realizations must be used instead.

[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

[10] As shown below, this correction is never rational when the consumer selects to search a product for which she already knows the realization of one of the two attributes. If this correction was costless, it follows immediately that searching (a_2, b_2) would be always optimal. Search costs in this environment, then, can be understood as the costs of "finding" the product with the correct combination of attributes.

The myopic value of searching (a_2, b_2) (that is, the expected utility of inspecting this product without accounting for the possible correction to (a_1, b_2) and (a_2, b_1)) given realizations a_1 and b_1 is:

$$m(a_2, b_2)|_{I=(\{a_1,b_1\})} = u_{1,1} \int_{-\infty}^{u_{1,1}-b_1} \int_{-\infty}^{u_{1,1}} dF(b)dF(a) + \int_{u_{1,1}-b_1}^{\infty} \int_{u_{1,1}}^{\infty} (a+b)dF(b)dF(a) - s \quad (1)$$

where m stands for myopic. This value simply captures the expected gain of searching (a_2, b_2) given the free recall assumption. Once this second search has taken place, the consumer learns the *ex post* utility of all available products. The consumer would, then, decide to keep product (a, b) such that:

$$(a, b) = \underset{(a_i, b_j) \in A}{\arg\max}(a_i + b_j + s \cdot \mathbb{1}(i \neq j)), \qquad i, j \in \{1, 2\}$$

To compute the expected utility of this inspection one must determine the probability of the individual attribute surpassing their respective counterparts, and their expected value conditional on it. In particular: a_1 would be ultimately kept if and only if:

$$a_2 < a_1 - s$$

or:

$$a_1 - s < a_2 < a_1 + s \wedge (a_1 + b_1) > (a_2 + b_2)$$

The first condition does not depend on b_2, and reflects the ability of correcting from (a_2, b_2) to (a_1, b_2). The second condition refers to the case in which a_2 is not high enough to be worth keeping, and (a_1, b_1) is overall better than (a_2, b_2). Similar conditions can be obtained for b_1. If these conditions are not satisfied, instead, it must be the case that a_2 is worth keeping. If both a_2 and b_2 are worth keeping at the same time, (a_2, b_2) is selected; otherwise, the consumer would perform a correction.

Given the above discussion, the following lemma reports the expected value of searching to discover two new attributes:

Lemma 1. *Define:*

$$q_a(a_1, b_1) = Pr[a < a_1 - s] \cup Pr[a_1 - s < a < a_1 + s \wedge (a + b) < a_1 + b_1]$$

$$q_b(a_1, b_1) = Pr[b < b_1 - s] \cup Pr[b_1 - s < b < b_1 + s \wedge (a + b) < a_1 + b_1]$$

Expected utility of inspecting (a_2, b_2) given realizations a_1, b_1 can be computed as:

$$\begin{aligned} E[u(a_2, b_2)]|_{I=(\{a_1,b_1\})} &= q_a(a_1, b_1)a_1 + q_b(a_1, b_1)b_1 \\ &+ (1 - q_a(a_1, b_1)) \int_{a_1}^{+\infty} a\, dF(a) \\ &+ (1 - q_b(a_1, b_1)) \int_{b_1}^{+\infty} b\, dF(b) \\ &- [1 + F(a_1 - s)[1 - F(b_1 + s)] + F(b_1 - s)[1 - F(a_1 + s)]]s \end{aligned} \quad (2)$$

Proof. All calculations and the explicit formulas for $q_a(a_1, b_1)$ and $q_b(a_1, b_1)$ can be found in

Appendix A ■

It is clear that for this inspection to be feasible it must hold:

$$E[u(a_2, b_2)]|_{I=(\{a_1,b_1\})} > u(a_1, b_1)$$

Even when this is satisfied, however, searching (a_2, b_2) is not necessarily the optimal choice. As mentioned, the individual values of a_1 and b_1 determine how attractive it is to search for (a_2, b_2). Intuitively, this option is most valuable when both a_1 and b_1 are not too high. If their realizations are very different, however, it might be preferable to keep the highest one and give up the ability to correct after the second search.

3.2. Keeping one attribute

Suppose now that the consumer decided to search a product sharing one of the two attributes with (a_1, b_1); without loss of generality, suppose that product is (a_1, b_2). Before this search is performed, the consumer already knows the realization for a_1. Since F is known, she also has expectations over the utility associated with searching (a_1, b_2). In particular:

$$m(a_1, b_2)|_{I=(\{a_1,b_1\})} = a_1 + b_1 \int_{-\infty}^{b_1} dF(b) + \int_{b_1}^{\infty} b\, dF(b) - s \tag{3}$$

Like before, $m(a, b)|_I$ represents the myopic value of searching a product given the information learned through past searches. It accounts for the cost of searching the product and the realization of known attributes (a_1 in this case). The free recall assumption implies that, if $b_2 < b_1$, (a_1, b_1) would be kept and (a_1, b_2) would be discarded. Unlike before, however, it can be shown that following this strategy does not allow for any correction:

Lemma 2. *If it is ever optimal to inspect a product characterized by an attribute with known realization, it is never optimal to inspect a product not characterized by it afterwards.*

Proof. The Proof can be found in Appendix A ■

In words, Lemma 1 says that once an attribute is deemed worth keeping by the consumer, it will never not be worth keeping. The proof follows a contradiction argument. Suppose the consumer decided to search keeping a previously sampled attribute. Intuitively, this attribute must represent high enough *ex post* utility. Any correction from this newly inspected product would necessarily correct away from this attribute. If it is ever worth it to do so, the product the consumer corrects to must have been better, in expectation, than the one that was actually inspected. In particular, it must be better in expectation in the attribute that was optimally elected to keep, which is a contradiction. The result follows.

The lemma implies that the value of searching along a previously inspected attribute is equivalent to its myopic value:

$$E[u(a_1, b_2)]|_{I=(\{a_1,b_1\})} = a_1 + b_1 \int_{-\infty}^{b_1} dF(y) + \int_{b_1}^{\infty} y dF(y) - s = m(a_1, b_2)|_{I=(\{a_1,b_1\})} \tag{4}$$

Once again, this search can only be feasible if:

$$E[u(a_1, b_2)]|_{I=(\{a_1,b_1\})} > u(a_1, b_1)$$

Whether it is the better to inspect a product sharing one or no attribute with (a_1, b_1) depends on the individual realizations a_1, b_1. Searching to discover two new attributes allows for correcting, keeping one discovered attribute does not. On the other hand: keeping an attribute from the first search minimizes the number of searches. For the latter strategy to be optimal, it must hold that one between a_1 and b_1 is particularly high (and thus worth keeping) and the other is particularly low (and thus easy to improve upon). Otherwise, it would be better to inspect (a_2, b_2) or stop at (a_1, b_1).

