

Consumer Search and Firm Strategy with Multi-Attribute Products*

Jacopo Gambato[†]

August 15, 2023

Abstract

I study a model of directed search where a consumer inspects products sharing attributes. In this framework, the order in which available products are inspected changes depending on past realizations, making the optimal search process history dependent. I embed said search process in an environment in which a multiproduct monopolist selects the menu of products offered and their prices. Consumers discover their preferences over 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 and profit off the information learned by the consumer through search. 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. These findings carry implications for recommendation systems and price discrimination based on consumers' search history.

Keywords: consumer search, directed search, learning, multiproduct monopoly, revealed preferences

JEL Codes: D42, D83, L12, L15

*I would like to thank the economics department at University of Mannheim and the IOEK department at ZEW Mannheim for guidance, and in particular Martin Peitz, Bernhard Ganglmair, Nicolas Schutz, Anton Sobolev, and Volker Nocke; I also thank Daniel Savelle, Axel Gautier, José Moraga-Gonzales, Justus Preusser, Sander Onderstal, Marco Haan, Julia Reimer, Heiko Karle, David Ronayne, Willy Lefez, as well as the participants to the 2023 MaCCI Annual Conference, 11th CRC Retreat, 2023 CLEEN Conference, 12th Workshop on Consumer Search and Switching Costs, 2023 CRESSE Conference, and 21st ZEW ICT Conference for their insightful comments.

[†]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.

To study the implications of learning in directed search when products are correlated through shared attributes, I embed the framework in an environment in which a multiproduct monopolist selects product menu and pricing. Multiproduct firms are widespread and, accordingly, have been object of thorough study. While much has been written about multiproduct firms' pricing strategies¹, the role of their menu composition has received significantly less attention, and the strategic interaction between the two even less so. In this paper, I present a framework that allows to capture the synergy between two similar products borne of consumers' learning of their own preferences, and how a firm that can coordinate the menu made available can profit off it.

The consumer in my framework observes, at the beginning of the game, only which products are made available and how these products are related, but is unaware of her preferences for the attributes that characterize them. Prices are posted, and contribute to determining the order in which consumer search for their preferred product². 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 the preferences they learn about through the search process.

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

¹Notable, among many others, are the early contributions by [Mussa and Rosen \(1978\)](#) and [McAfee et al. \(1989\)](#).

²Posted prices also resolve the standard hold-up problem of hidden prices in models of search as understood by [Anderson and Renault \(2006\)](#); A more in-depth analysis on the matter can be found in [Konishi and Sandfort \(2002\)](#) and [Robert and Stahl \(1993\)](#).

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 multiproduct 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.

Differential prices used to encourage search can arise, however, only if the product menu is relatively small. When the available products are few, a negative outcome of early inspections reduces the space of products from which the consumer can expect to obtain high utility. The value of inspecting new products decreases in the number of past inspections if their outcome is particularly bad, and more so the lower is the number of subsequent possible inspections (that is, the smaller is the menu of products made available). I show that when the product space is extended to encompass infinite varieties, setting differential prices can never be the optimal strategy for the monopolist. Differential prices make cheaper products prominent, which reduces the probability of selling more expensive alternatives. While with a finite menu this reduction in expected profits is compensated by a higher overall probability of trade, the same is not true for an infinitely large menu. The result suggests that the optimal pricing strategy of a multiproduct monopolist depends not only on his menu composition, but also on the menu's size.

I further 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 attribute³. 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 unique⁴. 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

³For example, [Hodgson and Lewis \(2020\)](#) shows evidence of "spatial learning" in search: consumers are shown to inspect more differentiated products early and to get closer to the eventually purchased option as search progresses.

⁴This is without loss of generality: if color is the only relevant attribute, two shirts sharing the same color are effectively the same product

of contributions, however, have instead focused on ordered search. In some cases⁵, the order in which the consumer searches is fixed exogenously; in others⁶ 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. Additionally, they assume that products have two attributes, one directly observable and one that needs to be discovered through search.

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 account 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 products available. Inspection of a product informs the consumer of her taste for other options. Inspecting multiple products not sharing any attribute indirectly informs the consumer of her preference for products that share attributes with them. A tension arises between the informational value of inspecting unrelated products, and the efficiency implied by keeping good realizations. I show that in this environment “backtracking” to a previously inspected and abandoned attribute can be optimal.

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. After each inspection, the value of searching further changes to account for the information learned, which in turn determines the possible search paths a consumer would take from that point onward. I show that the multi-attribute structure allows to apply a version of [Weitzman \(1979\)](#)’s optimal search in an environment with correlated products.

The dynamic studied in this paper has subtle implications for recommendation systems in environments in which an agent has control over the menu offered to consumers, and price discrimination based on consumers’ search history. The monopolist I study has no information about the consumer preferences but can induce self-selection based on taste through strategic pricing. If he did, however, an implication of the results is that he would optimally choose not to recommend the best match he could. Instead, he would have an incentive to recommend a subpar option first, and let the consumer find a better, more expensive match using the information acquired from the first inspection. Relatedly, if the seller could condition prices on search history, the same pricing outcome would emerge: consumers who search unrelated products signal that they obtained a bad realization on the first search, while consumers who search related products signal a positive early realization. It is clear then that the seller would want to react to the latter search path being taken by raising prices.

The rest of the paper is structured as follows: after reviewing the related literature, I present the

⁵[Arbatskaya \(2007\)](#), [Zhou \(2011\)](#)

⁶[Choi et al. \(2018\)](#), [Haan et al. \(2018\)](#), [Anderson et al. \(2020\)](#)

framework (Section 2) and characterize the optimal search process with multiple attributes and the learning process they imply in a simplified version of the model (Section 3). Afterwards, I solve the problem of a monopolist that selects which products to make available and their price (Section 4). I present a general version of the search model and provide the equilibrium pricing for the infinite products case in Section 5. After exploring extensions to the general model in Section 6, I conclude in Section 7.

Related literature This paper relates to several strands of literature. First, it contributes to the wide literature on multiproduct firms. Earlier contributions addressed several aspects of the possible strategies of this kind of firms and the products they offer. Some, like [Mussa and Rosen \(1978\)](#), focused on price discrimination with vertically differentiated products. Others, like [Eaton and Lipsey \(1979\)](#), discuss market pre-emption through introduction of horizontally differentiated options. Other notable examples relate to R&D expenditure ([Lin, 2004](#), [Lambertini and Mantovani, 2009](#)) and bundling of products ([McAfee et al., 1989](#)).

Fewer are the papers that address the strategic component inherent in selection of the menu composition. [Johnson and Myatt \(2002\)](#) study menu extension and pruning in response to entry, and focus on differentiated products; menu selection is, therefore, adaptive to changing market conditions. More closely related is work by [Brander and Eaton \(1984\)](#), which considers product introduction of more or less differentiated products. While the authors focus mostly on the duopoly setting, their multiproduct monopolist values minimizing substitutability of the products offered to avoid cannibalization. I show that, with learning and search frictions, this generally is not true: the monopolist I study wants to make available products that share attributes (and are, therefore, closer substitutes compared to products that don't) to encourage the consumer to search at higher prices.

An important result of this paper is the endogenous arising of differential prices for *ex ante* identical products. This result is in conflict with that of [Amir et al. \(2016\)](#): the author finds that a multiproduct monopolist facing linear demand always price his products independently. This result holds even if products are negatively correlated; I show here that the same is not true when products are positively correlated through shared attributes because the monopolist has the incentive to use cheap products to steer consumers towards more expensive ones. To the best of my knowledge, this paper is the first to study the implications of menu selection for consumer steering and, therefore, relates to this rapidly growing literature as well⁷.

The paper also 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 value, 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 stops 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. When

⁷In the steering literature, this paper is closest to [Ichihashi \(2020\)](#); other notable contributions, although less closely related, are [De Corniere and Taylor \(2019\)](#), [Teh and Wright \(2022\)](#), and [Heidhues et al. \(2023\)](#).

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 optimal path each realization would lead to. These paths in turn depend 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. Following the same intuition detailed in [Anderson et al. \(2021\)](#), I propose a tractable, history-dependent scoring system that incorporates the value of searching beyond the target of inspection. The score is determined only by the value of the paths that would be optimally taken after the realization they refer to and, therefore, reflect the full value of inspecting new attributes. Through this scoring system, I show that a dynamic version of [Weitzman \(1979\)](#)'s optimal search policy can be characterized in this environment with correlated products.

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. In a different spirit, [Ke and Lin \(2022\)](#) and [Bao et al. \(2022\)](#) study optimal search in a simple framework in which a discrete number of products share one of their two attributes. [Ke and Lin \(2022\)](#) returns conditions under which correlation is search leads to complementarity of the products available; [Bao et al. \(2022\)](#), instead, studies Bayesian updating when the consumer cannot distinguish the role of each attribute in the *ex post* utility each product grants.

More closely related to this paper is the aforementioned 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-routes 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, and indeed the mechanical intuition behind the two learning process is similar. 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, who always 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 multiproduct 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. Moreover, adaptation in search comes from endogenous learning through inspection of products and not as a separate action; this distinction separates the paper from earlier works, such as [Garcia and Shelegia \(2018\)](#) and [Greminger \(2022\)](#).

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. A notable exception is work by [Anderson and Renault \(2006\)](#), which studies monopoly pricing with search frictions. The authors focus on information disclosure through advertisement of price, product attributes, or both. Like in my setting, price advertisement is shown to resolve hold-up problems generated by search frictions. The paper also provides theoretical evidence of the limited interest of a firm to disclose too much information about his product's characteristic, since it would dissuade some valuable costumers from visiting. In the same spirit, [Mayzlin and Shin \(2011\)](#) shows that firms selling high quality products might choose to withheld information about it to encourage search. My approach is closer to the latter: differential prices arise to more efficiently capitalize on consumer taste, making the two strategies substitutes.

Unlike both, however, I study pricing as implemented by a multiproduct monopolist who can

coordinate prices of horizontally differentiated products. In most papers with the same premise⁸, consumers visiting a multiproduct seller 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 multiproduct 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 valuable 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 multiproduct 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. Framework

The products. I consider an industry with products differentiated with respect to two attributes⁹. A product (i, j) is identified by attributes $A_i \in A, B_j \in B$. In the simplified framework below, I assume A and B to come in two variants each; in the general model, I assume A and B to come in infinite variants¹⁰, so that the number of products available for purchase is infinite as well. Each attribute A_i can be found combined with all attributes $B_j, i, j \in \{1, 2, \dots\}$, and *vice versa*. One can visualize the products as displayed in a grid, with the rows representing the A attributes, the columns representing the B attributes, and the cells representing products defined by a specific combination of A and B attributes as depicted in [Figure 1](#). Notice that products are only differentiated through their attributes and are otherwise identical in quality¹¹.

The consumer. A representative, risk-neutral consumer (she) has unit demand, is aware of the available products and their attribute composition, and can inspect the products in any order she likes. The consumer has no prior knowledge of her preferences for the available attributes; she learns the realization of each attribute separately by inspecting a product characterized by

⁸[Zhou \(2014\)](#); [Rhodes \(2015\)](#); [Rhodes et al. \(2021\)](#)

⁹The framework is adapted from [Smolin \(2020\)](#)

¹⁰Relaxing this assumption is the topic of one of the extension.

¹¹As an example: shirts differ in their color (attribute A) and their fabric (attribute B).

it. In line with existing models¹² I assume that *ex post* utility generated by a generic product (i, j) takes the form:

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

where $u_{i,j}$ represents the realized *ex post* utility. I assume each attribute to be an i.i.d random variable distributed according to a cumulative distribution function F . Defining a generic attribute as $y \in A \cup B$, I further assume $F(y)$ to have support $[0, \hat{y}]$ for some positive \hat{y} , and to be twice-differentiable everywhere on it. The assumption that attributes enter $u_{i,j}$ additively crucially implies that there are no “synergies” between attributes: once an attribute is discovered, its realized value affects all products that are defined by it in the same way.

In the simplified framework below, I further 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 ¹³. Expected utility of an unsampled product (i, j) is then:

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

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

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

In this environment, I study the optimal sequential search process with free recall: a 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 both distribution F and search costs s , and learns the value of each attribute separately after each inspection. Finally, the consumer’s outside option is normalized to $u_0 = 0$.

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

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

1. Consumer and monopolist observe distribution F and search costs s ,

¹²For example: Choi et al. (2018) and Greminger (2022).

¹³In this specification, one can think of the search process as an effort to find an attribute that satisfies a need of the consumer.




Figure 1: *Graphical representation of products (i, j) . Products in the same row (resp. column) share attribute A_i (resp. B_j).*

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

For the pricing game, I consider Sub-game Perfect Equilibria¹⁴; 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, I assume that indifference is always resolved in favor of the monopolist – that is, the most profitable outcome is selected.

3. A Simple Model of Multi-Attribute Search

To fix ideas, I first characterize search in the simplified framework. The exercise highlights the role of shared attributes in determining the optimal search policy, and provides intuitions useful for the general model.

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

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

Further, assume attributes to be i.i.d random variables distributed according to $B(1, \alpha)$, where $\alpha \in (0, 1)$ refers to the probability that the consumer likes any inspected attribute y . In this specification, the cost of inspecting a product is $s \in (0, 2\alpha]$. Both distributional parameter and s are known. To illustrate the search dynamic in isolation, I start with the assumption that prices are exogenously set at zero; this assumption will be relaxed in the next section. The consumer can inspect any product in \tilde{N} ; for clarity, I start from the simple case in which $\tilde{N} \equiv N$. At any given point of the search sequence, the set of available products can be partitioned in the set of inspected products, I , and uninspected products, $\tilde{N} \setminus I$.

Updating expected utilities. Suppose the consumer already inspected one of the products. Since all products are *ex ante* identical, inspecting $(1, 1)$ first is without loss of generality¹⁵. More in general: whenever which product to inspect can be chosen randomly without loss of generality, I assume that products are inspected in increasing order of 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.

¹⁴Since the monopolist commits to menu and prices before the search process starts, and since prices are posted, there is no need to model beliefs explicitly in this environment.

¹⁵All products share each attribute that characterizes it with another product, and for all products there is one other product that shares no attributes with it. Therefore, all product are *ex ante* identical as long as prices are uniform or, in this case, abstracting from prices

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

$$E[u_{1,2}]|_{I=\{(1,1)\}} = A_1 + \alpha, \quad E[u_{2,1}]|_{I=\{(1,1)\}} = \alpha + B_1,$$

$$E[u_{2,2}]|_{I=\{(1,1)\}} = 2\alpha.$$

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

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

Which immediately returns the optimal follow-up search for each possible realization of $(1, 1)$:

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

Expected utility of searching. In this environment, different realizations lead to different search paths being taken every time a new product is inspected. These conditional search paths emerge predictably, and all realizations generate unambiguously an optimal path forward. In turn, this implies that a rational consumer would account for the likelihood of these different paths emerging, and the expected utility they are associated with, when deciding whether to start searching or not. From the above, then, we obtain the expected utility of searching given the available products and the optimal search paths that can emerge:

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

The first addendum refers to (i, j) being the best possible match ($u_{i,j} = 2$, with probability α^2). The second addendum refers to the eventuality of the consumer liking only one of the two

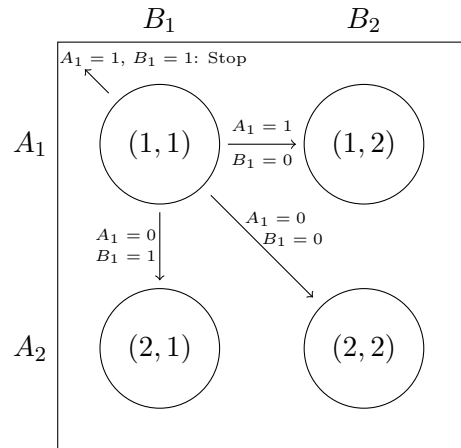


Figure 2: *Optimal search with binomial distribution and all products available, starting from $(1, 1)$.*

attributes, and the possible second search that outcome would entail, which only takes place if $s \leq \alpha$. The third refers to the case in which $u_{i,j} = 0$ so that the product sharing no attributes with it would be inspected next. When all products are available and uniformly priced (and therefore *ex ante* identical), inspecting products in increasing order of their index is without loss of generality. Figure 2 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 will highlight the effect of restricting the supply for the optimal order of search only referring to the specific cases and ways in which a multiproduct monopolist chooses to restrict the product menu.

4. Seller's Optimal Strategy

I embed the simple framework above in an environment in which a multiproduct monopolist selects which products to make available to the representative consumer (that is, he selects $\tilde{N} \subseteq N$), and their respective prices. I assume prices to be set before the search process starts, cannot be changed, and are observed costlessly by the consumer before she starts searching. I further assume distributional parameters and search costs to be known by the firm. Finally, I assume all production costs to be constant and set to zero. I show that in this environment the learning component allows the monopolist to optimally set different prices for *ex ante* identical products to profit off the learning process inherent to multi-attribute search.

4.1. The firm's problem

The monopolist's problem is twofold: he must set up prices to maximize profit, and must select \tilde{N} to generate trade opportunities. The two decisions are related. The consumer search path depends on the price she observes, and which prices would deter her from searching depend on the available products. In particular, the consumer is willing to search a product priced above its myopic expected value $2\alpha - s$ as long as the expected utility of searching from that point onward is non-negative. As shown above, this can be achieved when products that share attributes with each other are made available. A 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 – \tilde{N} and $\mathbf{p}(\tilde{N})$, where the latter represents the vector of prices associated with all products in \tilde{N} – interact in non obvious ways. Uniform prices, for example, cannot induce an order of search different from the one characterized above. If these uniform prices are too high, however, some search paths could end prematurely: even if products are identical *ex ante*, the first one searched ((1, 1) in the example above) carries more new information than every subsequent search that could arise. It follows that the highest price that would make two products not sharing attributes worth searching is different. If prices are not uniform, however, the consumer could adapt 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.

