Three Essays on Consumer Search and Platform Design

Inauguraldissertation zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften
der Universität Mannheim

vorgelegt von
Jacopo Gambato
im Frühjahrs-/Sommersemester 2024

Abteilungssprecher:	Prof. Klaus Adam, Ph.D
Referent:	Prof. Dr. Martin Peitz
Koreferent:	Prof. Dr. Bernhard Ganglmair
Vorsitzender der Disputation:	Prof. Nicolas Schutz, Ph.D
Tag der Disputation:	23.04.2024

Ackowledgements

First of all, I want to thank my supervisor, Martin Peitz. He was an inspiration ever since I started the PhD program, he taught me the importance of both rigor and simplicity, that sometimes less is more, and that keeping an open mind regarding problems and a careful thought process are essential to pursue useful, high quality research. Thank you, Martin, for everything. Thank you also to the other mentors I had the pleasure of confronting myself with throughout these years: Nicolas Schutz, Volker Nocke, Bernhard Ganglmair, Eleftheria Triviza, Anton Sobolev, and Atabek Atayev. A special thank you to Luca Sandrini, a wonderful co-author and friend, and a joy to debate economics, politics, risotto recipes, and videogames with.

I thank the ZEW, who gave me room to grow and nurtured me to thrive. To Georg Licht, who hired me despite not being a perfect match for the Department of Economics of Innovation and Industrial Dynamics, thank you for this chance. To Hanna Hottenrott, who just recently became head of the Department but immediately left her mark, thank you for your openness (and for one of the most entertaining road trips in recent memory); I wish I could see you have an even bigger impact in the years to come. Thank you also to Heidrun Förster, Heidi Halder, and Marion Schäfer, without whom the department could not have functioned, and whose kindness I will not soon forget. A thank you goes also to Ann-Katrin Daubner, my first point of contact for most problems and the occasional glass of wine.

I am grateful for the administrative support from Golareh Khalilpour, Ulrich Kehl, Marion Lehnert, Caroline Mohr, Susanne Putz, Sylvia Rosenkranz, and Kristina Kadel. I gratefully acknowledge financial support from Programm zur Förderung ausländischer Doktorandinnen und Doktoranden in strukturierten Promotionsprogrammen (DAAD, German Academic Exchange Service), and CRC TR 224 (Project B02).

The time spent working would have been much heavier without the company of my wonderful cohort: Giovanni Ballarin, Suzanne Bellue, Tommaso Gasparini, Lukas Mahler, Laura Montenbruck, Oliver Pfäuti, Andrés Plúas Lopez, Jasmina Simon, Jonathan Stäbler, and Aleksandra Khokhlova (an honorary member as far as I am concerned), as well as my colleagues at the ZEW, too many to mention all. A special thank you, however, goes to the members of the Junior Research Group Competition and Innovation, who were exceptional companions and friends: Bernhard Ganglmair (thanked already as a supervisor, now thanked as a friend), Nadine Hahn, Alexander Kann, Ilona Tsanko (the best office mate anyone could hope for), and Moritz Lubczyk and Marita Freimane, who I consider honorary members. Another big thank you to all that lent me their ear in these six years, that found the time to share a coffee, a beer, or a meal, and made my stay unforgettable: Carl-Christian Groh, Marco Reuter, Markus Trunschke, Fabian Seyrich, Yulia Evsyukova, So Jin Lee, Chiara Malavasi, Andrii Tarasenko, Matilde Cappelletti, Julia Krämer, and Max Riegel.

The most important acknowledgment goes to my family, who supported me all these years from my hometown, and always welcomed me with warmth, joy, and plenty of Italian food to replenish myself and to bring back up north. To my mother, who always got more excited than me for all my results and news. To my father, who took a scholarly interest to my work, questioned me with curiosity and sincerity, and proved to be a great sparring partner. Most of all, to my brother, younger than me and yet already a doctor, the best person I know, and the one I want to make proud more than anyone. I could not have done this without you all, and I am forever grateful.