3.3. The optimal search policy

The exercise above illustrates the different paths that arise when different products are searched. Conditional on previously acquired information, searching to discover two uninspected attributes is costlier in expectation, but rewards the consumer with more thorough information. Searching along a previously inspected attribute is cheaper in terms of expected search costs, but prevents the consumer from, potentially, discovering a better fit.

Once these differences are incorporated, the sequential search process can be fully characterized based on the updating expected values. In particular, whenever an inspection is performed, the best product to inspect is the one with the highest expected utility net of the total expected search costs. The difference between this process and a myopic one comes from the implication of past observations I on the consumer expectations over uninspected products. After an inspection is performed, the uninspected products can be reordered based on the information just learned. If the highest such expected value surpasses the highest known ex post utility, the product associated with it is inspected; otherwise, the search process stops and the most valued known product is kept.

Going backwards: the expected utility of the search process depends non-trivially on the composition of A. Once the available products are identified, however, the search patterns arising from all possible outcomes of the first search allow to determine $E[u(a, b)]|_{I\equiv\emptyset}, \forall(a, b) \in A$. In particular, $E[u(a, b)]|_{I\equiv\emptyset}$ is equal to the best expected outcome of the possible search paths, weighted for their probability of arising given the first realization, minus the total expected search costs.

Under the current specification all products are equivalent before any information is learned. This is the case whenever each attribute can be found in combination with all other available attribute: intuitively, this feature implies that all product carry information relevant for the same number of other products. When this is not the case, however, inspecting different products first lead to different search paths arising for the same outcome. While the logic of the optimal search survives, the actual search paths depend crucially on which attributes are available, and in which combination. Any precise mapping of the optimal search path in these cases is, therefore, game specific.

Under the restriction that all products are *ex ante* identical with respect to their distribution and their relative position to other products, the search process studied above can be intuitively generalized. In particular, one can consider an arbitrary number of attributes and, for each attribute, an arbitrary number of variants. For expositional purposes, suppose:

$$a \in \{a_1, a_2, a_3\} \qquad b \in \{b_1, b_2, b_3\}$$

Suppose once again that (a_1, b_1) has already been inspected. The choice to inspect next a product sharing one or no attributes with (a_1, b_1) is the same as before. Indeed, after the first search all products sharing the same attribute with (a_1, b_1) are identical to each other in expectation. So are all the products sharing no attributes with it. From Lemma 1: if a product sharing an attribute with (a_1, b_1) is searched right after, the consumer would proceed by choosing either inspected product or the remaining product sharing the same attribute with them. If a product sharing no attributes is searched instead, the remaining uninspected products would be reordered based on the information learned with the second search. The consumer would choose to search next the uninspected product with the higher (updated) expected utility if it is higher than:

$$\max_{(a_i,b_j)\in A}(a_i + b_j - s \cdot \mathbb{1}((a_i, b_j) \notin I)$$

which represents the (updated) best known product to keep. Notice that the best past realization is found among both inspected products and non inspected products that share all attributes with some inspected product.

The logic can be extended to more than two attributes as well. With three attributes a, b, c, for example, different degrees of discovery are possible. Searching a product sharing no attributes with any of the previously inspected products allows for correcting in up to two attributes rather than just one. Searching next a product sharing a single attribute with a previously inspected product stops the consumer from discovering her taste for alternatives along that attributes. Unlike before, such an inspection lets the consumer discover two new attributes and, potentially, correct along them. While the process gets more intricate, the core intuition is untouched. A more detailed discussion of the search process in these arbitrarily large games can be found in Appendix A.

Assuming that unknown attributes are in increasing order of their index, the optimal sequential search process can be summarized as such:

Proposition 1. *The optimal search process with $A \subseteq N$ multi-attribute products is such that:*

- *Search starts if $E[u(i)]|_{I\equiv\emptyset} - s \geq 0$ for some product $i \in A$,*
- *After every inspection, expected utility of all products $j \in A \setminus I$ updates based on attributes shared with products $i \in I$,*
- *Expected utility of a product j sharing all but one attributes with products $i \in I$ updates according to Equation 4; otherwise, it updates according to Equation 2,*
- *At every step, the next product $j = \arg\max_{i\in A\setminus I}(E[u(i)]|_I)$ to inspects satisfies:*

$$E[u(j)]|_I - s \geq \max_{i\in A}(u_i - s \cdot \mathbb{1}(i \notin I))$$

- *Search stops if no product $j \in A \setminus I$ satisfies the above condition.*

The process highlighted above has important implications: if products share attributes, the value of inspecting multiple unknown features is more than their myopic expected value. The information learned affect multiple products, which makes the consumer more willing to search than in a world in which optimal search is myopic. Further, because of this added value of information, when a consumer decides to keep a previously inspected attribute and keep searching, she reveals that she values it greatly. The optimal search path sees the consumer sorting herself towards attributes she finds valuable.

To investigate the consequences of this mechanism, I next study the incentives of a multi-product monopolist interested in making these products available. For the sake of tractability, I solve the menu selection and pricing game of the monopolist in a simplified environment in which products are still defined by two attributes ($a \in \{a_1, a_2\}$, $b \in \{b_1, b_2\}$), but each attribute follows a binomial distribution: $y \sim B(1, \alpha), \forall y \in \{a_1, a_2, b_1, b_2\}$. As shown in the example below, this restriction simplifies the analysis considerably.

3.4. Example: binomial distribution

Suppose, as before, that all products are available:

$$A \equiv N = \{(a_1, b_1), (a_1, b_2), (a_2, b_1), (a_2, b_2)\}$$

Since all attributes follow a binomial distribution, each product can only lead to *ex post* utility equal to 0, 1, or 2 depending on the number of attributes the consumer likes. Suppose again that all prices are normalized at 0, and that (a_1, b_1) has been already inspected by the consumer. She has learned the realization of a_1 and b_1 and needs to decide what to search next.

The problem of the consumer is much simpler than in the general model. Since each attribute can only have realization 0 or 1, the correction mechanism studied above can not rationally take place. This follows from the fact that if an attribute does not have *ex post* utility zero, the consumer cannot find a better fit. Therefore, if an inspected attribute is such that $y = 1$, it always kept; if $y = 0$, instead, it is always discarded. If a consumer were to inspect (a_2, b_2) and would want to correct towards one of the uninspected products, it would always have been better to search that product before (a_2, b_2).

Observation of a product still allows computation of the expected utility of the uninspected products according to which attributes they share. After sampling (a_1, b_1), the remaining products see their expected utility updated to:

$$E[u(a_1, b_2)]|_{I=\{a_1,b_1\}} = a_1 + \alpha \qquad E[u(a_2, b_1)]|_{I=\{a_1,b_1\}} = \alpha + b_1$$

$$E[u(a_2, b_2)]|_{I=\{a_1,b_1\}} = 2\alpha$$

These products can then be ordered after this update; the product with the highest expected utility is searched next as long as:

$$\max_{(a,b)\in A\setminus I} E[u(a,b)]|_I - s > \max_{(a,b)\in I} u(a,b)$$

Identifying the next optimal search conditional on the outcome of the first is straightforward:

- if $a_1 = b_1 = 0$, (a_2, b_2) is searched next; no correction can take place because $a_2 \geq a_1$ and $b_2 \geq b_1$,
- if $a_1 = b_1 = 1$, no follow-up search takes place since $a_1 \geq a_2$ and $b_1 \geq b_2$
- if $a_1 > b_1$, (a_1, b_2) is searched next (if $\alpha > s$),
- if $a_1 < b_1$, (a_2, b_1) is searched next (if $\alpha > s$).