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: $\tilde{N} \equiv N$. I then show the monopolist's incentive to restrict the menu and the effect this choice has on consumer search. Once again, it is assumed throughout that $s < 2\alpha$.

Consider a generic uniform 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 section above, the expected utility of performing the first inspection is:

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

that is: the value of inspecting $(1, 1)$ ¹⁶ is equal to the expected value generated by the search paths that are induced by the possible different realizations. These in turn depend on the relative value of s and α , over which the 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. The first search path to be prevented by high prices is the one that arises conditional on a bad first match. Indeed, given observation $u_{1,1} = 0$, $(2, 2)$ is searched as long as:

$$E[u_{2,2}]|_{u_{1,1}=0} = \alpha^2 \max\{2 - p^u, 0\} + 2\alpha(1 - \alpha) \max\{1 - p^u, 0\} - s \geq 0 \quad (2)$$

It is straightforward to show that there exists values p^u such that this condition is not satisfied but $E[u_{1,1}]|_{I=\emptyset}$ is positive: even if the consumer would not search after a bad realization of $(1, 1)$, the presence of products sharing attributes with it makes it more likely to find something worth purchasing. As long p^u is such that $E[u_{1,1}]|_{I=\emptyset}$ is non negative, the consumer can rationally start inspecting products. With this inspection, the consumer can discover that she likes both attributes, after which she always stop searching 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

¹⁶Once again, searching in increasing order of the indices is without loss of generality when all products are available and prices are uniform

additional search if and only if:

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

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

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

where the apex M stands for “myopic”, is the only feasible price when $s > \alpha$ ¹⁷.

Suppose now that $s \leq \alpha$. 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 2 with equality:

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

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

Alternatively, the monopolist can select higher prices that discourage search after a bad first realization. These prices must be strictly 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 $E[u_{1,1}]|_{I=\emptyset} = 0$, since for any higher price the consumer would not start searching:

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

where D stands for “discourage”, and L and H stand for “low” 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 $p \leq 1$: any higher price requires the consumer to like both attributes in a product to purchase it. Notice that it holds:

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

Therefore, the price that maximizes search and trade can be identified as the minimum between

¹⁷As I show in the next section, that correlation does not affect search when search costs are high enough holds generally.

p_L^E and 1. To simplify the notation, I define:

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

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

By plugging in the various (feasible) prices for the various combinations of α and s and following the search path different prices induce according to Equation 1, one can obtain the expected profit of the 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, when $s > \alpha$, expected profit is always:

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

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

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

which maximizes probability of trade and is always valid,

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

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

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

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

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

which does not allow for another search after a bad first realization. In all cases, expected profit is calculated as price times the probability of trade generated since production costs are assumed to be equal to zero.

The candidate prices reflect the relative importance of encouraging search and extracting rent conditional on search taking place. p_T 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_L^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 search costs and high probability of a match α . For such parameters the consumer is easily encouraged to start searching. If $s > \alpha$, there is only one candidate price, p^M , the lowest of the candidate prices. For $s \leq \alpha$, instead, which of the four candidate prices is selected depends on the relative value of s : the lower s is, the more aggressive 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 (1,1) 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 (1,1), inspection of (2,2) could not rationally take place. Indeed, (2,2) would only be inspected after a bad first realization, but $p^{u*} \in \mathbf{p}^D$ prevents this search altogether. 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 (2,2) were to be removed, search starting from (1,1) would be unaffected. Starting from any other product, however, would generate negative expected utility of search. Suppose for example that (2,2) was removed and that the consumer started from (1,2) or (2,1). Then, not only she would not rationally inspect the unrelated product, but she would not be able to inspect (2,2) after learning something positive about it. This cannot be optimal.

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

Proposition 1. *Consider a multiproduct monopolist selecting optimal menu $\tilde{N} \subseteq N$ 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 $s > \alpha$: $p^{u*} = p^M$, $\tilde{N} \subset N$, and the consumer can start searching from any available product,*
- *If $s < \alpha$ and $p^{u*} \in \mathbf{p}^E$: $\tilde{N} \equiv N$, and the consumer can start searching from any available product,*
- *If $s < \alpha$ and $p^{u*} \in \mathbf{p}^D$: $\tilde{N} \subset N$, and the consumer is steered toward a specific search path.*

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

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 hinder the probability of trade is often optimal. Raising prices is only worth it if the loss of a potential trade is compensated when trade does take place. In particular, α must be high enough that the chances of not liking the first product inspected are low, and s must be low enough that search is not discouraged. Whenever this is the case, the monopolist can raise price and choose not to introduce all possible variants; 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. Strategic menu selection can give rise to endogenous prominence based on the relative position of the products.

At uniform prices, the consumer retains some positive expected value from search when the 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 whenever fine tuning after a good, but not great, match is possible.

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

If the consumer, instead, learns that she likes an attribute inspected in the first search, she would like to search next along that attribute. This is clearly true if prices are uniform. Suppose however that $(1, 2)$ and $(2, 1)$ were priced slightly higher than $(1, 1)$. If the consumer has learned that she likes A_1 (resp. B_1), and if the price difference is not too high, she would still want to search the more expensive product. Going backwards: the consumer would start her search from the cheaper option given that products are *ex ante* identical. As long as the price differential is not too high, the consumer has no incentives to stop searching early, nor to deviate towards a different search path. By pricing $(1, 1)$ and $(2, 2)$ at $p = p_T$, and the remaining products at a

higher price, 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. When considering the equilibrium strategy of the monopolist, the following result emerges:

Proposition 2. *Consider a multiproduct monopolist selecting optimal menu $\tilde{N} \subseteq N$ and pricing $\mathbf{p}(A)$ of multi-attribute products. In equilibrium:*

- *Encouraging uniform prices are set for high search costs,*
- *Discouraging uniform prices are set for low search costs and high probability of a match*
- *All products are introduced if and only if prices are not set uniformly.*

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

Determining the optimal pricing vector with differential prices is challenging in this environment. In particular, the difference in prices can induce the consumer to adapt their search strategy to avoid the more expensive product and retain some expected utility. We are interested in finding out the optimal price spread from the 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 (1, 1) and (2, 2) have lower prices and therefore act as possible starting points; further, I keep the assumption of products over which the consumer is indifferent to be searched in increasing order of their indices.

First, consider the optimal price spread. The search rules determine two separate constraints. Prices must be such that search can start. Moreover, prices must be consistent with the search process as it unfolds. As the price increase being profitable 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 (1, 1) and observes $A_1 = 1, B_1 = 0$. Suppose moreover that the optimal base price selected by the monopolist is $p_T \leq 1$. Conditional on inspecting one attribute being worth the cost of inspection ($s < \alpha$), the consumer would want to search (1, 2) if:

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

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

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

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

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

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

Given search as characterized above, these pricing structure lead to the same probability of trade as their uniform counterparts. Compared to them, however, they lead to higher expected profit 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 that she likes something about it. In this case, the consumer would search (2, 2) hoping to find a good realization instead, and would only inspect the more expensive product if she knows she likes both of its attributes and nothing else. Consider again the candidate prices profile $p_{1,1} = p_{2,2} = p_T$, $p_{1,2} = p_{2,1} = p_T + \delta_L$. After realization $A_1 = 1$, $B_1 = 0$:

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

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

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

For α high enough and s low enough, inspecting (2, 2) before (1, 2) is a rational deviation: search in this case is cheap, and the likelihood of liking both attributes A_2 and B_2 is relatively high. This deviation is at the detriment of the 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 (2, 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 (2, 2) to induce his preferred order of search and keep the same prices for all other products.

The first reply further highlights the ability of the monopolist to condition prices on search behavior. If the consumer has an incentive to search (2, 2) after (1, 1) conditional on $A_1 + B_1 = 1$, 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 (1, 2) and (2, 1) cheaper. Since the consumer is interested in searching (2, 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 (2, 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 (2, 2) when the monopolist would want her to inspect (1, 2) or (2, 1).

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

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

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

- 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 (2, 2) to prevent the deviation deviation,

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

All three options are optimal for some combinations of α and s . The same exercise can be applied to the alternative pricing profile $p_{1,1} = p_{2,2} = p_H^E > 1$, $p_{1,2} = p_{2,1} = 2 - \frac{s}{\alpha}$: 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 (2, 2) always dominates 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 previous subsection. 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

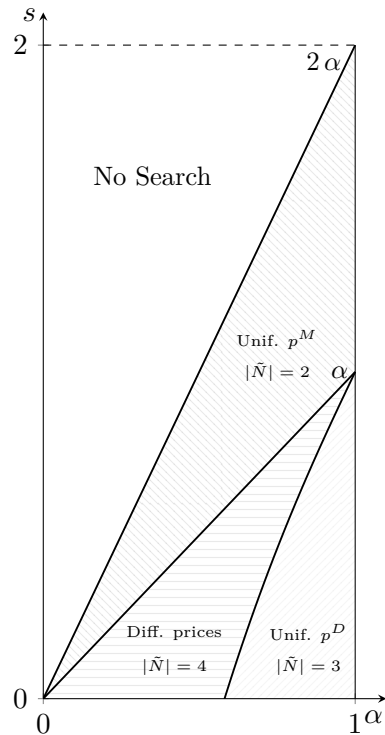


Figure 3: *Equilibrium monopoly menu selection and pricing for all feasible combinations of $\alpha \in (0, 1)$ and search costs $s \in (0, 2\alpha)$*

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.

Figure 3 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 for 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 patterns 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.

4.2. Discussion

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

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

A possible application of this framework relates to the practice of businesses to offer free samples of new products to attract interest¹⁸. In particular, by making some products prominent and easy to assess, a firm can encourage potential buyers to learn about their taste for novelties and alternatives that they might not consider otherwise. In doing so, the firm can use the positive experience associated with the sampling to increase the willingness to pay of consumers unaware of their preferences for said products. Together with the strategic ordering shown above, this

¹⁸Thanks to David Ronayne for pointing out this interaction.

points at the importance of menu selection and positioning of options in environments with search frictions.

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

Consumers self-selecting based on taste also creates the incentive to condition pricing on their search history¹⁹. Algorithmic pricing, the practice of pricing items automatically to adapt to the state of the market, are more and more commonly used in the digital world²⁰. The model highlights the role product's position in the attribute space plays in their pricing. Consumers are willing to search more expensive products only if they have already learned something positive about them by inspecting a different option. Equivalently, one can imagine a reactive pricing system that adapts as search unfolds. If two products sharing attributes are inspected in sequence, the ordering signals that the consumer has learned something positive about those attributes. The price of the second, then, can be safely raised by an algorithm trained to recognize these patterns. On the other hand, if two products not sharing any feature are inspected in sequence, both should be priced low to maintain the consumer engaged with the search. This interaction between spatial learning in search, evidence of which is presented in [Hodgson and Lewis \(2020\)](#), and algorithmic pricing suggests the need for further, more detailed investigation.

5. Optimal search with multi-attribute products

The simplified framework analyzed above hints at the mechanics underlying optimal search in this environment. A consumer can rationally inspect different products after different realizations, and the value of inspecting two products depends on the order in which they are inspected even if they are *ex ante* identical. Further, the value of inspecting a product depends on the other products that share attributes with it, since what is learned through inspection affects these products' expected utility.

I now generalize this intuition. More specifically, I obtain rules of search that apply generally in this environment. In what follows, I rethink the problem in a way that allows to apply [Weitzman \(1979\)](#)'s logic and therefore reduce the search problem to a set of rules reminiscent of Pandora's optimal search policy. [Weitzman \(1979\)](#)'s result is not immediately applicable to this

¹⁹Thanks to Willy Lefez for pointing out this interaction.

²⁰Airline companies and, more recently, e-commerce retailers are prime examples of this practice being widely in use.

environment due to the correlation between products: since products share attributes, it is not possible to assign to each product a score that only depends on the product itself. I show how this can be achieved by building “nests” of product to be scored as a single “box”²¹, and letting their score update following certain realizations to account for changes in the optimal search that would follow.

A similar environment is studied in [Bao et al. \(2022\)](#): in their framework, products share one attribute and are independent across the other; however, the authors study Bayesian updating when the two attributes cannot be observed separately. [Ke and Lin \(2022\)](#) also study a problem with products sharing one of their attributes. They characterize search with two correlated products and obtain conditions for complementarity between these products to arise. Instead, I propose a more general approach: by building fictitious boxes including all correlated products and tracing optimal follow-up inspections inside these boxes, I characterize search with correlated products in a way that preserves the intuition behind Pandora’s optimal search policy.

5.1. Towards a general framework

First steps Consider first a simpler case, illustrated in [Figure 4](#). The two products available share one attribute (A_1) and are independent along the other ($B_j, j \in \{1, 2\}$).

Suppose that the consumer already inspected $(1, 1)$: she has learned her valuation for A_1 , shared by both products, and B_1 . She still does not know her valuation for B_2 . At this stage, it is clear that choosing between stopping at $(1, 1)$ and costly inspecting $(1, 2)$ is governed by the standard myopic search process illustrated in [Weitzman \(1979\)](#)²². In particular, $u_{1,1}$ is known, and $(1, 2)$ ’s value is only unknown in B_2 . Therefore, we can express the value of inspecting $(1, 2)$ using [Weitzman \(1979\)](#)’s reservation value. In particular, the certain equivalent that makes a consumer indifferent between that value and costly discovering realization B_2 is the value z that solves:




Figure 4: *Two products available*

$$s = \int_z^{\hat{B}} (B_2 - z) dF(B_2).$$

Then, the reservation value of inspecting $(1, 2)$ when A_1 is known is simply²³:

$$r_{1,2} = A_1 + z.$$

Following [Weitzman \(1979\)](#), the consumer would inspect $(1, 2)$ if and only if $B_1 < z$, or $u_{1,1} < r_{1,2}$. [Figure 5](#) illustrates.

We cannot go backwards and apply the same myopic logic to the choice of inspecting $(1, 1)$: since the reservation value of each individual product depends on the other, we cannot apply

²¹Once again, this approach was inspired by the work contained in [Anderson et al. \(2021\)](#); I thank Daniel Savelle for his many helpful comments.

²²This intuition can be found, for example, in [Ke and Lin \(2022\)](#)

²³Notice that this is the same utility structure studied in [Choi et al. \(2018\)](#)

Pandora’s search algorithm.

Suppose however that the two products were in a bigger box, and that the consumer had to decide whether to open one box containing both (1, 1) and another box, containing (1, 2), or nothing at all. Following the intuition detailed in [Anderson et al. \(2021\)](#), opening this “compound” box, that I refer to as $X_{1,1}$, can be scored.

If the consumer opens the box, she discovers $u_{1,1}$, the implied reservation value $r_{1,2}$, and searches accordingly. Since we know how search takes place inside this box, we can give a score to $X_{1,1}$ that reflects not just the value of inspecting (1, 1) but also the value of the information learned through the possibility of correcting towards (1, 2). As shown below, this intuition can be generalized to an arbitrary number of products. When applied to each product separately, compound boxes generates an environment in which products sharing attributes can be appropriately scored to reflect the information they carry.

Notice that the consumer could also want to inspect (1, 2) first. We can then imagine another compound box, $X_{1,2}$, containing (1, 2) and a nested box containing (1, 1). The two are *ex ante* identical before either is opened, and once one is opened, the other becomes the smaller nested box contained in the one inspected first. Maintaining the assumption that the consumer inspects unknown attributes in increasing order of their index when indifferent is still without loss of generality.




Figure 5: *Myopic search after inspecting (1, 1)*

The value of a compound box. Consider the framework with infinite variants available for each attribute, and therefore with infinite products to, potentially, inspect. Let us build a compound box $X_{1,1}$ around product (1, 1). From the above, this box contains (1, 1) and infinitely many smaller boxes with all products characterized by either A_1 or B_1 . We must give a score to this compound box. Once $X_{1,1}$ has been opened, the consumer has learned her valuation for A_1 and B_1 . Suppose she decided to keep searching in this compound box by inspecting the nested boxes therein.

Suppose $A_1 > B_1$: it is clear that the consumer would choose to keep A_1 and possibly search to discover the best B_j out of the infinite many combinations available. Notice that all B_j s are independent. As this is the standard [Weitzman \(1979\)](#)’s environment, we can once again apply Pandora’s search. We know that each of the nested boxes can be assigned a score and searched in decreasing order or said score. Since A_1 is known, this environment perfectly mimics that analyzed in [Choi et al. \(2018\)](#). Further, since all attributes are i.i.d., all nested boxes share the same reservation value z satisfying:

$$s = \int_z^{\hat{y}} (y - z) dF(y),$$

where y represents a generic attribute. Notice that, given distribution F , z only depends on s . As shown in [Choi et al. \(2018\)](#), every reachable nested box has reservation value $r_{i,1} = A_1 + z$ if attribute A_1 is kept, or $r_{1,j} = B_1 + z$ if B_1 is kept instead.

Since there are infinite products, the consumer knows that she will keep opening boxes until she finds a B_j that beats the reservation value z of all of the boxes that would follow. Therefore, she expects to keep a product generating utility:

$$E[u_{1,j}] = A_1 + E[y|y \geq z],$$

where $A_1 > B_1$ is kept, and the latter is the expected valuation of an attribute that would only be kept if larger than z . To ease the notation, I define:

$$\bar{y} \equiv E[y|y \geq z],$$

which is immediately pinned down by the distribution F and search cost s through z . Notice that this value is only valid if the consumer does not open a compound box and “backtrack” to a previously inspected alternative: suppose that the consumer has opened compound boxes $X_{1,1}$ and $X_{2,2}$, and optimally decided to inspect product $(1, 2)$ next. When the decision is made, the value of this product is fully known and equal to $A_1 + B_2 - s$. I show below that this dynamic can be accounted for by updating the value of affected unopened compound boxes upon encountering such a realization. At the beginning of the game, before any box has been opened, this is clearly not relevant.