Finally, thank you to the wonderful people at "Delizioso", the best coffee in Mannheim, and the one place in town that felt like being back home.

Contents

List of Figures

List of Tables

Introduction

This dissertation contains three chapters on how information affects market outcomes, the role of information design, and learning with search frictions. Consumer data is accumulated in incredible volumes, and at evergrowing speed. The rise of data brokers testifies to the importance of data collection, manipulation, and analysis for profitability in the modern economy. Data, and information more in general, plays different roles in different contexts. For this reason, this dissertation focuses on the role of information in several settings.

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

The findings carry implications for recommendation systems: the model suggest that a monopolist with information on a consumer's taste has the incentive to bias recommendations away from her best match. Instead, the consumer would be optimally shown a cheap product she does not like in some dimension, which will encourage her to inspect more expensive products given the information she learned early in the search process. Further, the model warns about the ability of adaptive pricing algorithms to discriminate based on the search history of consumers: a consumer inspecting two products that share attributes in a sequence reveal that this attribute satisfies the consumer's needs. Then, it would be optimal for the products inspected after any other inspection to be priced differently based on the sequence of options sampled before.

A common concern in online markets is the ability of platforms connecting buyers and sellers to bias recommendations to maximize revenues using the information these agents have on consumer taste. A particularly interesting setting is that of streaming platform and their algorithmic recommendations. In the second chapter, which is joint work with Luca Sandrini, we study the incentives of a streaming platform to bias consumption when products are vertically differentiated. The platform offers mixed bundles of content to monetize consumers' interest in variety and pays royalties to sellers based on the effective consumption of the content they produce. When products are not vertically differentiated, the platform has no incentive to bias consumption in equilibrium. With vertical differentiation, royalties can differ; the platform always biases recommendations in favor of the cheapest content, which hurts consumers and the high-quality seller. Biased recommendation always diminishes the incentives of a seller to increase the quality of her content for a given demand. If a significant share of the users is ex-ante unaware of the existence of the sellers the platform can bias recommendations more freely, but joining the platform encourages investment in quality. The bias, however, can lead to inefficient allocation of R&D efforts. From a policy perspective, we propose this as a novel rationale for regulating algorithmic recommendations in streaming platforms.

In many online markets, exchange of information is at the center of economic interactions. Recent regulatory intervention has been focused on consumer privacy with the specific aim of empowering consumers in the way information about them is concealed and disclosed. A novel finding in the privacy literature by Ali, Lewis, and Vasserman (Review of Economic Studies, 2023) shows that when consumers have the ability to disclose private information strategically, they can do so in a way that leads to retention of higher gains from trade than without disclosure. In the third chapter, which is joint work with Martin Peitz, we analyze consumers' voluntary information disclosure in a platform setting. For given consumer participation, the platform and sellers tend to prefer limited disclosure of consumer valuations, in contrast to consumers. With endogenous consumer participation, seller and platform incentives may be misaligned, and sellers may be better off when consumers can disclose their valuations. A regulator acting in the best interest of consumers and/or sellers may want to intervene and force the platform to employ a disclosure technology that enables consumers to voluntarily disclose information from a richer message space.

Chapter I

Consumer Search and Firm Strategy with Multi-Attribute Products

1 Introduction

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

The consumer search literature has highlighted the role of search frictions as determinants of market outcomes.[2] The effect of these frictions for intra-firm search, however, has been so far understudied.[3] I contribute to the literature by incorporating correlation across the products offered by a single firm to study how this affects consumers' optimal search and how firms would condition their strategic decisions on it. In particular, I consider consumers that value products based on their attributes (Lancaster, 1966). Initially, consumers observe all products and respective attributes, but they do not know how much they value the different attributes. For example, laptops may differ in their processing speed and graphical capabilities, which depend on the processor and the graphic card that are installed.