Going backwards, expected value of the search process is then:

$$E[u(a,b)]|_{I\equiv\emptyset} = 2\alpha^2 + 2\alpha(1-\alpha)\max(1, 2\alpha + (1-\alpha) - s) + (1-\alpha)^2(2\alpha - s) - s$$

The first addendum refers to (a, b) being the best possible match ($u(a, b) = 2$, with probability α^2). The second addendum refers to the eventuality of the consumer liking only one of the two attributes, and the possible second search that outcome would entail. The third refers to the case in which $u(a, b) = 0$ so that the product sharing no attributes with it would be inspected next. Once again, when all products are available (and therefore *ex ante* identical), inspecting products in increasing order of their index is without loss of generality. Figure 1 exemplifies. When some products are not available this is generally not the case: missing products imply that inspecting different products leads to different search paths for the same outcome. In the next section, I highlight the effect of restricting the supply for the optimal order of search only referring to the specific cases and ways this restriction takes place. A more general discussion can be found in Appendix A.

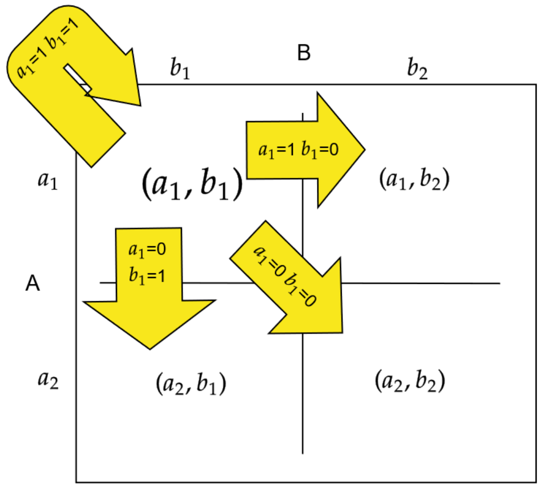

Figure 1: *Optimal search with binomial distribution and all products available, starting from* (a_1, b_1).

4. Monopoly pricing

The monopolist's problem is twofold: he must set up prices to maximize profit, and must select A to generate trade opportunities. The two decisions are related. The consumer search path depends on the price she observes; which prices would deter her from searching depend on the available products. This follows from the characterization above: the consumer is willing to search something with a negative myopic expected value if the information learned through search makes the inspection overall worth it. For this to happen, it is clearly necessary that some products that share attributes with each other are made available. A monopolist can, in principle, price products above their myopic expected value as long as he made available enough products to justify it.

The two decisions – A and $\mathbf{p}(A)$, where the latter represents the vector of prices associated with all products in A – 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 carries more new information than the last. It follows that the highest price that would make the two worth searching is different. If prices are not uniform, however, the consumer could adapt their 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. In either case, prices must be such that they optimally reply to the search path they induce in equilibrium; viceversa, the consumer must search optimally given observed prices.

To study these different interactions, I solve the menu and pricing game of the monopolist considering uniform and differential prices separately. I show that the monopolist can always manipulate prices to induce specific ordering of the consumer search. Moreover, I show that the monopolist 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 monopolist'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 monopolist 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: $A \equiv N$. I then show the monopolist's incentive to restrict the menu and the effect this choice has on consumer search. Throughout, I assume $s < 2\alpha$.

Consider a generic price level p^u. The monopolist 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 example above, the expected utility of performing the first inspection is:

$$\begin{aligned}
&E[u(a_1, b_1)]|_{I\equiv\emptyset} = \alpha^2 \max(2 - p^u, 0) - s \\
&+ 2\alpha(1-\alpha)\max(1 - p^u, \alpha\max(2 - p^u, 0) + (1-\alpha)\max(1 - p^u, 0) - s, 0) \\
&+ (1-\alpha)^2 \max(\alpha^2 \max(2 - p^u, 0) + 2\alpha(1-\alpha)\max(1 - p^u, 0) - s, 0)
\end{aligned} \quad (5)$$

that is: the value of searching (a_1, b_1)[11] is equal to the expected value conditional on which search paths are induced by different outcomes. These in turn depend on the relative value of s

[11] once again, searching in increasing order of the indices is without loss of generality when all products are available and prices are uniform

and α, over which the monopolist 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. It is straightforward to see that the first search path to be prevented by high prices is the one that arises conditional on a bad first match. Indeed, given observation $u_{1,1} = 0$, (a_2, b_2) is searched as long as:

$$\alpha^2 \max(2 - p^u, 0) + 2\alpha(1-\alpha)\max(1 - p^u, 0) - s \geq 0 \tag{6}$$

If instead the consumer observed $a_1 = 1$, $b_1 = 0$, the next search can be considered only if:

$$\alpha \max(2 - p^u, 0) + (1-\alpha)\max(1 - p^u, 0) - s \geq 0 \tag{7}$$

It can then be shown that p^u such that the latter condition binds breaks the former.

A price that discourages a follow-up search after a bad first realization does not prevent search from happening altogether: as long p^u is such that $E[u(a_1, b_1)]|_{I \equiv \emptyset}$ is non negative, the consumer can rationally start inspecting something. With this inspection, the consumer can discover that she likes both attributes, after which she always stop searching since 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:

$$u(a_1, b_1) = 1 - p^u \leq 1 + \alpha - s - p^u = E[u(a_1, b_2)]|_{I = \{(a_1, b_1)\}}$$

which is always satisfied if $\alpha > s$, that is, if inspecting a single attribute is worth the necessary search cost.

In general, the monopolist can select one of two pricing profiles: on one hand, he can elect to price products in a way that encourage a follow-up search after a first bad realization. These prices must make a product just myopically worth searching, or, they must solve equation 6 with equality:

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

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

Alternatively, the monopolist can select higher prices that discourage search after a bad first realization. These prices must be weakly higher than the encouraging counterpart 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 $E[u(a_1, b_1)]|_{I \equiv \emptyset} = 0$, since for any higher price the consumer would not start searching:

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

where ^{E} stands for "encourage", and L, I, and H stand for "low", "intermediate", and "high" respectively.