Further, notice that this search dynamic is only relevant if $z \geq 0$, or, if $s \leq E[y]$ ²⁴. If search costs are higher than $E[y]$, the consumer would never optimally choose to inspect nested products since it would hold:

$$u_{1,1} = A_1 + B_1 > \max\{A_1, B_1\} + z$$

for any realization A_1, B_1 . Then, if $s > E[y]$ the effective content of a compound box is simply the products immediately available before any inspection has taken place. For now, assume $s \leq E[y]$; the case in which $s > E[y]$ is characterized towards the end of the section.

Suppose search has yet to start (and $s \leq E[y]$). The compound box containing $(1, 1)$ and all nested products sharing an attribute with it can be scored. To do so, let’s decompose the possible scenarios. If both A_1, B_1 are above z , clearly the consumer would not search anything else inside the box. If only one is, say A_1 , she would keep that and search through the B attributes until she finds one that beats z . The same is true if B_1 is and A_1 is not. What happens if both are below z ?

Since there are infinite products, the consumer would keep the highest one, say A_1 without loss of generality, and search through the various B_j until she find something to keep. Therefore, before the compound box is opened, the consumer expects to gain from it utility equal to:

$$\begin{aligned} E[X_{1,1}] = & [1 - F(z)]^2(\bar{y} + \bar{y}) \\ & + 2F(z)[1 - F(z)](\bar{y} + \bar{y}) \\ & + F(z)^2 \left(\bar{y} + \frac{1}{F(z)^2} \int_0^z y dF(y)^2 \right), \end{aligned}$$

²⁴The knife edge case in which $s = E[y]$ generates $z = 0$. In this case, a consumer would be indifferent between keeping $(1, 1)$ and opening nesting boxes if and only if $\min\{A_1, B_1\} = 0$. As this is an event with probability zero, weakness of the inequality is irrelevant.

where the first two component reflect the expected value of both attributes being above z , either found immediately or eventually reached; the last component, instead, is \bar{y} , eventually found, and the highest of two realizations below z (and can therefore be expressed through order statistics). Rearranging:

$$E[X_{1,1}] = [2 - F(z)^2] \bar{y} + \int_{-\infty}^z y dF_{(2,2)}(y).$$

Following [Choi et al. \(2018\)](#), we can back out the reservation value of this compound box from the above formulation; in particular:

$$R_{1,1} = [2 - F(z)^2] \bar{y} + F(z)^2 z,$$

where, from the standard formulation contained in [Weitzman \(1979\)](#), z satisfies:

$$s = \int_z^z \left(\frac{y}{F(z)^2} - z \right) dF(y)^2.$$

The same can be done for all compound boxes: $R_{i,j} = [2 - F(z)^2] \bar{y} + F(z)^2 z$, $\forall i, j$ is the value of each compound box before search starts. Notice that, because products are *ex ante* identical, the value attached to each compound box must be the same. Notice also that this value is higher than the value of inspecting any of the products in isolation: just like in the simplified framework, the presence of uninspected alternatives sharing an attribute with a products makes inspecting it more valuable in expectation.

With infinite boxes, the consumer would open compound boxes in decreasing order of their index (because of the ordering assumption) and stop whenever either the highest realization or the reservation value r of the best nested box inside of it (since upon opening a compound box the reservation values of the nested boxes inside become visible) is higher than the reservation value of the next compound box. With infinite boxes, the consumer never stops until he finds something to keep. All possible combinations of values must be accounted for.

Notice that the value $R_{i,j}$ determined above relies on the consumer following the optimal search policy inside the fictitious box we built around each product and depicted in [Figure 6](#). Further, notice that the way these boxes were built is such that the same product is “contained” in multiple compound boxes. This is by design: in doing so, the value of the box can be updated following how realizations affect the optimal search from that point onward. In what follow, I show that the the value of all unopened compound boxes depends non-trivially, but predictably, on past realizations in a way that allows to determine cut-offs that depend only on distributional assumptions and search costs, above which the consumer always keeps a realized attribute. Further, I show that in this environment a consumer can rationally discard an attribute and then go back to it depending on the search costs. Finally, I show that for s high enough the correlation between products is immaterial, and the optimal search policy is exactly [Weitzman \(1979\)](#)’s search




Figure 6: *Graphical representation of optimal search inside compound box $X_{1,1}$.*

algorithm.

5.2. Optimal search in the general framework

Dynamic compound boxes. Before the consumer opens any box, all available products are contained in compound boxes $X_{i,j}$ with the score $R_{i,j}$ characterized above. This score is not static: since different compound boxes contain the same product, their values change as boxes are opened. Suppose that the consumer opened compound box $X_{1,1}$: she learned her valuation for A_1 and B_1 , and can now reach all products sharing one of those attributes through the nested boxes with value $r_{1,j}$ or $r_{i,1}$ inside of it. The value of opening $X_{2,2}$ depends on the realizations A_1 and B_1 . the formulation $R_{2,2} = [2 - F(z)^2]\bar{y} + F(z)^2\underline{z}$ crucially depends on the consumer not going back to either $(1, 2)$ or $(2, 1)$. Next, I determine when this is indeed the case and, when it is not, what are the implications. For now, I keep the assumption that $s \leq E[y]$: towards the end of the section I detail the implications of $s > E[y]$ for the search process.

If the consumer were to open $X_{2,2}$ and decided to keep one of the two attributes, she would start inspecting unopened nested boxes to discover her valuation for variants of the other attribute only if it was better to do so than to go back to a product defined by either A_1 or B_1 . In particular, if (without loss of generality) B_2 is kept, it must hold:

$$r_{2,j} = z + B_2 > A_1 + B_2 - s.$$

If both A_1 and B_1 are below $z + s$, search inside subsequent compound boxes follows the order characterized above, and the value $R_{2,2}$ assigned to $X_{2,2}$ before search started is still valid. In this case, the consumer would search according to standard search logic, selecting the best between $u_{1,1}$ (and stopping), $r_{1,j}$ (and open nested boxes characterized by A_1), $r_{i,1}$ (and open nested boxes characterized by B_1), and $R_{2,2}$ (and discover two new attributes) to proceed. Equivalently, if both A_1 and B_1 are below $z + s$, $(1, 2)$ and $(2, 1)$'s presence inside $X_{2,2}$ is immaterial: since the consumer would never inspect them, it is as if they are not there. Figure 7 illustrates.

Suppose now that $A_1 \geq z + s$. If this is the case, the consumer knows that opening $X_{2,2}$ would make her back track to $(1, 2)$ instead of inspecting $(3, 2)$ if she wanted to keep B_2 . In this case, the formulation $R_{2,2} = [2 - F(z)^2]\bar{y} + F(z)^2\underline{z}$ is invalid and must be updated. Updating this value requires to retrace the steps performed in the last subsection, and consider how the expected utility of different outcomes changed with the change in the optimal search path. Suppose $A_1 \geq z + s > B_1$ has already been learned; suppose further that the consumer decided to open $X_{2,2}$ and discovered her valuation for A_2, B_2 . How would she proceed?

If $B_2 > z$, $(2, 2)$ beats all unopened nested boxes containing a product sharing B_2 . Then, the consumer would optimally choose between keeping A_2 if $A_2 > A_1 - s$, or backtrack to $(1, 2)$, which has higher reservation value than all unopened nested boxes by construction. Instead, if $B_2 < z$, the consumer would always prefer opening nested boxes keeping A_2 then staying at $(2, 2)$. She would choose to do so if and only if:

$$A_2 + z > A_1 + B_2 - s \iff A_2 > A_1 - s - (z - B_2).$$

If A_2 is high enough, keeping it is better than backtracking; otherwise, backtracking is still




Figure 7: *Evolution of the search environment after inspecting (1,1). If $\max\{A_1, B_1\} < z + s$, they do not affect search in unopened compound boxes (on the left). Products characterized by attributes A_1 or B_1 are represented by their nested equivalent; products unrelated to (1,1) are unaffected (on the right).*

preferable. Notice that these are the relevant choice only because by construction compound boxes only contain products sharing attributes with the product around which they are built. This is crucial to appropriately define the value of opening the compound box compared to that of keeping a discovered attribute. From the above, after finding that $A_1 \geq z + s$, the reservation value of unopened compound box $X_{2,2}$ (and all other unopened compound boxes) updates to:

$$\begin{aligned} \tilde{R}_{2,2}^{A_1} = & [1 - F(z)][1 - F(A_1 - s)](E[y|y \geq A_1 - s] + \bar{y}) + \\ & [1 - F(z)][F(A_1 - s)](A_1 - s + \bar{y}) + \\ & [F(z)][1 - F(A_1 - s - (z - \underline{y}))](E[y|y \geq A_1 - s - (z - \underline{y})] + \bar{y}) + \\ & [F(z)][F(A_1 - s - (z - \underline{y}))](A_1 - s + \underline{z}), \end{aligned}$$

where $\underline{y} = E[y|y < z]$ is the expected value of a realization lower than z , \underline{z} solves:

$$s = \int_{\underline{z}}^z \left(\frac{y}{F(z)} - \underline{z} \right) dF(y),$$

and apex A_1 reflects the fact that $A_1 > z + s$ reroutes search towards itself.

Notice that since a product defined by A_1 can be found in all unopened compound boxes, the value of all of them update in the same way. More in general, since the highest such realization would be the one to be selected, I identify the value of a compound box in which an attribute y reroutes search towards itself as $\tilde{R}_{i,j}^A$ and $\tilde{R}_{i,j}^B$ if $\max\{y \in A \cap I\} > z + s$ or $\max\{y \in B \cap I\} > z + s$ respectively. Figure 8 illustrates the effect of finding $A_1 > z + s > B_1$. Whenever unambiguous, I refer to a box in which backtracking can be optimal as $\tilde{R}_{i,j}$.

If a consumer finds an attribute $y > z + s$ upon opening compound box $X_{1,1}$, she will proceed selecting the highest between $u_{1,1}$ (and stop), $r_{1,j}$, $r_{i,1}$ (and open nested boxes), and $\tilde{R}_{i,j}$ (and open compound boxes). If the latter is selected, the consumer will repeat the same choice




Figure 8: If $A_1 > z + s > B_2$, A_1 reroutes search towards itself (on the left). Since a product characterized by A_1 can be found in all unopened compound boxes, this affects the value of all of them (on the right).

updating as necessary, and possibly backtracking to a product whose attribute she has learned about during the search process. More importantly, assuming $A_1 > z + s > B_1$ is realized, the consumer is indifferent between opening new compound boxes and keeping A_1 if it holds:

$$A_1 + z = \tilde{R}_{2,2} \iff A_1 = \tilde{R}_{2,2} - z.$$

\tilde{R} depends non-trivially on s and on the realization $y > z + s$ that caused the updating. A higher search cost implies a lower threshold z (and $z + s$, since z decreases faster than s increases under our distributional assumption). This, in turn, implies that relatively low realizations are enough to trigger the updating for higher search costs. Further, it is clear that a higher realization $y > z + s$ increases the value of the unopened compound box, since it contains it. However, this increment is limited by the extra search cost to be paid to reach the relevant product inside this compound box, and the shrinking probability of finding a higher realization in the same dimension. For a high enough realization y , it can be shown that $y + z > \tilde{R}(y)$ always holds.

Of course, both A_1 and B_1 could be found to be above $z + s$ and, depending on s and the distributional assumption, not be jointly enough to beat the next unopened compound box. Notice however that if this is the case the update characterized above is incorrect, as B_1 beats all unknown variants B_j just like A_1 did before. The same procedure as above allows to obtain a different updated value, $\tilde{R}_{i,j}^{A,B}$, which becomes relevant from that point onward. With the same calculations as above:

$$\begin{aligned} \tilde{R}_{2,2}^{A,B} = & [1 - F(B_1 - s)][1 - F(A_1 - s)](E[y|y \geq A_1 - s] + E[y|y \geq B_1 - s]) + \\ & [1 - F(B_1 - s)][F(A_1 - s)](A_1 - s + E[y|y \geq B_1 - s]) + \\ & [F(B_1 - s)][1 - F(A_1 - s)](E[y|y \geq A_1 - s] + B_1 - s) + \\ & [F(B_1 - s)][F(A_1 - s)](\max\{A_1, B_1\} - s + \tilde{z}), \end{aligned}$$

where \tilde{z} solves:

$$s = \int_{\tilde{z}}^z \left(\frac{y}{F(\min\{A_1, B_1\} - s)} - \tilde{z} \right) dF(y).$$

No other forms of updating is possible.

Evolution of the search environment. Given the structure above, it is clear that keeping a discovered attribute (or, opening nested boxes instead of the next compound box) can be optimal only if the attribute by itself is high enough to beat the updated informational value attached to an undiscovered alternative. Formally, it can be shown that:

Proposition 3. *For any distribution $F(y)$ and search cost $s \leq E[y]$, there exist a unique value $y^T = \max\{z + s, \tilde{y}\}$, where \tilde{y} is the value that satisfies $\tilde{y} = \tilde{R}^{\tilde{y}} - z$, such that an attribute strictly above y^T is always kept. $z + s > \tilde{y}$ when search costs are below some threshold $0 < \tilde{s} < E[y]$. If $s < \tilde{s}$ (resp. $s > \tilde{s}$), y^T is decreasing (resp. increasing) in s .*

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

The result establishes that backtracking to a previously inspected attribute can be rational if search costs are high enough (but, for now, not higher than $E[y]$), that is, when $\tilde{y} > z + s$. In particular, if an attribute is high enough to beat an undiscovered substitute, but not high enough to beat the (updated) expected value of an unopened nested box by itself, a rational consumer would still prefer to discover two new attributes and, possibly, go back to it after finding a good enough attribute to go with it in the other dimension. On the other hand, if search costs are low enough, a discarded variant is never returned to, and the updating process does not bite: for s low enough, if a realization is high enough to affect the value of subsequent boxes it is also high enough to keep.

While apparently counter-intuitive, the result of Proposition 3 follows from the fact that the lower s is, the higher the threshold $z + s$ is. For small enough s , the value \tilde{y} that solves $\tilde{y} = \tilde{R}(\tilde{y}) - z$ is lower than $z + s$ and, therefore, is irrelevant. \tilde{y} grows in the search cost, while $z + s$ decreases. At \tilde{s} , the two are equivalent, and $\tilde{y} = z + s$ exactly. As s keeps growing, \tilde{y} keeps increasing and $z + s$ keeps decreasing: the distance between the two widens, and more realizations that change the value of unopened compound boxes without beating them in expectation become feasible. When this is the case, backtracking to a previously inspected attribute becomes possible provided that the realizations that follow make that choice superior, in expectation, than inspecting two undiscovered attributes.

Notice however that an attribute below $z < y < y^T$ could still be rationally kept if found in conjunction with another attribute such that, together, they beat $R_{i,j}$ for all unopened compound boxes $X_{i,j}$. In other words, Proposition 3 establishes a sufficient but generally not necessary condition for an attribute to be kept.

The updating process detailed here allows the value of unopened boxes to change depending on past realizations in a way that maintains consistency of the search process. In particular, at any “decision node” the value of every unopened box is pinned down by $R_{i,j}$, $R_{i,j}^y$, or $R_{i,j}^{A,B}$ depending on the highest past realization of inspected attributes $y \in A \cap I$ and $y \in B \cap I$. Notice that

these scores are mutually exclusive: Define $A^H = \max\{y \in A \cap I\}$, $B^H = \max\{y \in B \cap I\}$ the highest past realization in A and B ; then:

$$\mathcal{R}_{i,j} = \begin{cases} R_{i,j} & \text{if } \max\{A^H, B^H\} < z + s, \\ R_{i,j}^y & \text{if } y = \max\{A^H, B^H\} > z + s > \min\{A^H, B^H\}, \\ R_{i,j}^{A,B} & \text{if } \min\{A^H, B^H\} > z + s. \end{cases}$$

assigns the appropriate value to all compound boxes yet to inspect. Since by construction all unopened boxes update in the same way because all compound boxes contain at least one product characterized by every attribute available, it follows that:

Corollary 1. *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 (as long as one is available).*

Proof. All calculations can be found in Appendix B ■

By tracing optimal search along discovered attributes, $\mathcal{R}_{i,j}$ reflects not only the value of the product it represents, but that of all relevant information learned as well in terms of its implications for the search process that would follow. Then, at every decision node of the search process, Weitzman (1979)'s logic can be applied. While the whole search process cannot be determined *ex ante* because of the learning component introduced by correlation of products, at each decision node of the consumer there is no ambiguity regarding the value of her possible options. The main result of this section follows:

Proposition 4. *Define $A^H = \max\{y \in A \cap I\}$, $B^H = \max\{y \in B \cap I\}$ the highest past realization for A and B . Optimal search is characterized as follow:*

- **Search order:** *If a box is to be opened, it should be the reachable box with the highest reservation value $\mathcal{R}_{i,j}$ or $r_{i,j}$.*
- **Compound box selection:** *compound boxes are opened until all unopened compound box has reservation value $\mathcal{R}_{i,j}$ lower than the highest realized outcome or reservation value out of all reachable nested box.*
- **Stopping rule:** *Boxes are opened until the highest unopened (compound or nested) box has reservation value below the highest realized outcome.*

Where $\mathcal{R}_{i,j} = R_{i,j}$ if $\max\{A^H, B^H\} < z + s$, $\mathcal{R}_{i,j} = \tilde{R}_{i,j}^y$ if $\max\{A^H, B^H\} \geq z + s > \min\{A^H, B^H\}$, and $\mathcal{R}_{i,j} = \tilde{R}_{i,j}^{A,B}$ if $\min\{A^H, B^H\} \geq z + s$.