Consumers decide which products to search for and inspect, and then, based on their findings, adapt their strategy accordingly for their next search. The reasoning is as follows: if two products share an attribute, consumers value them identically with respect to that attribute. Through the search process, consumers learn their preferences for attributes and, depending on what they learn about specific attributes, can redirect their subsequent search because they know which products share the same attributes and which ones do not. The result of any given inspection makes consumers update her expectations for the remaining products based on which attributes they share. This, in turn, instructs the next inspection.

In many circumstances, these learning dynamics represent well consumer search behavior: if a consumer learns that she dislikes a certain attribute in a product, she would rationally try to avoid other products that share that attribute. For example, Hodgson and Lewis (2020) shows evidence of "spatial learning" in search: consumers inspect more differentiated products early and get closer to the eventually purchased option as search progresses. I show that this

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

[2] Consumers generally do not make consumption choices with perfect information; evidence of this can be found in the empirical industrial organization literature (Sovinsky Goeree, 2008) and in the marketing literature (Mayzlin and Shin, 2011).

[3] Prominent exceptions are Petrikaitė (2018) and Nocke and Rey (ming).

multi-attribute structure generates a version of Weitzman (1979)'s optimal search in an environment with correlated products. Further, I show that, in this environment, "backtracking" to a previously inspected and abandoned attribute can be optimal.

A multiproduct monopoly firm commits to a menu and posts products' prices anticipating the consumer optimal search process. Prices are posted and contribute to determining the order in which consumer search for their preferred option.[4] Because the outcome of each inspection instructs the next, each inspection reveals the consumer's learned preferences. The firm can price products differently to encourage consumers to self-sort based on the preferences they learn about through the search process, a mechanism reminiscent of that highlighted in Mayzlin and Shin (2011). Unlike in Mayzlin and Shin (2011), however, different prices can emerge in my framework even if products are *ex ante* identical from the consumer's perspective.

Differential prices might induce the consumer to deviate from the seller's preferred order of search. I show that in some cases, when the product menu is relatively small, the seller has an incentive to restrict the supply by removing specific products from the menu and, with them, alternative search paths available to the consumer.[5] The menu restriction induces the firm's preferred order of search to arise, and it is an optimal strategy when the likelihood of a positive realization is high and search is cheap. Whenever this is the case, the seller strictly prefers a uniform pricing strategy over setting different prices for different products. Therefore, both uniform and differential prices can arise in equilibrium.

The results highlight the ability of a multiproduct firm to steer consumers through strategic menu selection. By anticipating how a consumer would react after observing a product, the seller can encourage search towards better suited products, and profit off the consumer's incentive to find good matches. The seller wants the consumer to keep searching whenever possible: what is learned through inspection of a product makes the consumer fine-tune her selection.[6] The seller can increase profits by setting higher prices along paths consistent with positive realizations without discouraging the consumer to search on paths consistent with negative ones.

I propose strategic menu selection as a new mechanism through which consumers can be steered in their consumption choices. To the best of my knowledge, this paper is the first to study the implications of menu selection for consumer steering. In particular, the dynamics presented in this paper highlight the possibility to steer consumers through recommendation systems in environments when a single agent has control over the product menu. Additionally, the paper unearths a novel possible instance of price discrimination based on consumers' search history.

The rest of the paper is structured as follows: after reviewing the related literature, I present the framework (Section 2) and characterize the optimal search process with multiple attributes and the learning process they imply in a simplified version of the model (Section 3). Afterwards,

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

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

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

I solve the problem of a monopoly seller that selects which products to make available and their prices (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 and limitations in Section 6, I conclude in Section 7.