Two observations are in order. First, lower prices are always feasible whenever higher ones are: when p_H^D does not prevent search, all other options are still available to the monopolist. Second, the monopolist is not interested in his products being inspected, but in his products being purchased. Trade is maximized for $p \leq 1 = p_T$, where T stands for "trade": any higher price requires the consumer to like both attributes in a product to purchase it. Notice that it holds:

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

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

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

that is, p_L is the price that maximizes probability of trade. Overall, when selecting p^{u*} the monopolist chooses between maximizing search efforts, maximizing per-sale revenue, and maximizing probability of trade. Higher prices reduce 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 5, one can obtain the expected profit of the monopolist. These profits can then be directly compared and lead to a unique equilibrium price for all possible combinations of α and s. In particular, the prices above lead to expected profits:

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

which maximizes probability of trade and is always valid,

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

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.

The candidate prices reflect the relative importance of encouraging search and extracting rent conditional on the search taking place. p_L maximizes probability of trade: at this price level, the consumer is always encouraged to either purchase the first product found or to keep searching, and trade can take place as long as the consumer finds one attribute she likes. p_I^D prevents the consumer to search again after a bad first realization if s is high enough, but leads to higher profit conditional on trade taking place. p_H^E encourages search in the appropriate segment, but only generates trade if the consumer finds a product she likes in both attributes. Finally, p_H^D discourages search but leads to the highest profit conditional on trade taking place.

Intuitively, higher prices are preferable, for the monopolist, for low values of s and high values of α. For such parameters, the consumer is easily encouraged to start searching. In particular, it can be shown that for α low enough, p_L is always profit-maximizing. For high values of α, instead, which of the four candidate prices is selected depends on the relative value of s: the lower s is, the more aggressive the pricing can be.

Given prices, it is straightforward to obtain 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. In this case, starting from (a_1, b_1) is without loss of generality. However, for high enough prices not all products can be inspected after fixing a starting point. From the above discussion it emerges that if the monopolist optimally selects $p^{u*} \in \mathbf{p^D}$, conditional on the consumer starting from (a_1, b_1), inspection of (a_2, b_2) could not rationally take place. Indeed, (a_2, b_2) would only be inspected after a bad first realization, but $p^{u*} \in \mathbf{p^D}$ prevents this search from taking place. It follows that when the monopolist selects a price that prevents search after a bad first realization, introducing three or four products is equivalent from the monopolist's perspective. This equivalence, however, is a byproduct of the unrealistic assumption of zero production costs. It is then 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 (a_2, b_2) were to be removed, search starting from (a_1, b_1) would be unaffected. Starting from any other product, however, would generate negative expected utility of search. By removing a product, the monopolist effectively "locks" the consumer into a specific search path. The values α, s and p^u determine which search paths can be taken by the consumer. Given these search paths, products are introduced. For example: if it holds $s > \alpha$, 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 2. *Consider a multi-product monopolist selecting optimal menu $A \subseteq N$ and uniform pricing p^u of multi-attribute products. Define $\mathbf{p^E}$ the set of prices that* ***encourage*** *search after a bad realization, and $\mathbf{p^D}$ the set of prices that* ***discourage*** *it. In equilibrium:*

- *If $s < \alpha$ and $p^{u*} \in \mathbf{p^E}$: $A \equiv N$, and the consumer can start searching from any available product,*
- *If $s < \alpha$ and $p^{u*} \in \mathbf{p^D}$: $A \subset N$, and the consumer is restricted in where she can start searching ,*
- *If $s > \alpha$: $p^{u*} = \frac{2\alpha - s}{\alpha(2-\alpha)}$, $A \subset N$, and the consumer can start searching from any available product.*

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

The monopolist values higher probabilities of trade taking place: since prices are uniform, the monopolist is not concerned with which product is purchased as long as one is. Selecting prices that do not hinders 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 monopolist can raise price and choose not to introduce all possible varieties; as a consequence, consumers can be worse off in expectation. When the supply is restricted, moreover, the monopolist effectively induces a specific order of search.

At uniform prices, the consumer retains some positive expected value from search when the monopolist 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. The restriction on the monopolist's pricing structure seems sensible: as all products are *ex ante* identical for the consumer, they can be expected to be all priced at the same level. As I will show in the next section, however, the monopolist 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 off good ones.

4.2. Differential prices

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

The trade-off of the monopolist 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 likes at least partially the first option inspected. By pricing along these paths differently, the monopolist 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 uniform price p_L. When this price is optimally selected, it allows the consumer to keep searching after a bad first realization and trade is most likely to take place. In particular, what is needed is that the first product inspected, say (a_1, b_1), and the product that would be searched next conditional on $a_1 = b_1 = 0$, (a_2, b_2), to be priced at p^u. On this path, if the other products were priced higher than p_L, nothing would change since (a_1, b_2) and (a_2, b_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 (a_1, b_2) and (a_2, b_1) were priced slightly higher than (a_1, b_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 (a_1, b_1) and (a_2, b_2) at $p = p_L$, and the remaining products at a higher price, then, the monopolist can achieve both higher prices and higher probability of trade. In doing so, the monopolist erodes at the consumer expected utility without preventing search.

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 monopolist's point of view, in which cases this spread does not affect optimal search order, and, in the cases in which it does, what is the monopolist optimal reply. Henceforth, I assume that (a_1, b_1) and (a_2, b_2) have lower prices and therefore act as possible starting points; further, I keep the assumption of products being 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. As the price increment relies on the consumer learning about which attribute she likes, a higher price can only arise on a path dictated by the consumer finding an attribute to keep. Suppose the consumer inspects (a_1, b_1) and observes $a_1 = 1, b_1 = 0$. Suppose moreover that the optimal base price selected by the monopolist is $p_L \leq 1$. Conditional on inspecting one attribute being worth the cost of inspection ($\alpha > s$), the consumer would want to search (a_1, b_2) if:

$$u(a_1, b_1) = 1 - p_{1,1} \leq 1 + \alpha - s - p_{1,2} = E[u(a_1, a_2)]|_{I=\{(a_1,b_1)\}}$$

which implies $p_{1,2} = p_{1,1} + \alpha - s$, where $p_{1,1}, p_{1,2}$ are the observed prices for (a_1, b_1) and (a_1, b_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. As the monopolist is interested in the highest price that does not dissuade the search, the following candidate prices profile arises:

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

if $\alpha^2 < s < 2\alpha$, 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 exact steps as the former.