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

High search costs. It can be shown that the result of Proposition 4 applies generally. When search costs are so high that inspection of an attribute by itself is not worth its cost, that is, when $s > E[y]$, however, a few adjustments must be made. In particular, the consumer can never optimally inspect unknown nested products. This follows, once again, from the structure

of the compound box. Since $s > E[y]$ implies $z < 0$, any realization $u_{i,j}$ discovered opening a compound box beats all the adjacent nested boxes by construction:

$$u_{i,j} = A_i + B_j > \max\{A_i, B_j\} + z.$$

Therefore, before any search has taken place, the value of opening a compound box must be exactly the value of the product it represent and nothing else. Let $Y_{i,j} \equiv A_i + B_j$, and $G(\cdot) = 2F(\cdot)$ its distribution. Then, each unopened compound box can be represented by z_M that solves:

$$s = \int_{z_M}^{\hat{Y}=2\hat{y}} (Y - z_M) dG(Y).$$

This, however, does not mean that the value of compound boxes is static, just that searching forward inside a compound box cannot be optimal. However, going backwards can be optimal just like it was for lower search costs; in this case, updating is triggered by realizations $y \geq s > E[y]$.

It can be shown that the effect of this updating depends on the relative value of s and z_M . In particular, there exist a unique value $\bar{s} \in (E[y], E[Y])$, which solves $s = z_M$, such that for $s \geq \bar{y}$ correlation is immaterial to the optimal search policy from an outcome standpoint. While unopened compound boxes do update in their score following the logic detailed above, whenever this happens (that is, when some $y \geq s \geq \bar{s}$ is found), this realization is enough to induce the consumer to stop and keep it. Furthermore, since $s \geq z_M$, the optimal search process in this case is equivalent to that detailed in [Weitzman \(1979\)](#):

Corollary 2. *There exist a unique value $E[y] < \bar{s} < E[Y]$ such that, for $s \in [\bar{s}, E[Y]]$, correlation through attributes is immaterial, and the optimal search process is equivalent to [Weitzman \(1979\)](#)'s myopic search: consumers inspect uncorrelated products sequentially until a realization $Y > z_M$ is found.*

Proof. All calculations can be found in [Appendix B](#). ■

To understand the result of [Corollary 2](#), suppose A_1 was high, B_1 was low, and $A_1 + B_1 < z_M$. Suppose the consumer chose to open next $X_{2,2}$. Then, the consumer would choose between:

- keeping (1, 1),
- keeping (2, 2),
- inspecting (1, 2), with known value $r_{1,2} = u_{1,2} - s$,
- inspecting (2, 1), with known value $r_{2,1} = u_{2,1} - s$,
- opening $X_{3,3}$.

As it was the case for low search costs, the value of opening $X_{2,2}$ must reflect the ability of the consumer to combine A_1 and B_2 , or B_1 and A_2 . This gives rise to the same updating detailed above. In particular, updating only bites when some attribute $y \geq s$ is found. In [Appendix B](#), I show that for $s \in (E[y], \bar{s})$, this gives rise to the usual updating dynamic; in particular, in this interval, the sufficient (but not necessary) condition for an attribute to be kept is $y \geq z_M + \epsilon$ for




Figure 9: *Evolution of the search environment when $s > E[y]$. Before any inspection has taken place (on the left), all products are scored to reflect their myopic value; after inspection of (1, 1), all unopened boxes update to reflect the information learned. If A_1 and B_1 are below s , unrelated boxes are unaffected (in the center). Otherwise, they can reroute search towards themselves (on the right).*

some $\epsilon > 0$ pinned down by G . When $s \in [\bar{s}, E[Y]]$, instead, such updating does not bite and the sufficient condition for a product to be kept is simply $Y \geq z_M$. Therefore, for high enough search costs the optimal search policy is unaffected by correlation through shared attributes, and is equivalent to Pandora’s algorithm.

The result echoes the high search cost case analyzed in Section 4: in the binomial case it holds that if search costs are so high that inspection of a single attribute is not worth its cost, correlation between products does not bite. This needs not hold in general: while for the binomial case an attribute not good enough to be kept is always discarded, more general distributions allow for more nuanced updating. With high enough search costs however, the condition that leads to said updating is more stringent than the one necessary to keep a discovered attribute. In this case, optimal search is equivalent to Weitzman (1979)’s myopic optimal policy, and only products unknown in both attributes can be rationally inspected.

The model, then, can be used to rationalize the well known relation between the recent steady reduction in search costs fostered by the internet and the “long-tail effect”. As Bar-Isaac et al. (2012), among many others, argues, new and more efficient search technologies benefit fringe products in the market, and creates incentives for new fringe products to be introduced. The model suggests that the same is true for the variety of alternative version of the same product that cater to diverse tastes.

5.3. Optimal pricing with infinite products

The two point distribution assumption behind the simplified framework allowed to effectively ignore the updating dynamic that characterizes the general setting. Net of pricing effects, in the simplified framework an attribute is always either a match, and kept, or not a match, and discarded. The more general framework incorporates history-dependent updating that does not allow for this kind of clear pathing. However, it resembles, mechanically, the simplified framework. How these differences affect the seller’s strategy is the topic of this subsection.

In what follows I characterize pricing in the general model detailed above, that is, with two sets in infinite i.i.d. attributes defining products in an infinite grid. I once again focus my attention on the optimal strategy of a multiproduct monopolist with control over menu and pricing, and

assume him to post prices before the representative consumer starts inspecting products.

To understand pricing in this specification, consider again the structure of a compound boxes. The value of inspecting an unknown product depends on the expected value of the follow-up inspections along one of the attributes the product is defined by. The structure presented above can be readily adapted to incorporate prices. In particular, the value associated with each product must be reduced by the posted price; this new value, then, can be used to score compound boxes appropriately, accounting for the price of all products on the relevant search paths. Notation-wise, price of product (i, j) is henceforth indexed by $p_{i,j}$, which is the same convention used in Section 4.

Consider the compound box $X_{1,1}$ built around product $(1, 1)$ priced at $p_{1,1}$; the box contains all products $(1, j)$, priced at $p_{1,j}$, and all products $(i, 1)$, priced at $p_{i,1}$. Suppose the consumer opened $X_{1,1}$ and decided to keep searching in it keeping attribute A_1 . Then, she would inspect next the product $(1, j)$ that satisfies:

$$\max_j (A_1 + z - p_{1,j}) \geq A_1 + B_1 - p_{1,1},$$

where, once again, the value of inspecting any of the i.i.d. B_j attributes can be expressed by the score z .

Three things are worth noticing: first, if $p_{1,j}$ is not uniform, the consumer would always select to inspect products $(1, j)$ in increasing order of price. Second, for $X_{1,1}$ to be inspected before all other $(1, j)$ products, it must have been the cheapest of them. Third, if $p_{1,1} \neq p_{1,j}$, $(1, j)$ would be inspected next if and only if:

$$B_1 \leq z - (p_{1,j} - p_{1,1}) < z.$$

Of course, the same structure governs inspection of products $(i, 1)$.

In principle, all products $(1, j)$ could be priced differently. Suppose that prices were increasing in j and always strictly below z . Then, if the consumer decided to inspect $(1, 2)$ after discovering A_1, B_1 , he would expect to either keep it if it beats the reservation value of $(1, 3)$, or keep searching, and so on for all subsequent inspections. The total value associated with this path given vector of prices $\mathbf{p}_{1,k}$ of all products $(1, k > 1)$ is then:

$$y(\mathbf{p}_{1,j}) = \sum_{k=1}^{\infty} F(z - (p_{1,k+1} - p_{1,k}))^k \int_{z - (p_{1,k+1} - p_{1,k})}^{\hat{y}} (y - p_{1,k+1}) dF(y).$$

All compound boxes $X_{i,j}$ can be scored exactly as before assuming prices increase consistently. Define $\Delta_{i,k} \equiv p_{i,k+1} - p_{i,k}$ and $\Delta_{k,j} \equiv p_{k+1,j} - p_{k,j}$; further, define

$$\bar{y}_{i,k} = E[y|y > z - \Delta_{i,k}], \quad \bar{y}_{k,j} = E[y|y > z - \Delta_{k,j}].$$

Then:

$$\begin{aligned}
R_{i,j}(\mathbf{p}_{i,j}) = & [1 - F(z - \Delta_{i,i+1})][1 - F(z - \Delta_{j+1,j})](\bar{y}_{i,i+1} - \bar{y}_{j+1,j} - p_{i,j}) \\
& + [1 - F(z - \Delta_{i,i+1})]F(z - \Delta_{j+1,j})(\bar{y}_{i,i+1} + y(\mathbf{p}_{i,k})) \\
& + F(z - \Delta_{i,i+1})[1 - F(z - \Delta_{j+1,j})](y(\mathbf{p}_{k,j}) + \bar{y}_{j+1,j}) \\
& + F(z - \Delta_{i,i+1})F(z - \Delta_{j+1,j})\left(\underline{z}_{i,j} + \max\{y(\mathbf{p}_{i,k}), y(\mathbf{p}_{k,j})\}\right),
\end{aligned}$$

where $\underline{z}_{i,j}$ is the certain equivalent of opening a box distributed according to the highest of realizations $A_i < z - \Delta_{i,i+1}$ and $B_j < z - \Delta_{j+1,j}$ at cost s .

Obtaining an equilibrium pricing vector requires studying the impact of different price combinations on the optimal search process. While a high price that does not make a product never worth inspecting makes it more profitable to sell, it also pushes the product attached to it further away from the optimal starting point of the consumer. Suppose all products were priced at some uniform level p^u and one was slightly more expensive. Then, not only the more expensive product would have lower value in any search path in which it could be found, but all compound boxes that contain it would also have a lower $R(\mathbf{p})$ score. It follows that none of the boxes associated with this product would ever be inspected as there are infinite better alternative for the consumer.

Another difficulty relates to the updating described in the pages above. Attributes can still have realizations that reroute search towards themselves, and in a way that is much more cumbersome to keep track of when prices are accounted for. Moreover, since the relationship between R , \tilde{R}^y , and $\tilde{R}^{A,B}$ depends on the specific realization or realizations that triggered the update, the updating could lead to all unopened boxes to become less valuable than they originally were, which could lead the consumer to end his search prematurely.

Both concerns can be addressed, and the following result emerges:

Proposition 5. *Consider a multiproduct seller pricing infinite products defined by two infinite sets of i.i.d. attributes. It holds:*

- *there exist a unique equilibrium pricing vector,*
- *the equilibrium pricing vector is uniform,*
- *the uniform equilibrium price can be expressed as:*

$$p^* = \lambda^* R_{i,j} + (1 - \lambda^*) \underline{\tilde{R}}, \quad \lambda^* \in [0, 1],$$

where $\underline{\tilde{R}} \equiv \min\{\tilde{R}_{i,j}^y|_{y=z+s}, \tilde{R}_{i,j}^{A,B}|_{A=B=z+s}\}$.

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

Proposition 5 states that the only possible equilibrium features uniform pricing that can be expressed as a linear combination of different possible values of compound boxes. The first part of the result follows from the same intuition that generated differential pricing equilibria in the simplified model of Section 4. With only four products, different prices could arise specifically

because the second product searched was generally the last one: after inspecting (1, 1), if (2, 2) was inspected afterwards, the consumer would either purchase it or stop searching since no other products were available to inspect. To make it worth searching, it had to be priced accordingly.

This is still true for the infinite attribute case. However, since there are infinite products available, all compound boxes are of infinite size at the beginning of the search process. It follows that “discounting” some products to encourage search after a bad realization is unnecessary, as the value of inspecting products does not decrease mechanically after every inspection²⁵.

The fact that compound boxes do not shrink by itself does not imply that products cannot be priced differently. In principle, given the total value of a compound box $R_{i,j}$, different products could be priced differently to capitalize on the information learned through inspection just like it was the case for the simplified framework. In Appendix C, I show that this cannot be optimal. The intuition is as follows: suppose that, following the intuition driving the results of Section 4, compound box $X_{1,1}$'s products were priced according to $p_{1,1} = p$ for some $p > 0$ and $p_{1,j} = p_{i,1} = p + \delta$ for some $\delta > 0$ ²⁶. Plugging in these prices in the score of the compound box, one finds:

$$\begin{aligned} R_{1,1}(\mathbf{p}_{1,1}) = & [1 - F(z - \delta)]^2(2\bar{y}_\delta - p) \\ & + 2F(z - \delta)[1 - F(z - \delta)](2\bar{y}_\delta - (p + \delta)) \\ & + F(z - \delta)^2(\bar{y}_\delta + \underline{z}_\delta - (p + \delta)), \end{aligned}$$

where $\bar{y}_\delta \equiv E[y|y > z - \delta]$, and \underline{z}_δ solves:

$$s = \int_{\underline{z}_\delta}^{z-\delta} \left(\frac{y}{F(z - \delta)^2} - \underline{z}_\delta \right) dF(y)^2.$$

Studying $R_{1,1}(\mathbf{p}_{1,1})$ reveals that any positive δ would be detrimental to the expected profit of the seller. On one hand, the probability that the consumer finds a realization that induces her to keep searching after inspecting (1, 1) shrinks as δ increases since $F(z - \delta)$ is decreasing in δ . On the other hand, the participation constrain implied by the fact that the consumer must decide to open the first box becomes tighter as δ increases.

To see why, notice that the above can be rewritten as:

$$R_{1,1}(\mathbf{p}_{1,1}) = [2 - F(z - \delta)^2]y_\delta + F(z - \delta)^2\underline{z}_\delta - \left(p + \delta \left[1 - [1 - F(z - \delta)]^2 \right] \right).$$

It is clear that no pricing such that $R_{1,1}(\mathbf{p}_{1,1}) < 0$ can be optimal, since it would discourage the consumer from searching. On the other hand, notice that the value of opening a compound box net of prices is equivalent to that of opening the same box when search costs are higher, and in particular $s' > s$ such that $z' = z - \delta$. It follows that $\delta > 0$ makes starting the search process less valuable, which tightens the consumer participation constraint and, therefore, how high prices that do not discourage search can be. Since different prices tighten the consumer participation constraint, make it, all else equal, more likely to sell at a lower price than under uniform prices, and do not generate additional probability of trade, it is clear that a non-uniform

²⁵A more in depth discussion about the finite number of attributes case can be found in the Extensions.

²⁶In the Appendix, I show that if an equilibrium with differential prices exists, it must have prices following this structure.

pricing vector cannot be optimal.

The structure of p^* follows from the updating dynamic detailed in the previous subsection. Since any $y \geq z + s$ can change the value of all subsequent unopened boxes, and since $\tilde{R} \equiv \min\{\tilde{R}_{i,j}^y|_{y=z+s}, \tilde{R}_{i,j}^{A,B}|_{A=B=z+s}\}$ is the lowest value any unopened compound box can ever have at any point of the search process, any optimal price level must be between this value and $R_{i,j}$, the highest price that does not prevent search from taking place. The former must, by construction, lead to probability of trade taking place equal to one; the latter, instead, leads to the lowest probability of trade sustainable in equilibrium.

To understand how different pricing levels affect expected profits of the monopolist one must identify the expected probability of trade for any price. This is straightforward. Trade does not happen if, after an update is triggered, the consumer prefers to stop searching rather than go on. For this to happen, prices must be high enough, and the realization triggering the update must be low enough, that the new reservation value of all unopened boxes is lower than the respective price.

Formally, Define:

$$p_\lambda = \lambda R_{i,j} + (1 - \lambda)\tilde{R}, \quad \lambda \in [0, 1].$$

a generic candidate equilibrium price, and:

$$y = y_\lambda : \tilde{R}_{i,j}^y|_{y=y_\lambda} = p_\lambda.$$

The realization y that equates the updated value it triggers to its price. Since $\tilde{R}_{i,j}^y$ is increasing in y , y_λ must be increasing in p_λ . Then, the probability of trade taking place if $p^u = p_\lambda$ is the probability of a realization such that the update is triggered and no available product, inspected or uninspected, is kept or discovered. Expected profits of the monopolist $p^u = p_\lambda$ is then:

$$\pi_\lambda = p_\lambda \left[1 - Pr \left(y_{2,2} \in (z + s, y_\lambda) \wedge y_{1,2} < \tilde{R}_{i,j}^y|_{y=y_\lambda} - y_{2,2} \right) \right],$$

where $y_{2,2}$ and $y_{1,2}$ are the highest and lowest of two independent i.i.d. realizations, which represent the realized utility of an inspected product.

Notice that the price increases and the probability of trade decreases in λ in ways that are directly pinned down by $F(y)$ and s . The equilibrium price must then be:

$$p^* = p_\lambda : \lambda = \arg \max_{\lambda \in [0,1]} (\pi_\lambda).$$

The discussion establishes two things. First, even with infinite products available and uniform prices, probability of trade might not be equal to one. This follows from the updating dynamic, which could be triggered by a good, but not good enough, realization. Second, that the number of available products, the wideness of the menu, affects optimal pricing structure. Comparison of the simplified and the infinite attribute framework suggest that the wider the range of options available to the consumer, the more likely products are to be priced uniformly. Ease of substitution generated by abundance of search options, in principle, can deter the seller from steering consumers towards specific search paths. Conversely, a wider selection lets the consumer obtain more bad realizations without being deterred from searching without the need

of discounting products, which in turn suggests that the monopolist would have no incentive to restrict the supply.

Finally, notice that the probability of trade taking place for any given price p_λ is decreasing in s . This is intuitive: the higher s is, the lower $z + s$ is, which implies that for high search costs more realizations can make the consumer stop searching before finding something to keep. This suggests that, for s low enough, the monopolist has a stronger incentive to price more aggressively. Therefore:

Corollary 3. *For s low enough, in the unique equilibrium it holds $p_{i,j} = p^* = R_{i,j}, \forall (i,j)$, and the expected utility of search for the consumer is equal to zero.*

As search costs increase, the probability that the consumer terminate her search prematurely increases if prices are set too high. When search costs are high enough, the monopolist has the incentive to reduce them to increase the probability of trade, leaving some expected utility to the consumer as a result. The observation implies that a reduction of search costs does increase the value of search for the consumer net of prices, but allows the monopolist to extract rent more efficiently as well.

6. Extensions

I consider a variety of extensions to the general model. In what follows, I discuss how the model changes when the number of available attributes is finite and sketch the structure of three-attributes compound boxes. Further, I contextualize the assumption of realized but uninspected products require a search cost to be paid to be purchased. Finally, I discuss limitations and directions forward.

6.1. Finite number of attributes

In the baseline model I considered an infinite number of variants for each attribute. The advantage of this approach is that once the consumer starts searching in one direction, she can continue to do so without ever changing until she finds something to keep, which happens with probability one. This feature made building compound boxes reasonably straightforward. Restricting the environment to finite sets of attributes introduces new challenges. The logic underneath the structure of the compound boxes, and the search process itself is, however, unchanged.

Consider a box like the one in Figure 10. This box can be scored following the same logic used for the infinitely large boxes: if upon opening the box $\min\{A_1, B_1\} > z$, the consumer would stop. If $\max\{A_1, B_1\} > z > \min\{A_1, B_1\}$, the highest would be kept. If $z > \max\{A_1, B_1\}$, the highest would be kept as well. Unlike before, we must account for the eventuality that an attribute that beats z is never found: in this case, both variants would be inspected, and the highest would be kept.

Further, when scoring compound boxes we must account for the eventuality of inspecting keeping one attribute fixed and, after exhausting all products sharing that attribute, switching to products characterized by the other. Consider again Figure 10, and suppose A_1 and B_1 both had very low realizations such that $z > A_1 > B_1$. From the above, we know the consumer would inspect (1, 2) next. In the infinite attributes case, the consumer would never run out of (1, j) products to inspect. Now, this is not the case. Would the consumer ever optimally inspect (2, 1) next?




Figure 10: *Three products available*

For this to be the case, it must be that:

$$B_1 + z > A_1 + \max\{B_1, B_2\}.$$

If $B_1 > B_2$, this is trivially true since $z > A_1 > B_1$. Otherwise, if it holds:

$$B_2 < z - (A_1 - B_1),$$

then the consumer would optimally inspect (2, 1) next and keep the highest of all three realizations. This, too, needs to be accounted for when scoring unopened compound boxes.

Another difference with the infinitely large boxes of the baseline model follows from the fact that unlike before, when compound boxes are opened and discarded, this implies that following boxes “shrink” by one variant per attribute. Suppose $X_{1,1}$ contained products characterized by n variants of A and $m - 1$ variants of B . Further, suppose $\max\{A_1, B_1\} < z + s$. Then, $X_{2,2}$ would effectively contain products characterized by $n - 1$ variants of A and $m - 1$ variants of B . Assuming consumers search in increasing order of the index²⁷, then, the size of each subsequent compound box $X_{i,i}$ would have $n + 1 - i$ variants of A and $m + 1 - i$ variants of B .

The corner represent another complication. While it is clear that the value of a compound box containing only a product and nothing else, relevant if all other options have been exhausted, can be obtained as a standard reservation value, what happens if the box is asymmetric is less obvious. In particular, suppose there were two variants of A and three of B . Then, after opening two boxes sharing no attributes, all remaining products would be either known or with only one unknown. In this case, standard Weitzman (1979) logic would once again dictate what to inspect next.

Finally, notice that finding one or more attributes above $z + s$ affect all subsequent boxes in the same way as they did before: such an attribute beats all remaining unopened boxes in the same dimension, and reroutes search towards itself in every unopened compound box. This, in turn, allows the same exact updating detailed in the baseline model. Accounting for all of the above translates to rather unwieldy expressions for \mathcal{R} ; both R and updating logic can be found in Appendix D. However, the discussion above motivates the following:

²⁷this is, to an extent, without loss of generality if we imagine each consumers to reorder the attributes according to some private infinitesimally small taste shock; in this sense, the labeling is just an arbitrary convention since attributes are i.i.d. by assumption.

Corollary 4. *After fixing an order of indifference (e.g.: increasing order of indices), optimal search as defined in Proposition 4 (or Corollary 2) applies to the finite attribute case.*

As a final note, notice that Corollary 4 is only valid if the grid is complete: if some combination of attributes was not available, the compound boxes would not be *ex ante* identical anymore, since some would necessarily be “larger” than others. Accounting for this additional feature generally makes the problem exceedingly complex, and is therefore left to future reaserch.

6.2. More than two attributes

The baseline model features products defined by two attributes, A and B . In this extension I show that the intuition of the two attributes case can be generalized to accommodate specifications with more attributes. In particular, I obtain starting reservation values of compound boxes with three attribute products, describe the search dynamic that the structure entails, and discuss implications.

In the baseline model, two attribute products are scored by building fictitious boxes including the product itself and closed boxes with all other products sharing attributes with it. The same logic can be applied to three attribute products. The difference is that, with three attributes, two different kinds of closed boxes must be included in the larger compound box together with the product it represents. On one hand, all products sharing exactly one attribute with the central product can be represented as small nested boxes equivalent to the ones contained in the two attributes case. This follows from the same intuition: if a product shares two attributes with the visible one, it is unknown in only one dimension after opening the compound box.

On the other hand, products that share only one attribute with the central one are unknown in two dimensions, and must be placed in different, larger boxes. If attributes come in infinite variants as assumed in the baseline model, the three attributes compound box contains infinite small nested boxes and infinite larger, two-dimensional boxes equivalent to the compound boxes of the baseline model. These “intermediate” boxes, moreover, contain infinite smaller compound boxes as well. For visual aid, one can imagine multiple grids representing two attribute products side by side to resemble a cube, with the intermediate boxes representing search along one of the sides, and the small boxes representing search along one of the edges.

To score such a box, one must consider how search inside of it would proceed. In particular, thresholds over which an attribute, or two, are kept must be defined. The problem is more complex than in the two attributes case because we must compare different objects, small and intermediate boxes. However, recall from Proposition 3 that each attribute can be assigned the same threshold value y^T above which it is always kept in a two dimensional box. This implies that the optimal search inside the three attributes box depends on how many products are found that surpass this unique threshold.

Recall that $G(Y) = 2F(y)$ is the cumulative distribution of the sum of two attributes. Since $F(y)$ is assumed to be twice-differentiable everywhere on support $[0, \hat{y}]$, so is $G(Y)$ on support $[0, 2\hat{y}]$. Then, we can write the reservation value of this box based knowing that:

- if three attributes are found above y^T , the consumer will stop,




Figure 11: *Graphical representation of a three attribute compound box centered around $(1, 1, 1)$. Products that share two attributes with $(1, 1, 1)$ can be displayed along the edges of a cube (north for products sharing A_1, B_1 , south-west for products sharing A_1, C_1 , south east for products sharing B_1, C_1); Products that share one attribute with $(1, 1, 1)$ can be displayed along the sides of the cube (north-west for products sharing only A_1 , north-east for products sharing only B_1 , south for products sharing only C_1).*

- if two attributes are found above y^T , small boxes will be opened next until something that beats z is found,
- if one attributes are found above y^T , intermediate boxes will be opened next until something that beats \mathcal{R}_2 is found,
- if no such attribute is found, the highest will be kept and intermediate boxes will be opened as well.

Notice that, mechanically, this is the same outcome anticipated by the consumer in the two attribute case.

This directions can be used to score three attribute products just as for the two attributes one. The added difficulty compared to the smaller case lies both in the higher number of options and on the existence of several different thresholds determining the value of the box conditional on different outcomes. In particular, notice that $\bar{y} = E[Y|Y > \mathcal{R}_2]$ represents the expected value of opening intermediate boxes given the first set of realizations and, therefore, follows updating as described in the main model depending on the number of attributes found above $z + s$ but below y^T . A more detailed analysis can be found in Appendix E as well.

As before, these are the scores of all three attribute boxes before search starts. Solving the optimal search process from now on requires going through the same steps as before: when deciding what to search next, the consumer compares the reservation value of unopened compound boxes with that of past realizations and visible nested boxes. Finding attributes above certain thresholds can make the consumer backtrack to a previously discarded option, and search can be represented in terms of Pandora's optimal search policy at every step.

The discourse above hints at the major limitation of this approach: generalizing the model to products defined by an arbitrary number of attributes requires finding the equivalent to threshold y^T for all boxes of smaller size. For example, a ten attribute product would be represented by a box containing smaller boxes with unknown realizations in up to nine dimensions, each of which must be characterized. Solving the three attributes case is necessary to characterize the four attribute case, which is in turn necessary to characterize the five attributes case, and so on. A more nuanced approach might be possible; finding out is left for future research.

6.3. Purchase without inspection

Currently, it is assumed that consumers must expend a search cost to inspect any uninspected product. Since products in this environment share attributes, however, after inspecting multiple products some uninspected products would be fully revealed. If search is understood as the physical action of finding a product, this distinction does not matter. If, however, one were to interpret search as the time and effort necessary to ascertain the quality of the match of a product, it would be sensible to suggest that products uninspected but nonetheless known in their realization should not need search costs to be expended. In this extension I explore the implications of this alternative interpretation.

First, it is clear that if taking a product whose attribute have been fully independently discovered is free, the only optimal search process would be one that involves searching new attributes in pairs until the highest realization for each attribute is such that they, together surpass the value of all uninspected products. Mechanically, this implies that the reservation value $R_{i,j}$ or compound boxes should account for the fact that keeping one attribute fixed and searching along the other is never optimal. This can be achieved by modifying the way reservation values update after each observation. In particular, the lowest realization that reroutes search towards itself inside all unopened compound boxes is (without loss of generality) $A_1 > z$ rather than $A_1 > z + s$. With this change, the choice of keeping an attribute is always dominated since all products sharing an attribute with an inspected product would be contained, at zero additional cost, in all unopened compound boxes. The updating still takes place, but since the realization must not be discounted by the search cost, all realizations A_1 update future compound boxes to higher values. Therefore, all else equal, y^T , the threshold above which an attribute is always kept, would need to be higher than in the baseline model when $s < E[y]$.

The pricing game in Section 4 is also affected by this change. Remember that the price the multiproduct monopolist can impose is restricted by the search cost and by the opportunity of searching additional products. If the consumer would always search on the diagonal before selecting a combination of known attributes to keep, the monopolist would have an incentive to increase the price of all products off the diagonal to capitalize on the consumers' ability to correct his choice for free. Because consumers would need to be induced to start searching, however, this also implies that the products on the diagonal would need to become cheaper too. The change in interpretation does not affect the result qualitatively, but the mechanical change to the search process suggests that prices would be even more dispersed in equilibrium under this alternative interpretation of search costs.

6.4. Limitations and directions forward

On the Eventual Purchase Theorem Since the problem can be solved by scoring products in a way that takes into account the relevant information each carries through the expected search paths that they open up, one might wonder whether a version of the eventual purchase theorem exists in this environment. The eventual purchase theorem (henceforth, EPT), first proposed in [Armstrong \(2017\)](#) and [Choi et al. \(2018\)](#), states that the outcome of a search process to find one out of independently distributed products can be obtained through a simple statistic. In particular, the product i that is ultimately selected by a consumer will be the one with the highest statistic:

$$W_i = \min\{r_i, u_i\},$$

or, the highest minimum between reservation and match value of a product.

Obtaining a similar statistic in this environment comes with a few challenges. First, a product is kept as long as its match value surpasses that of different objects, namely the closed compound boxes and the closed reachable nested boxes. These objects are associated with different scores, one reflecting the value of the implied search paths that would follow from it, one reflecting the value of the one inspection it involves. A statistic like the one governing the EPT, then, should account for both²⁸.

Another difficulty is the threshold over which attributes are kept. In the main analysis I show that an attribute is kept as long as it has realization above y^T . However, this is only true for the first attribute kept. Once this is selected, and the consumer starts looking for a variant of the other attribute to go with it, the threshold relevant for this additional component is z rather than $y^T \geq z + s$. Hence, the statistic above represents a sufficient but not necessary condition. This change in the relevant threshold further complicates the problem.

Finally, it should be pointed out that some distributional restriction could help generate a viable statistic. Suppose that F was such that:

$$\tilde{R} < z + s < R = [2 - F(z)^2]\bar{y} + F(z)^2z,$$

or, such that no realization below $z + s$ would beat a closed compound box without the updating discussed in the main model. Intuitively, this requires s to be small enough (so that $z + s > y^T$), F to be such that \bar{y} is sufficiently larger than z , and z to be relatively high. In this case, $y \geq z + s$ would be a necessary and sufficient condition for the first attribute to be kept. Then, the relevant thresholds would be defined by $z + s$ and z . Further, since a consumer would search to discover two attributes on the diagonal and adjust along rows or columns based on the realization on the diagonal itself, the indices i, j would determine the position of a product (i, i) on the diagonal (that would be kept if and only if its highest realization was above $z + s$ and its lowest was above z) and that of a product $[i, j \neq i]$ off the diagonal (whose attribute $j > i$ would be kept if above z).

If the condition above is not satisfied, however, the condition $y > y^T$ is not necessary for a

²⁸This is clearly only true for $s < \bar{s}$: as shown above, for higher search costs the optimal search process is unaffected by correlation and, therefore, the EPT as characterized by [Armstrong \(2017\)](#) and [Choi et al. \(2018\)](#) trivially applies in this case.

product to be kept; in this case, the thresholds should account for the eventuality of an attribute being discarded and then picked again, and for two attributes below y^T to beat closed boxes without updating.

Different distributions In principle, removing the assumption of attributes following the same distribution can be accommodated. One can easily imagine a variant of the model above in which all A attributes were i.i.d and all B attributes were too, but the two sets followed a different distribution. This does not affect the analysis significantly. What is far more challenging is accounting for different distributions across different variants of the same attribute in the general framework. The reason stems from the way compound boxes are constructed: with different distributions come different reservation values z for the same search cost s , which means that the expected value of searching along one dimension is not straightforward to compute. A possible solution might be to use the EPT as characterized by [Armstrong \(2017\)](#) and [Choi et al. \(2018\)](#) to pin down said value, and the value of all other dimensions. This, too, is left open for future research.

7. 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. The framework highlights a novel interaction between pricing and optimal order of inspection in directed search: consumers have an incentive to find better matches in their search process as they learn what they like. This dictates their strategy predictably. On one hand, this allows to rethink the problem in a way that generates threshold values of searching different available options in a way reminiscent of [Weitzman \(1979\)](#)'s optimal search policy. On the other, it highlights that a multiproduct monopolist is 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. Additionally, the monopolist studied here is interested in eliciting specific search patterns, but he is indifferent regarding which product acts as starting point. 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.

The framework's predicted search patterns align well with recent evidence of spatial learning in search: [Hodgson and Lewis \(2020\)](#) reports evidence of search for digital cameras to be

characterized by a learning process consistent with the one model in this framework. Cameras are identified by their unique combination of brand, pixel, zoom, and display; the authors show, among other things, that the distance (in terms of identified attributes) between two products inspected sequentially shrinks as the search process continues. Consumers are shown to search to inspect a broader set of attributes early only to close in on their preferred alternatives in later stages, getting closer and closer to the product they ultimately choose to purchase. This pattern cannot be easily reconciled with standard search models, but is well in line with the prediction of this framework. Further, the model can more easily rationalize the pervasive tendency of consumers to retrace their steps while searching for products.

The implications of this model for recommendation systems and algorithmic pricing schemes has been addressed in an earlier section. It is worth stressing out, however, that these implications go beyond the specific market structure studied here. Coordination of menu and pricing allows a multiproduct monopolist to induce specific search paths to arise. Equivalently, one can imagine e-commerce platforms to do the same through manipulation of the options presented to consumers and the information therein. This is especially true in a world in which data on consumers' decisions, consumption and search patterns is abundant, and algorithmic pricing and recommendation systems are ever more effective at predicting human behavior.

From this perspective, 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, restricts the possible search paths that can feasibly arise. If consumers adapt their search strategy as suggested in this model and in line with the spatial learning described above, 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 e-commerce 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

- Amir, Rabah, Jim Y Jin, Gerald Pech, and Michael Tröge**, “Prices and deadweight loss in multiproduct monopoly,” *Journal of public economic theory*, 2016, 18 (3), 346–362.
- Anderson, Simon P and Régis Renault**, “Advertising content,” *American Economic Review*, 2006, 96 (1), 93–113.
- Anderson, Simon P, Maxim Engers, and Daniel Savelle**, “Ordered Search: Equilibrium and Optimum,” *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**, “Ordered consumer search,” *Journal of the European Economic Association*, 2017, 15 (5), 989–1024.
- 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.
- Bao, Renjie, Sanxi Li, and Jun Yu**, “Search with Interdependent Values,” *Available at SSRN 4135148*, 2022.
- Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat**, “Search, design, and market structure,” *American Economic Review*, 2012, 102 (2), 1140–60.
- Blander, James A and Jonathan Eaton**, “Product line rivalry,” *The American Economic Review*, 1984, 74 (3), 323–334.
- Choi, Michael, Anovia Yifan Dai, and Kyungmin Kim**, “Consumer search and price competition,” *Econometrica*, 2018, 86 (4), 1257–1281.
- De Corniere, Alexandre and Greg Taylor**, “A model of biased intermediation,” *The RAND Journal of Economics*, 2019, 50 (4), 854–882.
- Doval, Laura**, “Whether or not to open Pandora’s box,” *Journal of Economic Theory*, 2018, 175, 127–158.
- Eaton, B Curtis and Richard G Lipsey**, “The theory of market pre-emption: the persistence of excess capacity and monopoly in growing spatial markets,” *Economica*, 1979, 46 (182), 149–158.
- Garcia, Daniel and Sandro Shelegia**, “Consumer search with observational learning,” *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.
- Heidhues, Paul, Mats Köster, and Botond Köszegi**, “Steering fallible consumers,” *The Economic Journal*, 2023, 133 (652), 1430–1465.
- Hodgson, Charles and Gregory Lewis**, “You can lead a horse to water: Spatial learning and path dependence in consumer search,” *Available at SSRN 3667788*, 2020.
- Ichihashi, Shota**, “Online privacy and information disclosure by consumers,” *American Economic Review*, 2020, 110 (2), 569–595.