Given search under uniform prices, these prices lead to the same probability of trade as their uniform counterparts. Compared to them, however, they lead to higher expected profit since the more expensive products are purchased with positive probability. Notice that this deviation preserves the internal consistency of the search process since the consumer would always inspect the cheapest product first if she has no information on any of the available products. These prices, however, 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 she like it. In this case, the consumer would search (a_2, b_2) hoping to find a good realization instead, and would only inspect the more expensive product if she knows she likes both its attributes. Consider again the candidate prices profile $p_{1,1} = p_{2,2} = p_L$, $p_{1,2} = p_{1,2} = p_L + \delta_L$. After realization $a_1 = 1$, $b_1 = 0$:

$$u(a_1, b_1) = 1 - p_{1,1} \quad E[u(a_1, b_2)]|_{I=\{(a_1,b_1)\}} = 1 + \alpha - s - p_{1,2}$$

$$E[u(a_2, b_2)] = \alpha^2(2 - p_{2,2}) + (1 - \alpha)(1 - p_{2,2}) + \alpha(1 - \alpha)(2 - s - p_{1,2}) - s$$

For α high enough and s low enough, inspecting (a_2, b_2) is a feasible 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 monopolist: the more expensive products now are reached with lower probability. The monopolist can optimally reply in three ways:

- the monopolist can let the consumer search (a_2, b_2) first, and further increase $p_{1,2}$ and $p_{2,1}$ to $(p_{1,1} + 1 - s)$,
- the monopolist can reduce prices $p_{1,2}$ and $p_{2,1}$ to encourage his preferred order of search to arise,
- the monopolist can remove (a_2, b_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 monopolist to condition prices on search behavior. If the consumer has an incentive to search (a_2, b_2) after (a_1, b_1) conditional on $a_1 + b_1 = 1$, he 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 monopolist can make (a_1, b_2) and (a_2, b_1) cheaper. Since the consumer is interested in searching (a_2, b_2) first 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

(a_2, b_2) forces the consumer to take the path the monopolist want 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 (a_2, b_2) when the monopolist would want her to inspect (a_1, b_2) or (a_2, b_1).

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

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

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

- if the monopolist removes (a_2, b_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 monopolist must react accordingly as well. IN particular, for this alternative profile, removing (a_2, b_2) always dominate the other two strategies.

The feasible expected profits under differential prices must be compared to the highest expected profit under uniform prices obtained in the section above. Two results emerge. First, whenever the monopolist had 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. As 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 monopolist'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. Overall, the following equilibrium results emerge:

Proposition 3. *Consider a multi-product monopolist selecting optimal menu $A \subseteq N$ and pricing $\mathbf{p}(A)$ of multi-attribute products. In equilibrium:*

- *Low, uniform prices are set if $s > \alpha$,*
- *High, uniform prices are set for high values of α and low values of s,*
- *All products are introduced if and only if prices are not set uniformly,*
- *if $s < \alpha$, the consumer is always restricted in where she can start searching.*

Proof. All calculations and precise cut-off values for α and s can be found in Appendix B. ■

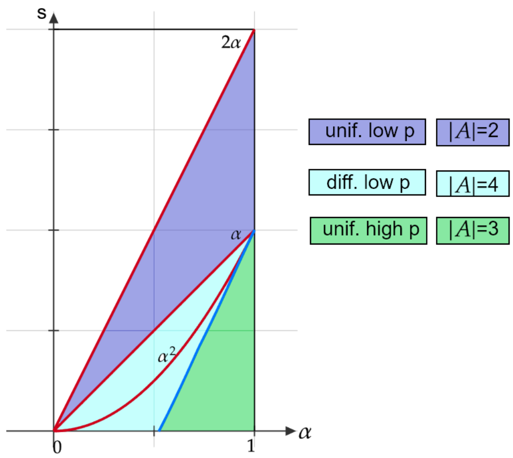

Figure 2: *Equilibrium monopoly menu selection and pricing for all feasible combinations of* $\alpha \in (0, 1)$ *and search costs* $s \in (0, 2\alpha)$

Figure 2 summarizes the equilibrium menu and pricing selection for all feasible combinations of α and s. The two decisions are intertwined. The monopolist has an incentive to make all products available only if they can all be reached, and purchased, with positive probability.

When search costs are very high, only a bad first realization induces the consumer to keep searching: introducing more than two products allow the consumer to randomize her starting point but at no benefit from the monopolist. On the other hand, when search costs are very low the monopolist 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. Moreover, whenever all products are introduced, they are never priced uniformly. This, in turn, implies that the monopolist always has an incentive to induce specific search patters to arise. Further, it implies that in no case the consumer is free to select where to start searching: menu selection and pricing always cause an asymmetry in per-product sale. The monopolist is indifferent between which product or products can act as starting point. *Ex post*, however, the learning dynamic analyzed above leads to some products becoming endogenously prominent despite the fact the all products are equivalent before any inspection has taken place.

4.3. Limitations and discussion

The model comes with some strong simplifying assumptions. First, I assume that only two agents are active: the multi-product monopolist and a representative consumer. Extending the argument to multiple consumers with identical distribution of "taste" for the possible attributes is straightforward. The monopolist commits to product menu and prices before the search starts. Multiple consumers with common search costs and common distribution would all behave as the representative consumer, choosing among the feasible starting point at random but otherwise following the same search paths conditional on the same outcome. This approach does not lead to new insights.

Consumer taste is, however, heterogeneous in many ways. Different consumers can be expected to like different attributes. They can also be expected to be more or less picky about any given attribute. To incorporate this additional heterogeneity, one could think of a population of consumers characterized by the same search costs but with α extracted from some distribution. As shown above, monopoly prices depend crucially on the relative size of α and s. If a price induces a consumer with probability $\underline{\alpha}$ of liking an attribute to search, all consumers with a higher α would start searching as well. This additional level of heterogeneity might affect the ability of the monopolist to make consumers self-sort through pricing. For example: the monopolist would want to encourage consumers with a high α to search to find the best match available by pricing in a way that discourages consumers with low α to start searching at all. Given any active segment of consumers, however, the monopolist would try to target the more profitable consumers by preventing the least profitable ones to start. This dynamic could iterate indefinitely and lead to the equilibrium unraveling.

The results in the main text hold in a world in which products have no vertical differentiation and utility from consumption is only driven by horizontal preferences. With vertical differentiation, if consumers can be assumed to prefer higher quality all else equal, the interests of consumers and seller would naturally align. As shown in Petrikaitė (2018), the monopolist would likely want consumers to start from the highest quality product and price the others to be searched only conditionally on the first realization. Since the dynamic studied in this paper relies on self-sorting on the consumer side, it is less clear what the equilibrium outcome would be if the product the consumer wants to buy and the one the monopolist wants to sell are not the same.

The intuition presented here is likely to hold under the more general distributional assumptions studied above. The simplicity of the binomial distribution implies that, once an attribute is found acceptable, it never needs to be exchanged for an unsampled one under uniform prices. As shown in Section 3, this is still the case for more general distributions, but how good an attribute needs to be in order to be kept is not as trivial. In particular, sampling more informative products (that is, products defined by not yet inspected attributes) is more rewarding in the general model than under the binomial distribution assumption. In different environments, this change in the value of information might lead to either lower prices, to encourage buyers to select an attribute to keep, or to higher prices, to fully profit off consumers finding their best match as they sample different options. Which effect would dominate is likely to depend on the cost of search: if inspection is very costly, searching to discover as many attributes as possible is discouraged. A monopolist, then, is likely to encourage minimal expected searches, since s erodes at his profits through the constraint it imposes on prices.