- Johnson, Justin P and David P Myatt**, “Multiproduct Quality Competition: Fighting Brands and Product Line Pruning.,” *Available at Oxford University Research Archive*, 2002.
- Ke, T. Tony and Song Lin**, “Informational Complementarity,” *Management Science*, 2022, 66 (8), 3699–3716.
- Konishi, Hideo and Michael T Sandfort**, “Expanding demand through price advertisement,” *International Journal of Industrial Organization*, 2002, 20 (7), 965–994.
- Lambertini, Luca and Andrea Mantovani**, “Process and product innovation by a multiproduct monopolist: a dynamic approach,” *International Journal of Industrial Organization*, 2009, 27 (4), 508–518.
- Lancaster, Kelvin J**, “A new approach to consumer theory,” *Journal of political economy*, 1966, 74 (2), 132–157.
- Lin, Ping**, “Process and product R&D by a multiproduct monopolist,” *Oxford Economic Papers*, 2004, 56 (4), 735–743.
- Mayzlin, Dina and Jiwoong Shin**, “Uninformative advertising as an invitation to search,” *Marketing science*, 2011, 30 (4), 666–685.
- McAfee, R Preston, John McMillan, and Michael D Whinston**, “Multiproduct monopoly, commodity bundling, and correlation of values,” *The Quarterly Journal of Economics*, 1989, 104 (2), 371–383.
- Mussa, Michael and Sherwin Rosen**, “Monopoly and product quality,” *Journal of Economic theory*, 1978, 18 (2), 301–317.
- Nocke, Volker and 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.
- Robert, Jacques and Dale O Stahl**, “Informative price advertising in a sequential search model,” *Econometrica: Journal of the Econometric Society*, 1993, pp. 657–686.
- 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.
- Teh, Tat-How and Julian Wright**, “Intermediation and steering: Competition in prices and commissions,” *American Economic Journal: Microeconomics*, 2022, 14 (2), 281–321.
- Weitzman, Martin L**, “Optimal search for the best alternative,” *Econometrica: 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. Simplified framework: monopoly pricing

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

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

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

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

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

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

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

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

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

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

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

If $\frac{3\alpha^2 - 2\alpha^3}{1 + 2\alpha - 2\alpha^2} \leq s < \alpha$, both p_L^D and p_L^E are feasible. Moreover, it holds $p_M = p_L^E < p_L^D$ for the whole segment. Therefore, it is sufficient to compare expected profits under p_L^E and p_L^D . Notice

that p_L^D is such that searching again after a bad first realization is not possible. In this segment:

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

Therefore, the monopolist compares:

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

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

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

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

p_L^E is selected otherwise.

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

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

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

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

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

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

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

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

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

π_H^E is optimal for:

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

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

All feasible combinations of $\alpha \in (0, 1)$ and $s \in (0, 2\alpha)$ are then accounted for when restricting the 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:

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

Which is valid for $p_L^E < 1$ or, $\alpha^2 < s$. As shown in the main text, for $s > \alpha$ the consumer has no reason to search again after finding something she likes, and indeed would lead to a lower, rather than higher, price level for $p_{1,2}$ and $p_{2,1}$. In this segment ($\alpha^2 < s < \alpha$), such prices lead to strictly higher expected profits. Indeed, when the consumer starts from (1, 1) (equivalently, (2, 2)), she only searches the more expensive product if she already knows that she likes it in some attribute. The consumer cannot start from any other product: if she starts from the more expensive product, her expected utility of search in this segment is negative since 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 (1, 2): inspecting (2, 2) leads to utility equal to two with probability α^2 , and allows to correct to (1, 2) if she learns that she likes B_2 but not A_2 , which happens with probability $\alpha(1 - \alpha)$. The expected utility along this alternate path is equal to:

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

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

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

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

The first reaction re-establishes the equilibrium: the consumer now inspects the more expensive product only if he knows it is the only product that leads to utility equal to two. 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:

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

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

The second reaction also re-establishes the equilibrium: by setting a lower price for (1, 2), the 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 (2, 2) prevents the deviation from taking place at all. Since no follow-up search in case of a bad first realization is possible without (2, 2), however, overall probability of trade decreases. Expected profits in this case are:

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

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

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

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

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

while $\bar{\pi}$ 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 (1, 1) and learned $A_1 = 1$, $B_1 = 0$. Then, she would want to inspect (1, 2). She does so as long as:

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

which implies:

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

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

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

Once again, the 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 (2, 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 $p^{u*} = p_T$, 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 $p^{u*} = p_L^D$ and $0 < s < \alpha^2$.

Two results emerge: when selecting p_T as base product and the consumer does not adapt their search strategy, this is always optimal. Second, when there is adaptation by consumer and 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 $\bar{\pi}$ and $\underline{\pi}$ when $p^* = p_T$. It holds:

$$\begin{aligned}\underline{\pi} > \pi_H^D &\iff \frac{\alpha^4 - 8\alpha^3 + 12\alpha^2 - 4\alpha}{2\alpha^3 - 6\alpha^2 + 4\alpha + 1} < s < \alpha^2 \\ \bar{\pi} > \pi_H^D &\iff \frac{\alpha^4 + 4\alpha^3 - 10\alpha^2 + 4\alpha}{2\alpha^4 - 4\alpha^3 + 4\alpha^2 - 2\alpha - 1} < s < \alpha^2\end{aligned}$$

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

B. General model: search dynamic

The value of infinite compound boxes We show here how to obtain the value of unopened compound boxes conditional on past realizations. Consider first the value of these boxes before anything is observed. All compound boxes $X_{i,j}$ contain products (i, j) , readily available, $[i, j' \neq j]$, in closed boxes, and $[i' \neq i, j]$, also in closed boxes.

Suppose the box was opened and A_i, B_j was observed. From the above, we can give reservation value $r_{i,j'}$ to all closed boxes containing products $[i, j' \neq j]$. Since all attributes are assumed be i.i.d, following Choi et al. (2018):

$$r_{i,j'} = A_i + z, \quad z : s = \int_z^{\bar{y}} (y - z) dF(y),$$

and in the same fashion:

$$r_{i',j} = B_i + z, \quad z : s = \int_z^{\bar{y}} (y - z) dF(y),$$

Unopened boxes inside $X_{i,j}$ are independent as they share no attribute by construction. Then, a consumer would search among these alternatives following Weitzman (1979)'s optimal search policy: the box with the highest reservation value is opened first, then the second, until a realized

outcome is higher than all unopened boxes. From the above:

$$r_{i,j'} > r_{i',j}, \forall \{i', j'\} \iff A_i > B_j$$

Therefore, the consumer would always want to keep the highest of the two realizations and inspect variants of the lowest one if $\min\{A_i, B_j\} < z$, and stop there otherwise. Suppose (w.l.o.g.) that $A_i > z > B_j$. The consumer keeps A_i and searches variants of B until she finds $B_j > z$ to keep. Then, in expectation, she expects to get:

$$A_i + [1 - F(z)]E[y|y > z] + F(z) \left([1 - F(z)]E[y|y > z] + F(z)(\dots \right.$$

or

$$A_i + [1 - F(z)] \sum_{k=1}^{\infty} F(z)^k E[y|y > z] = A_i + E[y|y > z].$$

The same applies if B_j is kept instead.

If $\min\{A_i, B_j\} > z$, the consumer would stop and obtain $A_i + B_j$. If $\max\{A_i, B_j\} < z$, the highest of the two would be kept, and the consumer would expect to get:

$$\max\{A_i, B_j\} + E[y|y > z].$$

These paths depend on whether the realizations of A_i, B_j are above or below z . Combining the various cases:

$$\begin{aligned} E[X_{i,j}] = & [1 - F(z)]^2 (\bar{y} + \bar{y}) \\ & + 2F(z)[1 - F(z)](\bar{y} + \bar{y}) \\ & + F(z)^2 \left(\bar{y} + \frac{1}{F(z)^2} \int_0^z y dF(y)^2 \right), \end{aligned}$$

where $\frac{1}{F(z)^2} \int_0^z y dF(y)^2$ is the expected value of the highest of two i.i.d. independent realizations, obtained through order statistics. Then, $R_{i,j} = [2 - F(z)^2] + F(z)^2 \underline{z}$, where \underline{z} is the certain equivalent of searching a box with distribution $F_{2,2}(\cdot) = F(\cdot)^2$ with support $[0, z]$.

We now want to characterize the value of closed boxes conditional on the past realizations. Assuming products are inspected in increasing order of their index, inside any unopened box $X_{i,j}$ we find the product, all unopened nested boxes representing products with higher indices in one attribute, and products that have been already discovered, with reservation value:

$$r(i, j'' < j) = A_i + B_{j''} - s, \quad r(i'' < i, j) = A_{i''} + B_j - s,$$

respectively. Clearly, any of these boxes beats the unopened ones if:

$$A_{i''} + B_j - s > B_j + z \iff A_{i''} > z + s,$$

and equivalently for B .

When one such attribute is found, the search inside future compound boxes must be updated, since consumers will not opened unknown boxes in that dimension but go back. If only one such

attribute ($A_{i''}$, w.l.o.g.) is found in previous searches:

$$\begin{aligned}\tilde{R}_{i,j} = & [1 - F(z)][1 - F(A_{i''} - s)](E[y|y \geq A_{i''} - s] + \bar{y}) + \\ & [1 - F(z)][F(A_{i''} - s)](A_{i''} - s + \bar{y}) + \\ & [F(z)][1 - F(A_{i''} - s - (z - \underline{y}))](E[y|y \geq A_{i''} - s - (z - \underline{y})] + \bar{y}) + \\ & [F(z)][F(A_{i''} - s - (z - \underline{y}))](A_{i''} - s + \underline{z}).\end{aligned}$$

The first two lines refer to case in which $B_j > z$ and either $A_i > A_{i''} - s$, and is therefore kept, or $A_i < A_{i''} - s$, and is discarded in favor of product $[i'', j]$. The second two refer to the case in which $B_j < z$ and either A_i is kept and search continues (if $A_i > A_{i''} - s - (z - B_j)$), or discarded in favor of $[i'', j]$ once again. The above reflects expected values before the box is opened just as in the first value found. Notice that this value is only relevant conditional on finding an attribute $A > z + s$: the score above is valid because, at the time of relevance, A_i is known.

Finally, if two attributes $A_{i''} > z + s$, $B_{j''} > z + s$ are found, they both beat the respective sets of closed boxes by construction. Following the same steps we obtain:

$$\begin{aligned}\tilde{R}_{i,j}^{A,B} = & [1 - F(B_{j''} - s)][1 - F(A_{i''} - s)](E[y|y \geq A_{i''} - s] + E[y|y \geq B_{j''} - s]) + \\ & [1 - F(B_{j''} - s)][F(A_{i''} - s)](A_{i''} - s + E[y|y \geq B_{j''} - s]) + \\ & [F(B_{j''} - s)][1 - F(A_{i''} - s)](E[y|y \geq A_{i''} - s] + B_{j''} - s) + \\ & [F(B_{j''} - s)][F(A_{i''} - s)](\max\{A_{i''}, B_{j''}\} - s + \tilde{z}),\end{aligned}$$

where \tilde{z} solves:

$$s = \int_{\tilde{z}}^z \left(\frac{y}{F(\min\{A_{i''}, B_{j''}\} - s)} - \tilde{z} \right) dF(y),$$

since, in expectation, the highest between $A_{i''}$ and $B_{j''}$ is expected to be kept if $\max\{A_i, B_j\} < z$.

Notice that at any given time either no, one, or two attributes above $z + s$ can have been found, and that these conditions are mutually exclusive. Therefore, at any given time, only one of these values are relevant for all boxes. If no such attribute is found, search in unopened boxes is unchanged and, therefore, the attributes become irrelevant after being discarded; this follows from the infinite attributes assumption, as all unopened boxes will be of infinite size no matter how many boxes were opened beforehand. Moreover, notice that all closed compound boxes contain exactly one product characterized by $A_{i''}$ and one characterized by $B_{i''}$, so all unopened boxes are affected in the same way, and update all in the same way at the same time. Finally, notice that if more than one variant for the same attribute is found above $z + s$, the highest of them beats all others in the follow-up searches and is, therefore, the only relevant one.

Threshold result Proposition 3 states that there exists a unique value $y^T = \max\{z + s, \hat{y}\} \in (z + s, \hat{y})$ such that an attribute found above it is always kept. In particular, this value is such that:

$$y^T + z \geq \tilde{R}_{i,j}, \quad \hat{y} + z = \tilde{R}_{i,j},$$

for all unopened boxes. Notice that $\tilde{R}_{i,j}$ is the relevant one to be considered since it is the value accounting for the attribute rerouting search towards itself in subsequent boxes. Notice also that, because of it, the condition in Proposition 3 is sufficient but not necessary, as a lower realization could be kept if found, for example, in an optimally inspected nested box (in which case the lower threshold would be $z < z + s$).

To prove the result, we must study the shape of $\tilde{R}_{i,j}$ and $z + s$ separately. Consider first $z + s$: since z solves:

$$s = \int_z^{\hat{y}} (y - z) dF(y),$$

and since we assumed F to be well-behaved, z is continuous and decreasing in s . In particular, it can be shown that z is convex in s and therefore decreases faster than s increases for $s \in [0, E(y)]$. $z + s$, then, is decreasing and such that:

$$\lim_{s \rightarrow 0} z + s = \hat{y}, \quad \lim_{s \rightarrow E(y)} z + s = [E(y)].$$

Next, we must study how the attribute that re-routs search towards itself affects $\tilde{R}_{i,j}$ of unopened boxes. Remember that, after finding $A_{i''} > z + s$, it holds:

$$\begin{aligned} \tilde{R}_{i,j} = & [1 - F(z)][1 - F(A_{i''} - s)](E[y|y \geq A_{i''} - s] + \bar{y}) + \\ & [1 - F(z)][F(A_{i''} - s)](A_{i''} - s + \bar{y}) + \\ & [F(z)][1 - F(A_{i''} - s - (z - \underline{y}))](E[y|y \geq A_{i''} - s - (z - \underline{y})] + \bar{y}) + \\ & [F(z)][F(A_{i''} - s - (z - \underline{y}))](A_{i''} - s + \underline{z}). \end{aligned}$$

First, it's clear that $\tilde{R}_{i,j}$ is increasing in the realization $A_{i''} > z + s$: indeed, for all realizations $A_{i''} > z + s$, the value of searching the next compound box in the dimension A is a linear combination of $A_{i''} > z + s$ and the expected value of a realization strictly above $A_{i''} > z + s$, both discounted by the search cost associated with backtracking. Further, it can be shown that:

$$\lim_{s \rightarrow 0} \tilde{R}_{i,j} < 2\hat{y}, \quad \lim_{s \rightarrow E(y)} \tilde{R}_{i,j} > 2\bar{y}.$$

To show the first limit result, notice that $\lim_{s \rightarrow 0} z = \hat{y}$ (thus, $1 - F(z) = 0$), and, therefore, $\lim_{s \rightarrow 0} \bar{y} = \hat{y}$. Further, $A_{i''} > z + s$ implies $A_{i''} = \hat{y}$. Substituting all of the above in the expression for $\tilde{R}_{i,j}$, defining $\mu \equiv E[y]$:

$$\begin{aligned} \tilde{R}_{i,j}|_{s \rightarrow 0} = & 0 + 0 \\ & [1 - F(\mu)](E[y|y \geq \mu] + \hat{y}) + \\ & [F(\mu)](2\hat{y}) < 2\hat{y}. \end{aligned}$$

The result follows: if an attribute $A_{i''} > z + s$ is found, the following nested boxes will have score $\lim_{s \rightarrow 0} A_{i''} + z = 2\hat{y}$, which beats $\tilde{R}_{i,j}$ of all unopened boxes.

For the second result, notice that as s goes to $[E(y)]$, z converges to zero and, at the limit,

$F(z) = 0$, and $\bar{y} = \mu$. Further, $A_{i''} > z + s$ can take any value above s . Substituting:

$$\begin{aligned}\tilde{R}_{i,j}|_{s \rightarrow \tilde{y}} &= [1 - F(A_{i''} - s)](\mu + E[y|y > \mu]) \\ &\quad + F(A_{i''} - s)(\mu + A_i - s) \\ &\quad + 0 + 0 > 2\mu > \mu + z = \mu = z + s.\end{aligned}$$

As s increases, then, $\tilde{y} : \tilde{y} = \tilde{R}_{i,j} - z$ increases. For s low enough, $\tilde{y} < z + s$ and, therefore, is irrelevant since it does not affect the value of the boxes that follow. Then, any realization $y > \max\{z + s, \tilde{y}\}$ is always kept.

Search dynamic We want to show that the structure above can be used to obtain the optimal search policy in the spirit of [Weitzman \(1979\)](#). First, notice that by scoring unopened boxes as above, all follow-up optimal searches are accounted for and contribute to the score of the box overall. In this sense, $\mathcal{R} \in \{R, \tilde{R}^y, \tilde{R}^{A,B}\}$ reflects the value of searching a product and the value of the information this inspection implies in terms of optimal follow-ups given only the information learned with the inspection itself.

When a compound box is opened, two things happen: first, the relevant attributes are observed; second, all nested boxes inside become reachable in their “myopic” form. On the grid, these nested boxes substitute the compound boxes sharing an attribute with the opened one. This is true as long as once the consumer starts opening nested boxes, she never wants to abandon the attribute she elected to keep as per [Corollary 1](#):

Proof. The proof of [Corollary 1](#) is as follows: 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. Therefore, if such a product is inspected, either its realization is low, and the consumer makes the same choice again, or is high, and the consumer cannot have the incentive to inspect anything else.

From [Proposition 3](#), after inspecting (i, j) , A_i is kept and $[i, j + 1]$ is inspected next if:

$$A_i + z > \max\{A_i + B_j, z + B_j, \mathcal{R}_{i+1,j+1}\}.$$

Since $z > B_j$, $\mathcal{R}_{i+1,j+1} = \tilde{R}_{i+1,j+1}^A$, that is, A_i reroutes search towards itself in all unopened compound boxes and $X_{i+1,j+1}$ in particular.

Suppose the consumer inspected $[i, j + 1]$: B_{j+1} is now known, and A_{i+1} is unknown. If $B_{j+1} > z$, the consumer would always stop there since $A_i + B_{j+1} > A_i + z$. Suppose that is not the case, and that the consumer found it optimal to inspect something other than $[i, j + 2]$.

Clearly, inspecting $[i + 1, j]$ cannot be optimal by standard [Weitzman \(1979\)](#)’s logic, since $z + B_j < A_i + z$. Since after inspection of $[i, j + 1]$ $B_{j+1} < z$ is known, the same must apply to $[i + 1, j + 1]$. Finally, all other unopened compound boxes are unaffected, since $B < z < z + s$ cannot trigger an update. Therefore, $[i, j + 2]$ is the only possible optimal search.

The proof does not rely on products being characterized by two attributes. Each individual

attribute, if kept after a search, must be strictly better than the unopened compound boxes that follow. Since the argument applies to each attribute in isolation, it is robust to an environment with a higher number of attributes defining products. ■

Backtracking can only be possible after discovering multiple realizations of a product. The optimal search path with infinite product is instructed by comparison of:

- $u_{i,j}$, the readily available product in the latest compound box opened,
- $r_{i,j'}$, the nested boxes that maintain A_i fixed, which if selected rules out any new attributes A being ever inspected,
- $r_{i,j'}$, the nested boxes that maintain B_j fixed, which if selected rules out any new attributes B being ever inspected,
- $\mathcal{R}_{i,j'}$, the unopened compound boxes, the score of which reflects optimal search that follows depending on the realization.

Assuming once again that attributes are inspected in increasing order of their index, it is clear that only compound boxes $X_{i,i}$ can be optimally opened, while all nested boxes that can be opened have value $r_{i,j}$, $i \neq j$. Opening a compound box informs the consumer of $u_{i,i}$, which can be compared with both kinds of unopened boxes. if $u_{i,i} > \mathcal{R}_{i+1,i+1}$, either the product is kept or nested boxes are opened. Both can happen, depending on the individual realizations of A_i and B_i .

$u_{i,i} < \mathcal{R}_{i+1,i+1}$, instead, the comparison is between $\max\{r_{i,i+1}, r_{i+1,i}\}$ and $\mathcal{R}_{i+1,i+1}$. If no individual attribute is high enough to beat, in expectation, $\mathcal{R}_{i+1,i+1}$, another compound box should be optimally opened, and the same inequalities dictate the next. Otherwise, one attribute is kept and nested boxes are opened, and no other variant of the same attribute can be discovered as per Corollary 1. This proves the result in Proposition 4.

High search costs. We want to show that for $\bar{s} < s \leq E[Y]$, where $\bar{s} = z_M$, the optimal search process is equivalent to a myopic search process as per Corollary 2.

First, it is straightforward to show that a consumer would not rationally inspect two products sharing an attribute. Suppose by contradiction that that was the case. Without loss of generality, suppose the consumer inspected product (i, j) and then chose to inspect $[i, j + 1]$. Then, it must hold:

$$A_i + z > A_i + B_j.$$

Since $B_j \geq 0$ and since $s > E[y] \rightarrow z < 0$, this cannot happen. It follows that the score of any compound box $X_{i,j}$ depends only on product (i, j) ; $R_{i,j}$, then, is the reservation value of a box containing an unknown realization from distribution $G(Y)$, which can be identified by the value z_M that solves:

$$s = \int_{z_M}^{\hat{Y}} (Y - z_M) dG(Y).$$

Next, we show that there exist a value $\bar{s} \in (E[y], E[Y])$ such that:

$$\bar{s} = \int_{\bar{s}}^{\hat{Y}} (Y - \bar{s}) dG(Y).$$

In particular we want to show that $\bar{s} > E[y]$ and $\bar{s} < E[Y]$. The latter is true by definition of z_M . We prove the former by contradiction. Suppose $\bar{s} = E[y]$. Then, by definition of z_M , $\bar{s} = E[y]$ implies:

$$E[y] = -E[y] + E[y] \int_0^{E[y]} dG(Y) + \int_{E[y]}^{\hat{Y}} Y dG(Y).$$

Since it holds:

$$E[y|y < E[y]] \int_0^{E[y]} dG(Y) < E[y] \int_0^{E[y]} dG(Y),$$

and, by linearity of $E[\cdot]$:

$$\int_0^{E[y]} Y dG(Y) + \int_{E[y]}^{\hat{Y}} Y dG(Y) = E[Y] = 2E[y],$$

it must hold:

$$E[y] < -E[y] + E[y] \int_0^{E[y]} dG(Y) + \int_{E[y]}^{\hat{Y}} Y dG(Y).$$

Therefore, if $s = E[y]$, it holds $z_M > s$. Further, since z_M is decreasing in s , there exist a unique value $\bar{s} = z_M$.

Suppose $s \in (E[y], \bar{s})$: it can be shown that in this case updating does affect the value of unopened compound boxes. Consider the updating process of compound boxes described above. A consumer would ever backtrack to a previously discovered attribute A_i if and only if $A_i - s > 0$. Suppose inspection of (i, i) leads to realization $A_i \leq s > B_i \geq 0$. Then, expected gain of inspecting $[i + 1, i + 1]$ is:

$$E[u_{i,j}|A_i > s > B_i] = E[B_{i+1}] + [1 - F(A_i - s)]E[y|y > A_i - s] + G(A_i - s)(A_i - s) - s.$$

Suppose first $A_i = s$, $B_i = 0$. (i, i) beats the next unopened compound box if it holds:

$$s > E[A_{i+1}] + E[B_{i+1}] - s = \int_0^{\hat{Y}} Y dG(Y) - s.$$

Notice that the rhs has certain equivalent equal to $z_M > s$. Therefore, $y \geq E[y]$ is not a sufficient condition for y to be kept.

Suppose now $A_i = z_M > s$, $B_i = 0$. (i, i) beats the next unopened compound box if it holds:

$$z_M > E[B_{i+1}] + [1 - G(A_i - s)]E[y|y > A_i - s] + G(A_i - s)(A_i - s) - s = z_M + G(A_i - s)(A_i - s - E[y < A_i - s]),$$

which is again a contradiction. The updating increases the value of all subsequent unopened compound boxes when $y > s$ is found.

Finally, suppose $A_i = \hat{y}$, $B_i = 0$. (i, i) beats the next unopened compound box if it holds:

$$\hat{y} > E[B_{i+1}] + (\hat{y} - s) - s,$$

which is satisfied since $s > E[y]$. Therefore, there exist a threshold value $y^T = z_M + G(y^T - s)(y^T - s - E[y|y < y^T - s])$ which represents a unique sufficient but not necessary condition for y to be kept.

We can apply the same argument to prove that the opposite is true for $s > \bar{s}$. In particular, since this condition implies $s > z_M$, $y \geq s$ is a sufficient condition for attribute y to be kept. Then, $y \geq z_M$ is necessary (by definition of z_M) and sufficient. For $s > \bar{s}$ correlation is immaterial, and the optimal search process is as described in Corollary 2.

C. General model: monopoly pricing

The proof of Proposition 5 comes in three steps. First, I show that if a non-uniform equilibrium price vector exists, it must be such that lower uniform prices are set for exactly one product characterizing all attributes, and higher uniform prices are set for all other products. Next, I show that for all such price vector there exist an equivalent uniform vector that generates the same probability of trade but leads to higher expected profits. Finally, I show that the formulation of p^* as per Proposition 5 is indeed optimal.

Optimal differentiated price vector Suppose the monopolist wanted to set differential prices for his infinite products. First, it is obvious that at least one product must be priced differently than all others. For notational clarity, I define $p_1 < p_2 < p_3$ as a set of three price levels. I show that any optimal differential price vector must be such that a set of products sharing no attributes with each other must be priced at p_1 and all other products must be priced at either p_2 or p_3 , but there cannot be any vector with more than two price levels.

First suppose that more than one product sharing an attribute A_i has price set at p_1 . The geometry of the product space implies that there must be one attribute B_j for which the same applies. For example, if $(1, 1)$ and $(1, 2)$ were priced at p_1 , $(1, 2)$ and $(2, 2)$ would also need to be. Then, the consumer would optimally start her search process from $(1, 2)$ since compound box $X_{1,2}$ contains the most cheap products. If the consumer then wanted to open a new compound box, she would optimally select $X_{3,3}$ and proceed along the diagonal.

If p_1 is such that the consumer would want to open $X_{1,2}$ but not $X_{3,3}$ without updating, the monopolist would have the incentive to set a lower p_1 to all products on the diagonal and increase the price of $(1, 2)$; on the other hand, if the consumer is willing to open $X_{3,3}$ without updating, then $p_{1,2} = p_1$ implies that with positive probability the consumer will choose to keep either A_1 or B_2 and purchase $(1, 1)$ or $(2, 2)$ at a lower price that he would have been willing to. Therefore, the monopolist would have the incentive to increase $p_{1,2}$ to re-establish the canonical order of search. This intuition extends to any number $n > 1$ of products for each attribute, and to all attributes. Therefore, at most one product per attribute can be optimally set to be cheaper than the others.

Suppose now that a strict subset of attributes has all associated products priced at either p_1 or

p_2 , while all other attributes follow the pricing detailed above. If the selected price is p_1 , all products with such attributes are cheaper than all others, and are therefore more valuable to the consumer. If the consumer is willing to exhaust these products and still inspect the attributes with differentiated prices, with positive probability the monopolist sells at a lower price than the consumer was willing to pay. If the selected price is p_2 , all such attributes would be pushed to the end of the search order and never reached since the consumer has infinite better alternatives available²⁹.

Finally, suppose that exactly one product per attribute is priced at p_1 and all others are priced at either p_2 or p_3 . Suppose first that a finite subset of attributes has products priced at either p_1 or p_2 and all other attributes have products priced at either p_1 or p_3 . A consumer that optimally decides to start searching will search first the compound box or boxes in which the most cheap products can be found. If he is willing to keep searching the boxes until only the ones with the highest number of expensive products and stop without updating, having the latter group cannot be optimal, and all products should belong to the former group. If the consumer is still interested in searching, instead, all products should belong to the latter group.

Suppose now that all attributes are such that one product is priced at p_1 , a finite subset of products is priced at p_2 and all others are priced at p_3 . If the consumer optimally elected to keep an attribute after inspecting a product priced at p_1 , she would select to inspect the ones priced at p_2 first. If after exhausting them she would stop, all other products should also have been priced at p_2 . Otherwise, all products should have been priced at p_3 . The result immediately extends to any number of price levels larger than two. The result follows.

Optimality of uniform prices Next, I show that for any vector of differential prices structured as above, there exist an uniform price vector that preserves probability of trade and returns strictly higher expected profit. As discussed in the main text (and detailed in the next part of the proof), probability of trade conditional on the consumer starting to search depends on the probability of finding realizations such that the resulting updating of unopened compound boxes makes the consumer stop searching and not purchase anything. The highest uniform price is such that:

$$R_{i,j}(\mathbf{p}^{unif}) = [2 - F(z)^2]\bar{y} + F(z)^2\bar{z} - p^{unif} = 0, \quad \forall(i, j).$$

Henceforth, I refer to $R_{i,j}(\mathbf{p}^{unif})$ and $R_{i,j}$ as the initial value of compound box $X_{i,j}$ with and without prices. For any uniform price p^{unif} , probability of trade is equal to:

$$q^{unif} = \left[1 - Pr \left(y_{2,2} \in (z + s, y(p^{unif})) \wedge y_{1,2} < \tilde{R}_{i,j}^y|_{y=y(p^{unif})} - y_{2,2} \right) \right],$$

where $y(p^{unif}) : p^{unif} = \tilde{R}_{i,j}^y|_{y=y(p^{unif})}$ is the realization that, after triggering the update, generates an updated value of compound boxes exactly equal to p , while $y_{1,2}$, $y_{2,2}$ are the first and second order statistics of two of an i.i.d. random variable distributed according to $F(y)$. It's clear that this probability is decreasing in p^{unif} : the higher p^{unif} is, the wider the range of realizations that make the consumer stop searching. Then, we can define a generic price $p^{unif} = R_{i,j} - k(p^{unif})$, for some positive k , to define the expected probability of trade.

²⁹This implies that the pricing scheme with infinite products has infinite payoff-equivalent equilibria in which the monopolist sets a high price for all products defined by a finite subset of attributes which are never reached.

To obtain the expected probability of trade in the differential case, we repeat the same steps for a price vector such that $p_{i,i} = p$, $p_{i,j \neq i} = p + \delta$. Define $R_\delta = [2 - F(z - \delta)]\bar{y}_\delta + F(z - \delta)^2 \underline{z}_\delta$ the initial value of an unopened compound box with differential prices as obtained in the main text. It holds:

$$R_{i,j}(\mathbf{p}^{\text{diff}}) = R_\delta - \left(p + \delta \left[1 - [1 - F(z - \delta)]^2 \right] \right).$$

For a differential price vector to have the same probability of trade as some p^{unif} , the two price vectors need to make the consumer stop searching for the same realizations. Therefore, defining probability defining q^{diff} the probability of trade under differential prices, it holds:

$$q^{\text{unif}} = q^{\text{diff}} \iff R_\delta - \left(p + \delta \left[1 - [1 - F(z - \delta)]^2 \right] \right) = R_{i,j} - p^{\text{unif}} = k(p^{\text{unif}}).$$

Notice that $p + \delta \left[1 - [1 - F(z - \delta)]^2 \right]$ is the expected profit of the seller if a purchase takes place: with probability $[1 - F(z - \delta)]^2$, the consumer purchases a product along the diagonal at p , and with the remaining probability she purchases a product priced at $p + \delta$. Equivalently, expected profit of the seller if the purchase takes place under uniform prices is trivially equal to p^{unif} . The condition above is equivalent to:

$$p + \delta \left[1 - [1 - F(z - \delta)]^2 \right] = R_\delta - k.$$

To prove the claim, I show that $R_\delta < R_{i,j}$ whenever $\delta > 0$. Remember that it holds:

$$\begin{aligned} R_\delta &= [1 - F(z - \delta)]^2 (2\bar{y}_\delta) \\ &\quad + 2F(z - \delta)[1 - F(z - \delta)](2\bar{y}_\delta) \\ &\quad + F(z - \delta)^2 (\bar{y}_\delta + \underline{z}_\delta). \end{aligned}$$

This is equivalent to the value of opening a compound boxes $R_{i,j}$ with some $s' > s$ under uniform prices. In particular, define:

$$s' : s' = \int_{z-\delta}^{\hat{y}} (y - (z - \delta)) dF(y).$$

Since reservation value of a single unknown attribute is decreasing in the search cost, if $\delta > 0$ it must be that $s' > s$ since $z - \delta < z$. Since $R_{i,j}$ is decreasing in s by construction, this implies that for any $\delta > 0$, there exists a uniform price vector that generates the same probability of trade $q^{\text{unif}} = q^{\text{diff}} = q$ but such that expected profit conditional on trade taking place is:

$$p^{\text{unif}} = R_{i,j}|_s - k(p^{\text{unif}}) > R_{i,j}|_{s' > s} = R_\delta = p + \delta \left[1 - [1 - F(z - \delta)]^2 \right],$$

which proves the result.

Optimal uniform prices vector Finally, I show that the formulation of p^* as per Proposition 5 is indeed optimal. As explained above, each price level pins down the probability of trade. The highest possible optimal price is:

$$\bar{p} = R_{i,j},$$

since if $p > R_{i,j}$, the consumer would not start searching. To find the lowest possible optimal price, we find the price that leads to probability of trade equal to one. Since trade can only not happen if the update of a compound box is such that the consumer does not purchase anything and stops searching, we find the lowest possible realization after an update.

Suppose $A_i = z + s$, $B_j < z + s$ is found after opening $X_{i,j}$. Then:

$$\begin{aligned}\tilde{R}_{i',j'}^{y=z+s} &= [1 - F(z)]^2(2\bar{y}) + \\ &\quad [1 - F(z)][F(z)](z + \bar{y}) + \\ &\quad [F(z)][1 - F(\underline{y})](E[y|y \geq \underline{y}] + \bar{y}) + \\ &\quad [F(z)][F(\underline{y})](z + \underline{z}).\end{aligned}$$

Suppose now $A_i = z + s$, $B_j = z + s$ is found after opening $X_{i,j}$. Then:

$$\begin{aligned}\tilde{R}_{i',j'}^{A,B=z+s} &= [1 - F(z)]^2(2\bar{y}) + \\ &\quad 2[1 - F(z)][F(z)](z + \bar{y}) + \\ &\quad [F(z)]^2(z + \underline{z}).\end{aligned}$$

While $\underline{z} > \underline{z}$ always holds, the relationship between $E[y|y \geq \underline{y}]$ and z depends on the distribution. If p is equal to the minimum of this too expression, pinned down by $F(\underline{y})$, then, p is the lowest candidate optimal price, since it returns probability of trade equal to one:

$$p = \underline{R} \equiv \min\{\tilde{R}_{i,j}^y|_{y=z+s}, \tilde{R}_{i,j}^{A,B}|_{A=B=z+s}\}.$$

For any price $p \in (\underline{R}, R_{i,j}]$, probability of trade is lower than one and can be pinned down precisely by the price. Following the main text, define with:

$$p_\lambda = \lambda R_{i,j} + (1 - \lambda)\underline{R}, \quad \lambda \in [0, 1].$$

a generic candidate equilibrium price, and:

$$y = y_\lambda : \tilde{R}_{i,j}^y|_{y=y_\lambda} = p_\lambda.$$

Then, probability of trade q is:

$$\left[1 - Pr\left(y_{2,2} \in (z + s, y_\lambda) \wedge y_{1,2} < \tilde{R}_{i,j}^y|_{y=y_\lambda} - y_{2,2}\right)\right].$$

This follows from the search dynamic. Suppose a box $X_{i,j}$ is opened and (W.L.O.G.) $A_i \in (z + s, y^\lambda) > B_i$. By construction, the next closed compound box has now negative value, so the consumer would not inspect it. We want to see when the consumer also decides to stop searching. First, it is clear that for this to be relevant, this needs to be the first update that made the consumer update the value of unopened boxes, since otherwise he would have already stopped and purchased something. Therefore, internal consistency requires all previously inspected boxes to be below p_λ . The consumer then must choose between:

- purchasing the last opened box, which happens only if $A_i + B_j \geq p_\lambda = \tilde{R}_{i,j}^y|_{y=y_\lambda}$,
- inspecting nested boxes optimally, which now have score $r_{i,j'} = A_i + z - p_\lambda$.