Finally, a rather strong assumption comes in the form of independent valuation of attributes by the consumer. It is easy to argue that some attributes go better together than others, and that in some cases products can be greater than the sum of their parts. In the current specification, I abstracted away from these cases to focus on the implications of learning in a vacuum. Introducing additional variation in *ex post* utility, for example in the form of a random component ε capturing correlation across attributes, is beyond the scope of this analysis. Nonetheless, this additional generalization would better represent problems of product design and their relationship with consumer search as understood in this model. As per all of the above, this dimension of the problem remains to be investigated.

5. Conclusion

In this paper, I study the implications of product correlation through shared attributes for directed search and the associated incentives of a monopolist to introduce different products and prices to capitalize on consumer learning. Consumers have an incentive to find better matches in their search process as they learn what they like. This dictates their strategy predictably. The monopolist is then able to profit off the learning process by setting differential prices to let consumers self-select based on their preferences.

It is natural to ask what would be the features of an hypothetical equilibrium with competing firms. The results above rely strongly on the monopolist's ability to coordinate product menu and pricing. Restricting the supply by means of strategic de-listing would clearly not be possible when products are introduced and priced by separate agents. Competition should lead to more variety as a consequence.

Furthermore, the monopolist is interested in eliciting specific search patterns, but he is indifferent regarding which product acts as starting point. In the analysis above one product was always selected as first in the consumer search query for consistency but without loss of generality. Conditional on a certain variant being the first one visited, however, the remaining available products do not generate the same expected profit. While this is irrelevant for a monopolist, competing sellers would likely try to gain prominence through undercutting strategies. If an equilibrium with competing sellers exists, it should then feature lower, uniform prices when consumers have the same prior considered here. These observation imply that competition would allow consumers to more freely explore the options made available in the market.

Finally, the paper has implications for environments in which order of search is selected by an agent different from the consumer. Inspecting different products leads to consumers learning about different attributes. This, in turn, leads to different search paths emerging. If consumers adapt their search strategy as suggested in this model, strategic selection of what consumers are allowed to sampled at the beginning of their search would lead to specific products being inspected, and others never being reached. If one agent selling a product variety could dictate what the consumer can see, as is the case for modern hybrid platforms, it seems clear that he would try to induce a favorable order of search. In line with recent work on consumption steering and self-preferencing, 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, Maxim Engers, and Daniel Savelle, "Ordered Search: Equilibrium and Optimum," *Working Paper*, 2020.

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

Arbatskaya, Maria, "Ordered search," *The RAND Journal of Economics*, 2007, *38* (1), 119–126.

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

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

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

Doval, Laura, "Whether or not to open Pandora's box," *Journal of Economic Theory*, 2018, *175*, 127–158.

Garcia, Daniel and Sandro Shelegia, "Consumer search with observational learning," *The 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.

Lancaster, Kelvin J, "A new approach to consumer theory," *Journal of political economy*, 1966, *74* (2), 132–157.

Nocke, Volker and Patrick Rey, "Consumer Search and Choice Overload," *Working Paper*, 2022.

Petrikaitė, Vaiva, "Consumer obfuscation by a multiproduct firm," *The RAND Journal of Economics*, 2018, *49* (1), 206–223.

Preuss, Marcel, "Search, Learning and Tracking," *Learning and Tracking (April 9, 2021)*, 2021.

Rhodes, Andrew, "Multiproduct retailing," *The 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.

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

Smolin, Alex, "Disclosure and pricing of attributes," *Available at SSRN 3318957*, 2020.

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

Wolinsky, Asher, "True monopolistic competition as a result of imperfect information," *The 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. Search

Lemma 1

Proof. It must be shown that, given realizations a_1, b_1 of the first search, the expected utility of inspecting (a_2, b_2) is as in lemma 1. The explicit form of the expected utility of such an inspection is:

$$
\begin{aligned}
&E[u(a_2,b_2)]|_{I=(\{a_1,b_1\})} = \\
&= a_1\left[\int_{-\infty}^{a_1-s} dF(y) + \int_{a_1-s}^{a_1+s}\int_{-\infty}^{b_1-s} dF(y)dF(y) + \frac{1}{2}\int_{a_1-s}^{a_1+s}\int_{b_1-s}^{b_1+s} dF(y)dF(y)\right] \\
&+ b_1\left[\int_{-\infty}^{b_1-s} dF(y) + \int_{-\infty}^{a_1-s}\int_{b_1-s}^{b_1+s} dF(y)dF(y) + \frac{1}{2}\int_{a_1-s}^{a_1+s}\int_{b_1-s}^{b_1+s} dF(y)dF(y)\right] \\
&+ \int_{a_1}^{+\infty} ydF(y)\left[\int_{a_1+s}^{\infty} dF(y) + \int_{a_1-s}^{a_1+s}\int_{b_1-s}^{\infty} dF(y)dF(y) + \frac{1}{2}\int_{a_1-s}^{a_1+s}\int_{b_1-s}^{b_1+s} dF(y)dF(y)\right] \\
&+ \int_{b_1}^{+\infty} ydF(y)\left[\int_{b_1+s}^{\infty} dF(y) + 2s\int_{a_1-s}^{\infty}\int_{b_1-s}^{b_1+s} dF(y)dF(y) + \frac{1}{2}\int_{a_1-s}^{a_1+s}\int_{b_1-s}^{b_1+s} dF(y)dF(y)\right] \\
&- [1 + F(a_1-s)[1-F(b_1+s)] + F(b_1-s)[1-F(a_1+s)]]s
\end{aligned}
$$

To show this is true, one can split the space of possible realizations (a, b) according to G based on two thresholds for each attribute. Specifically, the final product selected depends on the realization of (a_2, b_2) with respect to $a_1 \pm s$ and $b_1 \pm s$. These threshold segment the space in nine subset. When a realization falls in one of these subsets, the optimal product to keep can be identified.

If $a_2 < a_1 - s$, a_1 is kept for all realizations of b_2. Similarly, if $b_2 < b_1 - s$, b_1 is kept for all realizations of a_2. a_1 is also kept if $a_1 - s < a_2 < a_1 + s \wedge b_2 < b_1 - s$. The inverse is true for b_2. Finally, a_1 and b_1 are both kept when:

$$a_1 - s < a_2 < a_1 + s \wedge b_1 - s < b_2 < b_1 + s \wedge (a_1 + b_1) < (a_2 + b_2)$$

The latter condition splits the subset in which $a_1 - s < a_2 < a_1 + s \wedge b_1 - s < b_2 < b_1 + s$ in two.

Conditional on these conditions not being met, the remaining probability can be split similarly to determine when a_2 and b_2 would be kept. Their expected value can be directly computed conditional on their realization being higher than the respective variants. Overall, then one can associate each combination of possible outcomes of a_2 and b_2 to determine the conditional expected value of each. When aggregated, the resulting expecting value of searching (a_2, b_2) can be computed as above.