The latter can be shown to never be a relevant option by contradiction. Recall that the consumer is indifferent between keeping attribute A_i and opening the next compound box if $A_i = \tilde{y} = \tilde{R}_{i,j}^{\tilde{y}} - z$. Since at most $A_i = y_\lambda - \epsilon$ for $\epsilon > 0$ arbitrarily small, the highest possible updated value $\tilde{R}_{i,j}^{\tilde{y}} = p_\lambda - \epsilon'$ for $\epsilon' > 0$ arbitrarily small.

If $A_i = \tilde{y}$, $r_{i,j'} = \tilde{R}_{i,j}^{\tilde{y}} - p_\lambda$, but $\tilde{R}_{i,j}^{\tilde{y}} - p_\lambda > 0$ requires $A_i = \tilde{y} > y_\lambda$. Therefore, inspecting nested boxes is never an option if a realization that would make the consumer stop opening compound boxes is found. Then, the consumer stops without purchasing anything if and only if the last product (i, j) inspected is such that $B_j < A_i \in (z + s, y^\lambda)$ and $A_i + B_j < p_\lambda = \tilde{R}_{i,j}^y|_{y=y_\lambda}$, which is the expression for q obtained above.

Then, expected profit as a function of λ is:

$$\pi_\lambda = p_\lambda \left[1 - Pr \left(y_{2,2} \in (z + s, y_\lambda) \wedge y_{1,2} < \tilde{R}_{i,j}^y|_{y=y_\lambda} - y_{2,2} \right) \right].$$

Since $\frac{\partial p_\lambda}{\partial \lambda} > 0$, $\frac{\partial q_\lambda}{\partial \lambda} < 0$, the choice of λ determines the the expected profit and, therefore:

$$p^* = p_\lambda : \lambda = \arg \max_{\lambda \in [0,1]} (\pi_\lambda)$$

is the optimal uniform price given distribution $F(y)$ and search cost s .

D. Extension - Finite number of attributes

We want to show that the logic of the infinite attribute case is unaffected in the finite attribute one. Suppose products were defined by two attributes. Suppose further that A came in n variants, and B in m . Following the same intuition employed for the infinite attributes case, we can build compound boxes accounting for the number of viable products that share attributes with each.

Suppose the first box $X_{1,1}$ is opened. Inside there are $(1, 1)$, $n - 1$ nested products characterized by B_1 , and $m - 1$ nested products characterized by A_1 . As before, it is clear that the highest between A_1 , B_1 would be kept. W.l.o.g., suppose $A_1 > B_1$. Then, the consumer has $m - 1$ possible products to inspect before running out. Suppose $A_1 > B_1 > z$: then, the consumer would stop. Suppose instead $A_1 > z > B_1$, the consumer expects to find:

$$A_1 + [1 - F(z)]\bar{y} + [1 - F(z)] \sum_{k=1}^{n-1} \bar{y} + F(z)^n \left(\frac{1}{F(z)^n} \int_0^z y dF_{n,n}(y) \right).$$

Since at every step she will either find something above z or keep searching, and if she finds nothing she expects to keep the highest of n realizations below z . The same is true if $B_1 > z > A_1$, and if $\max\{A_1, B_1\} < z$, in which case the highest would be kept. In this last case, however, it is possible for the consumer to exhaust all $(1, j)$ products and then optimally decide to inspect $(i, 1)$ products instead. As this is the more involved case, let us focus on that in isolation.

We want to back out the value of searching inside a compound box conditional on the first

realization inside of it, $u_{1,1} = A_1 + B_1$ to be such that $\max\{A_1, B_1\} < z$. Suppose first $A_1 > B_1$. Then, the consumer would search through the $m - 1$ nested products characterized by A_1 in order. If she ever finds an attribute B_j that beats z , she would stop. If she does not, after exhausting all $m - 1$ nested products, she might decide to start inspecting the $n - 1$ nested products characterized by B_1 if:

$$\max_{j \in \{2, \dots, m\}} [B_j] < z - (A_1 - B_1).$$

This implies that with every new inspection keeping A_1 fixed, with probability $[1 - F(z)]$ the consumer would find $B_j > z$ and stop, with probability $[F(z) - F(z - (A_1 - B_1))]$ the consumer would keep searching but not drop A_1 even if nothing above z is found, and with probability $F(z - (A_1 - B_1))$ the consumer would still decide to inspect $(i, 1)$ after exhausting products $(1, j)$. Therefore, she expects to find:

$$\begin{aligned} E[U_{1,1}|z > A_1 > B_1] &= [1 - F(z)] \left(\sum_{k=0}^{m-2} \sum_{h=0}^k [F(z) - F(z - (A_1 - B_1))]^h F(z)^{k-h} \right) \bar{y} \\ &+ \left(\sum_{k=0}^{m-2} [F(z) - F(z - (A_1 - B_1))]^{m-k-1} F(z)^k \right) \max\{\bar{y}_{(1,2)}, \bar{y}_{(m-k-1, m-k-1)}\} \\ &+ F(z - (A_1 - B_1))^{m-1} \left([1 - F(z)] \sum_{k=0}^{n-1} F(z)^k (\bar{y} - (\bar{y}_{(2,2)} - \bar{y}_{(1,2)})) \right. \\ &\left. + F(z)^{n-1} \max\{\bar{y}_{(2,2)} + \bar{y}_{(m-1, m-1)}, \bar{y}_{(1,2)} + \bar{y}_{(n-1, n-1)}\} \right), \end{aligned}$$

where $\bar{y}_{(\cdot, \cdot)}$ is the expected value of the order statistic conditional over the relevant constrained support. In the same fashion, we can obtain $E[U_{1,1}|z > B_1 > A_1]$, which is different as long as $n \neq m$.

Combining all scenarios, expected utility of opening the compound box $X_{1,1}$ is:

$$\begin{aligned} E[U_{1,1}] &= [1 - F(z)]^2 (\bar{y} + \bar{y}) \\ &+ [1 - F(z)] F(z) \bar{y} + [1 - F(z)] \sum_{k=1}^{n-1} F(z)^k \bar{y} + F(z)^n \left(\frac{1}{F(z)^n} \int_0^z y dF_{n,n}(y) \right) \\ &+ [1 - F(z)] F(z) \bar{y} + [1 - F(z)] \sum_{k=1}^{m-1} F(z)^k \bar{y} + F(z)^m \left(\frac{1}{F(z)^m} \int_0^z y dF_{m,m}(y) \right) \\ &+ F(z)^2 \left(\bar{z}_2 + \frac{1}{2} E[U_{1,1}|z > A_1 > B_1] + \frac{1}{2} E[U_{1,1}|z > B_1 > A_1] \right), \end{aligned}$$

where:

$$\bar{z}_2 : s = \int_{\bar{z}_2}^z \left(\frac{y}{F(z)^2} - \bar{z}_2 \right) dF_{2,2}(y).$$

Suppose now that the box has been opened and the consumer found two attributes below $z + s$. Then, when the consumer decides to open the next compound box, she knows she will not go back to either A_1 or B_1 . In practice, these products do not matter for the optimal search that would occur inside the second box. Then, the second box has value equal to $R_{1,1}$ if it had $n - 1$ variants of A and $m - 1$ variants of B inside of it. On the other hand, if one or two attributes are found above $z + s$, they re-route search towards themselves in the same way they would have

done in the infinite attribute case. In this case, the same updating takes place. Finally, notice that if $n = m$, it would be possible to open $n - 1$ boxes and wanting to discard them all. In this case, the reservation value of the last compound box would just be the value implied by distribution $G(Y) = F(y) + F(y)$ since attributes are i.i.d.

At every step, the consumer can choose one of many options. Suppose some compound boxes on the diagonal³⁰ were still unopened, after opening box (i, i) , the consumer chooses between:

- $\max_{j \in \{1, 2, \dots, i\}} [u_{j,j}]$, the highest realized value,
- $\max\{r_{i,i+1}, r_{i+1,i}\}$, the value of closed nested boxes,
- $\mathcal{R}_{i+1,i+1}$, the value of the next compound box,
- $\max_{j \neq j'} [u_{j,j'}] - s$, the highest known but uninspected product for some $j \neq j'$, $j, j' \leq i$.

Compared to the infinite case, two differences emerge: first, $\mathcal{R}_{i+1,i+1} = R_{i+1,i+1}$ becomes smaller and smaller as the consumer proceeds on the diagonal if no attributes $y > z + s$ are found. If one is, the box still shrinks in one dimension (the one for which such attribute is not found). If two are, boxes update to $\tilde{R}_{i,j}^{A,B}$ and become all the same “size”, since the consumer would not explore now individual attributes in either direction. If boxes do shrink, a box explored earlier has more value than one explored late, in expectations. Therefore, some realizations might be too low to beat a subsequent large compound box, but large enough to beat a subsequent small box. Notice that this was not the case in the infinite products case: a product could not be abandoned and then recovered unless it rerouted search towards itself since, if he did not, search inside the box would not have changed.

Thus, one must consider the outside option of the consumer. While with infinite products no combination of attributes could be optimal to keep unless rerouting took place, it is possible now since boxes shrink. Then, the relevant outside option of the consumer is not necessarily the highest realization found, but rather the highest between the highest realization and the highest realization not yet inspected, minus the search costs. This distinction was not relevant for the infinite case, but it is for the finite one. However, since the certain equivalent of such a product is just the known utility of the product minus its search cost, the result of Proposition 4 is unaffected.

Finally, we notice that if $n \neq m$, a consumer could inspect a full diagonal and not find anything good enough to keep. In this case, however, all realizations for the attribute with less variants would be exhausted by construction. After the last element on the diagonal is inspected, all products remaining will have at most one unknown component. In this case, [Weitzman \(1979\)](#)’s result trivially applies.

E. Extension - More than two attributes

We are interested in extending the intuition of the main model to products defined by more than two attributes. Suppose each product was characterized by a combination of three attributes, A ,

³⁰Once again, I assume the consumer breaks indifference by searching in increasing order of the index when indifferent.

B , and C , each coming in infinite variants. We want to build compound boxes around these products. A compound box built around, for example, $(1, 1, 1)$ will contain all products defined by at least one between A_1 , B_1 , and C_1 . In particular, this box will contain:

- product $(1, 1, 1)$, readily available,
- small nested boxes each containing one between products defined by:
 - $[1, 1, k]$, with reservation value $r_{1,1,k} = A_1 + B_1 + z$,
 - $[1, j, 1]$, with reservation value $r_{1,j,1} = A_1 + C_1 + z$,
 - $[i, 1, 1]$, with reservation value $r_{i,1,1} = B_1 + C_1 + z$,
- “intermediate” nested boxes centered around products defined by:
 - $[1, j, k]$, with reservation value $\mathcal{R}_{1,j,k} = A_1 + \mathcal{R}_{j,k}$,
 - $[i, 1, k]$, with reservation value $\mathcal{R}_{i,1,k} = B_1 + \mathcal{R}_{i,k}$,
 - $[i, j, 1]$, with reservation value $\mathcal{R}_{i,j,1} = C_1 + \mathcal{R}_{i,j}$,

each containing the product readily available and small nested boxes that share one of the two unknown attributes with it.

The logic underlying the value of opening this bigger box is the same underlying the two attributes compound box, though the presence of multiple different objects makes it more complex. The relevant question is: for which observations is the consumer expected to open which of the smaller boxes contained inside the compound box?

From the two attributes case, we know that there are several relevant thresholds for each attribute to account for. We know that an individual attribute beats all closed compound boxes that are affected by it if above threshold $\max\{z + s, y^T\}$. Suppose the consumer opened $X_{1,1,1}$ and found $A_1 > \max\{z + s, y^T\}$. Then, A_1 beats all closed boxes defined by a different variant for the A attribute. This implies that the consumer would not want to open nested boxes $(i, 1, 1)$ nor $(i, 1, k)$, or $(i, j, 1)$. Which one she would open between the remaining ones depend on realizations B_1 and C_1 . Since this applies for the other attributes as well, high realizations can be resolved immediately: if all three attributes are above $\max\{z + s, y^T\}$, the consumer would stop and gain $3E[y|y > \max\{z + s, y^T\}]$. If exactly two attributes are, the consumer would want to keep those two attributes and search along the third one, expecting to gain $2E[y|y > \max\{z + s, y^T\}] + E[y|y > z]$. By the same logic, if only one is the consumer would want to keep that one and open intermediate boxes instead. If none are, the highest would be kept and intermediate boxes would be opened as well.

We know from the two attributes case that if s is high enough it holds: $z + s < y^T$ so that some realizations are not good enough to be kept but are able to reroute search towards the respective attribute. We must distinguish two cases: if $s \leq \tilde{s}$, this cannot happen, and all viable unopened compound boxes have unknown component of their value equivalent to $R_{i,j,k} = [2 - F(z)^2]\bar{y}_2 + F(z)^2\bar{z}_2$, where the subscript “2” indicates that we are referring to the equivalent objects in the two attributes case. If $s > \tilde{s}$, instead, we must distinguish between realizations that are kept and realizations that update the expected value of searching forward.

Suppose $s \leq \tilde{s}$ and define $G(Y) = 2F(y)$ the cumulative distribution of the sum of two attributes, which is well defined thanks to the i.i.d. assumption. Then, expected value of opening compound box $X_{1,1,1}$ is:

$$\begin{aligned} E[U_{1,1,1}] &= [1 - F(z + s)]^3 (3\bar{y}_3) \\ &\quad + [1 - F(z + s)]^2 ([F(z + s) - F(z)](2\bar{y}_3 + E[y|z < y < z + s]) \\ &\quad + F(z)(2\bar{y}_3 + E[y|y > z])) \\ &\quad + [1 - F(z + s)]F(z + s)^2(\bar{y}_3 + \bar{Y}) \\ &\quad + F(z + s)^3 \left(\bar{Y} + \frac{1}{F(z + s)^3} \int_0^{z+s} y dF_{(3,3)}(y) \right), \end{aligned}$$

where:

$$\bar{y}_3 \equiv E[y|y > z + s], \quad \bar{Y} \equiv E[Y|Y > R_{i,j,k}],$$

and the final component is the order statistic representing the expected value of the highest of three realizations below $z + s$. Then:

$$\begin{aligned} R_{1,1,1} &= [1 - F(z + s)]^3 (3\bar{y}_3) \\ &\quad + [1 - F(z + s)]^2 ([F(z + s) - F(z)](2\bar{y}_3 + E[y|z < y < z + s]) \\ &\quad + F(z)(2\bar{y}_3 + E[y|y > z])) \\ &\quad + [1 - F(z + s)]F(z + s)^2(\bar{y}_3 + \bar{Y}) \\ &\quad + F(z + s)^3 \left(\bar{Y} + \bar{z}_3 \right), \end{aligned}$$

where:

$$\bar{z}_3 : s = \int_{\bar{z}_3}^{z+s} \left(\frac{y}{F(z + s)^3} - \bar{z}_3 \right) dF(y)^3.$$

Suppose now that $s > \tilde{s}$. Now, there are realizations $y \in (z + s, y^T)$ that cause an update in the boxes they affect. This update follows the same logic of the two attributes case: an attribute in this interval search towards itself rather than allowing for exploration of other variants in that dimension once a variant in the other dimension is picked. In particular, once an attribute such as this is found, the boxes affected have updated value $\mathcal{R}_{i,j,k} = y + \tilde{R}_2$ rather than $\mathcal{R}_{i,j,k} = y + R_2$, where y is the (at that point) known attribute, and \tilde{R}_2 is the value of the closed boxes with two unknowns after a variant of one of the unknown components is found such that $y \in (z + s, y^T)$. If two such attributes are found, the update uses \tilde{R}_2 instead. Then, $R_{1,1,1}$ must change to reflect these possibilities:

$$\begin{aligned}
R_{1,1,1} = & [1 - F(z + s)]^3 (3\bar{y}_3) \\
& + [1 - F(z + s)]^2 \left([F(z + s) - F(z)](2\bar{y}_3 + E[y|z < y < z + s]) \right. \\
& \left. + F(z)(2\bar{y}_3 + E[y|y > z]) \right) \\
& + [1 - F(z + s)] \left([F(y^T) - F(z + s)]^2 (\bar{y}_3 + E[Y|Y > \tilde{R}_2]) \right. \\
& + [F(y^T) - F(z + s)]F(z)(2\bar{y}_3 + E[Y|Y > \tilde{R}_2]) \\
& \left. + F(z)^2(2\bar{y}_3 + E[Y|Y > R_2]) \right) \\
& + F(z + s)^3 \left(z_3 + [F(y^T) - F(z + s)]^2 (E[Y|Y > \tilde{R}_2]) \right. \\
& \left. + [F(y^T) - F(z + s)]F(z)(E[Y|Y > \tilde{R}_2]) + F(z)^2(E[Y|Y > R_2]) \right).
\end{aligned}$$

This score reflect the value of inspecting a product and then searching optimally based on only the information learned through it, consistently with all the information available. Based on the number of attributes found to be high enough to keep, different paths can be optimal. Comparing search paths across different three attribute boxes, then, would lead to the same updating dynamic highlighted for the two attribute case, and similar threshold values determining how good a realization must be to beat the next three attribute compound box, preserving the logic of the search process.