■

Lemma 2

Proof. It must be shown that if a product characterized by a previously inspected attribute is ever searched, it is never optimal to search a different variant of that attribute. I start with the two attributes case, and then generalize the argument.

The proof is by contradiction: suppose at any point of the sequential search process it is optimally selected to search a product characterized by a previously inspected attribute. Then, it must be the case that this product is better, in expectation, than a product sharing no attributes with any inspected products. It must also be the case that it is better than the best available realization.

For simplicity, refer to the inspected product with the highest realized utility as (a_i, b_i). The next product sharing no attributes with (a_i, b_i) that would be inspected would then be (a_{i+1}, b_{i+1})[12]. If instead (a_i, b_{i+1}) (resp. (a_{i+1}, b_i)) is searched next, it must hold:

$$a_i + b_i < a_i + b_i \int_{-\infty}^{b_i} dF(y) + \int_{b_i}^{\infty} ydF(y) - s$$

Moreover, since (a_i, b_{i+1}) (resp. (a_{i+1}, b_i)) is the optimal next search, it must be that:

$$E[u(a_i, b_{i+1})]_I > E[u(a_{i+1}, b_{i+1})]_I > m(a_{i+1}, b_{i+1})_I$$

Where $m(a_{i+1}, b_{i+1})_I$ is the myopic value associated with (a_{i+1}, b_{i+1}). Therefore:

$$a_i + b_i \int_{-\infty}^{b_i} dF(y) + \int_{b_i}^{\infty} ydF(y) < a_i \int_{-\infty}^{a_i} dF(y) + \int_{a_i}^{\infty} ydF(y) + b_i \int_{-\infty}^{b_i} dF(y) + \int_{b_i}^{\infty} ydF(y)$$

Suppose now that, after (a_i, b_{i+1}) is inspected, the consumer optimally selects to inspects (a_{i+1}, b_{i+1}). Since now b_{i+1} is known, this search is only feasible if:

$$a_i \int_{-\infty}^{a_i} dF(y) + \int_{a_i}^{\infty} ydF(y) + b_{i+1} - s > \max\{a_i + b_i, a_i + b_{i+1}\}$$

Both options lead to a contradiction: if $b_i > b_{i+1}$, this search being optimal would imply:

$$\left(a_i \int_{-\infty}^{a_i} dF(y) + \int_{a_i}^{\infty} ydF(y) + b_{i+1} - s\right) - a_i > b_i - b_{i+1} > 0$$

if $b_i < b_{i+1}$, instead:

$$\left(a_i \int_{-\infty}^{a_i} dF(y) + \int_{a_i}^{\infty} ydF(y) - s\right) - a_i > 0$$

Which contradicts (a_i, b_{i+1}) being optimally searched first.

The contradiction does not rely on products being characterized by two attributes. Each individual attribute, if kept after a search, must be strictly better than the expected realization of an unknown variant, minus the associated search costs. Since the argument applies to each attribute in isolation, it is robust to an environment with a higher number of attributes defining products. ■

[12]For clarity, I maintain that attributes are discovered in order based on their indices.

Extension: No Learning It can be readily shown that the process nests the standard Pandora search in environments in which there is no learning induced by shared attributes. Two different cases must be considered. First, suppose all products shared the same attribute a_1, and differed in their second attribute b. it is clear that all products are *ex ante* identical under the usual assumption of attributes being independently distributed. Moreover, since a_1 is shared between all products, each sequential search reduces to a myopic search process to find a variant b_i such that:

$$b_i > b_i \int_{-\infty}^{b_i} dF(y) + \int_{b_i}^{\infty} y dF(y) - s$$

which is equivalent to the myopic Pandora search as in Weitzman (1979).

The same is clearly true if no product share any attribute with each other: suppose that no product $(a_i, b_j), i \neq j$ was available. Then, inspecting a product can never lead to any correction since no product to correct to are available. Then, the value of searching a new product is equivalent to the myopic one: search continues until product (a_i, b_i) such that:

$$a_i + b_i > u_{i,i} \int_{-\infty}^{u_{i,i}-b_i} \int_{-\infty}^{u_{i,i}} dF(b)dF(a) + \int_{u_{i,i}-b_i}^{\infty} \int_{u_{i,i}}^{\infty} (a+b)dF(b)dF(a) - s$$

is found.

Extension: More Varieties It can be shown that the optimal search process is well defined when attributes come in a different number of varieties, as long as all combinations of said attributes are available. Consider again the optimal search following the first inspection. Once again, attributes are assumed to be discovered in order based on their indices. Suppose products are defined by two attributes:

$$a \in \{a_1, a_2, ..., a_n\} \qquad b \in \{b_1, b_2, ..., b_m\}$$

with $n < m$. Maintaining the assumption of identical, independent distributions across attributes, each search aimed at an unknown product can either aim at discovering two new attributes until product (a_n, b_n) is reached, or can at any point select a variant a_i or b_i to keep.

It is clear from Lemma 1 that once an attribute is kept while fully undiscovered products are available, the difference in the number of variants for the two attributes is irrelevant. In this case, products would be sampled myopically until one is deemed worth keeping, or all products sharing that attribute are discovered. Issues may arise if the first n searches all target fully undiscovered products, and none of them induces the consumer to stop (with or without correcting). Suppose this happens. After the n^{th} inspection, the consumer has learned *ex post* utility of n^2 products. She can choose to pick any of them (paying the search cost if it the best match requires a correction), or keep searching.

The unknown products, however, are all partially known: since no undiscovered attribute a is left, the $(m-n) \cdot m$ unknown products are only unknown in their attribute b. It follows that the best next product to search is the one with the highest realization along its a attribute. The myopic value arguments, then, apply.

Extension: More Attributes Increasing the number of attributes defining a product makes the calculation of the expected utility of products characterized by multiple uninspected attributes quickly get out of proportion. It is however, still feasible to find them to compare them. Consider an environment in which products are defined by three attributes:

$$a \in \{a_1, a_2\} \qquad b \in \{b_1, b_2\} \qquad c \in \{c_1, c_2\}$$

so that eight different products are available (assuming $A \equiv N$). As usual, suppose attributes are discovered in order: the first product inspected would then be (a_1, b_1, c_1). To determine what should searched next, the same general consideration of the main model apply.

First, it is clear that if after the first search two out of three attributes are to be kept, the expected utility of the product sharing them with (a_1, b_1, c_1) is equal to the myopic expected gain of discovering one attribute, which can be adapted from Equation 3. This allows to find reservation value of products (a_1, b_1, c_2), (a_1, b_2, c_1), and (a_2, b_1, c_1).

Suppose now that the consumer considers to keep only one attribute, say a_1 without loss of generality. Then, the only product that can be inspected next is (a_1, b_2, c_2), from which the consumer can correct to either (a_1, b_1, c_2) or (a_1, b_2, c_1), depending on the outcome of the second search. Since a_1 is known, the value of searching (a_1, b_2, c_2) is equivalent to the value of searching an hypothetical product (b_2, c_2) after (b_1, c_1) if products were only defined by $b \in \{b_1, b_2\}$ and $c \in \{c_1, c_2\}$. The expected utility of inspecting this product, then, can be constructed following Equation 2. Equivalent considerations lead to the expected utility of inspecting products (a_2, b_1, c_2) and (a_2, b_2, c_1).

Finally, if the consumer chooses to keep no attributes, he would search (a_2, b_2, c_2) next. From here, the consumer can correct to any of the uninspected products since she now knows the realization of all individual attributes. Once again, the expected utility associated with this inspection can be constructed following Equation 2. All these values can then be compared given realizations a_1, b_1, c_1 to find the optimal follow-up search. Going backwards, this in turn allows to find reservation price of all products before the search process starts.

Extension: Missing Combinations of Attributes - Binomial example When not all combinations of attributes are available to inspect, which product to inspect first is not without loss of generality. To fix ideas: consider the simplified environment in which attributes' realizations are extracted from a binomial distribution $B(1, \alpha)$. Suppose that:

$$A = \{(a_1, b_1), (a_1, b_2), (a_2, b_1)\}$$

Further suppose that $2\alpha > s > \alpha$: after a product such that $u(a, b) = 1$, searching again is not worth the cost of inspection. All products are, however, myopically worth inspecting.

It is straightforward to show that searching (a_1, b_1) first is strictly worse than any of the other two. Indeed, all three have the same probability of generating ex post utility $u(a, b) > 0$, which would make the consumer stop searching. Suppose however $u(a_1, b_1) = 0$. Expected utility of the remaining products updates to:

$$E[u(a_1, b_2)]|_I = E[u(a_2, b_1)]|_I = \alpha - s < 0$$

Both products share an attribute that is known to be disliked by the consumer. In this case, search would stop as well. Suppose now the consumer inspected (a_1, b_2) (resp. (a_2, b_1)). If $u(a_1, b_2) = 0$ (resp. $u(a_2, b_1) = 0$), expected utility of the remaining products updates to:

$$E[u(a_2, b_1)]|_I = 2\alpha - s > 0 \qquad (\text{resp.} E[u(a_1, b_2)]|_I = \alpha - s > 0)$$

$$E[u(a_1, b_1)]|_I = \alpha - s < 0$$

Going backwards, then:

$$E[u(a_1, b_2)]|_{I\equiv\emptyset} = E[u(a_2, b_1)]|_{I\equiv\emptyset} = 2\alpha - s + (1-\alpha)^2(2\alpha - s)$$

$$E[u(a-1, b_1)]|_{I\equiv\emptyset} = 2\alpha - s$$

Similar arguments can be constructed for $\alpha > s$: different outcomes lead to different search paths arising, which imply different expected utility values relevant for the selection of the first product to inspect. The same logic can be applied to larger, more complex search problems. The logic described in Proposition 1 is intact, but a proper mapping of the optimal search when not all combinations are represented is necessarily problem specific.

B. Monopoly pricing

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

The monopolist is interested in finding prices that maximize probability of trade times price. Given expected utility of search as per Equation 5:

$$\begin{aligned} &E[u(a_1, b_1)]|_{I\equiv\emptyset} = \alpha^2 \max(2 - p^u, 0) - s \\ &+ 2\alpha(1-\alpha)\max(1 - p^u, (1-\alpha)\max(1-p^u, 0) + \alpha\max(2-p^u, 0) - s, 0) \\ &+ (1-\alpha)^2 \max(\alpha^2 \max(2-p^u, 0) + 2\alpha(1-\alpha)\max(1-p^u, 0) - s, 0) \end{aligned}$$

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

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

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_L^E is feasible among the candidates above. Furthermore, it can be shown that:

$$\alpha \leq s < 2\alpha \;\rightarrow\; p_L^E < 1$$

By plugging in p_L^E in equation 5, one sees that at this prices the consumer stops and purchase if $u(a,b) \neq 0$, and is willing to search again if $u(a_1, b_1) = 0$. It is clear that no deviation from p_L 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, $p^{u*} = p_L$.

If $\frac{3\alpha^2-2\alpha^3}{1+2\alpha-2\alpha^2} \leq s < \alpha$, both p_I^D and p_L^E are feasible. Moreover, it holds $p_L < p_I < p_T$ for the whole segment. Therefore, it is sufficient to compare expected profits under p_L^E and p_I^D. Notice that p_I^D is such that searching again after a bad first realization is not possible. In this segment:

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

Therefore, the monopolist compares:

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

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

Direct comparison indicates that p_I^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, $p_L < p_T = 1 \iff \alpha^2 < s < \alpha$. 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 $0 < s < \alpha^2$, $p_L > p_T$; 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 monopolist 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 monopolist can introduce differential prices profitably. In particular, consider pricing such that:

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

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 (a_1, b_1) (equivalently, (a_2, b_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 since the she purchases the cheaper product with higher probability if she starts from it.

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 (a_1, b_2): inspecting (a_2, b_2) leads to utility equal to two with probability α^2, and allows to correct to (a_1, b_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 (a_1, b_1) directly if $s > \alpha^2$. Therefore, no deviation is possible in this segment.

If $s < \alpha^2$, two changes must be accounted for. First, $p_T < p_L$ is the preferred option, since $p_L > 1$ does not lead to trade taking place. Further, the consumer would want to search the cheaper (a_2, b_2) first, since search costs are low. The monopolist can react by:

- letting the consumer do so, increase the price of (a_1, b_2) to $p_T + 1 - s$

- reducing the price (a_1, b_2) to induce his preferred order of search
- removing (a_2, b_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. Since this is the case, its price can be increased, since 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 (a_1, b_2), the monopolist makes sure that the consumer has no incentive to deviate. Since $s < \alpha^2$, 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 (a_2, b_2) prevents the deviation from taking place at all. Since no follow-up search in case of a bad first realization is possible without (a_2, b_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 monopolist can choose an intermediate strategy between uniform prices at p_H^E and uniform prices at p_H^D. Suppose the consumer inspected (a_1, b_1) and learned $a_1 = 1$, $b_1 = 0$. Then, she would want to inspect (a_1, b_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 (a_2, b_2) instead of (a_1, b_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 monopolist 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 (a_2, b_2).

Unlike in the previous case, the latter option is always optimal. When the monopolists selects differentiated prices, then, for α high and s low the consumer has an incentive to adapt in a way that makes the monopolist 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 $p^{u*} = \min\{p_L, p_T\}$, 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 monopolist can optimally select, direct comparison of the profit is sufficient. The same applies to the case in which $p^{u*} = p_I$ and $0 < s < \alpha^2$.

Two results emerge: when selecting p_L as base product and the consumer does not adapt their search strategy, this is always optimal. Second, when there is adaptation by consumer and monopolist, 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 monopolist 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 blue area in graph 2 in the main text.