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

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

Other authors have incorporated correlation in products in the presence of search frictions: Shen (2015) and Armstrong and Zhou (2011) embed the search process in a Hotelling framework so that, in both settings, the available products are perfectly negatively correlated. 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) provides conditions under which correlation in 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. Conditional search order is also at the core of a recent contribution by Doval (2018): The paper extends Weitzman's search process by allowing the consumer to consider all uninspected products as viable outside options, which changes the relative value of the available options and, therefore, the optimal consumer search process.

Weitzman (1979)'s result relies on the assumption that boxes are independently distributed. I relax this assumption and propose a tractable, history-dependent scoring system that incorporates the value of searching beyond the target of inspection. The score is determined accounting for the paths that would be optimally taken by the consumer after the realization they refer to and, therefore, reflect the full value of inspecting new attributes and the respective continuation value. Through this scoring system, I show that a dynamic, adaptive version of Weitzman (1979)'s optimal search policy can be characterized in this environment. Therefore, the paper relates to the growing literature of learning in search (Garcia and Shelegia, 2018; Greminger, 2022; Preuss, 2023).

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

This paper contributes to the literature on the interaction between menu selection and pricing (Brander and Eaton, 1984; Johnson and Myatt, 2002; Nocke and Schutz, 2018). Novel to the literature is the inclusion of correlation across the products offered by the firm. Correlated products allow consumers to learn their preferences as they inspect options and, therefore, the presence of correlated products affects the value of inspecting each product in isolation.

Finally, the paper contributes to the literature of pricing in search. The seminal Wolinsky (1986) model, and most of the literature that followed, focuses on competitive settings.[7] Instead, I study within-firm directed search in a monopoly setting as in recent contributions by Petrikaitė (2018) and Nocke and Rey (ming). The latter studies the incentives of a multiproduct seller to "garble" product information to induce consumers to search longer. Because search costs are assumed to be fixed, the firm has no incentive to price discriminate. Petrikaitė (2018), instead, shows that a multiproduct seller can steer consumers towards expensive products by obfuscating cheaper options. In my framework, steering can arise in the form of differential prices being optimally set by the seller without strategic obfuscation of the products made available. The paper, then, is related to the growing steering literature as well (e.g. Ichihashi, 2020, who also considers a monopoly setting).[8]

2 Simplified Framework

The products. I consider an industry with products differentiated with respect to two attributes.[9] A product (i, j) is identified by attributes $A_i \in A$ and $B_j \in B$. I consider first a simplified framework in which A and B come in two variants each and follow a simple binomial distribution. In Section 5, I propose a more general framework to build on the intuition of the simplified one. Each attribute A_i can be found combined with all attributes B_j, $i, j \in \{1, 2\}$, and *vice versa*. One can visualize the products as displayed in a grid, with the rows representing the A attributes, the columns representing the B attributes, and the cells representing products defined by a

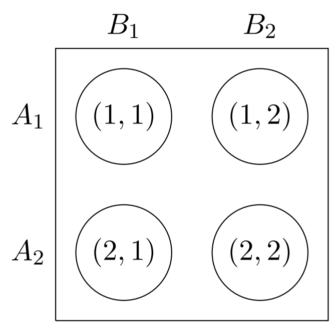

Figure 1: *Products in the same row (resp. column) share attribute A_i (resp. B_j).*

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

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

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

specific combination of A and B attributes as depicted in Figure 1. Notice that products are only differentiated horizontally through their attribute compositions and are otherwise identical in quality.

The consumer. A representative, risk-neutral consumer (she) has unit demand, is aware of the available products and their attribute composition, and can inspect the products in any order she likes. The consumer has no prior knowledge of her preferences over the available attributes; she learns the realization of each attribute separately by inspecting a product characterized by it. In line with existing models,[10] I assume that *ex post* utility generated by a generic product (i, j) takes the form:

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

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

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

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

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

In this environment, I study the optimal sequential search process with free recall: a consumer can always go back to a previously inspected product at no additional cost. The cost of inspecting a product is indexed by the constant $s \in (0, 2\alpha)$. The consumer learns the value of each attribute separately after inspecting a product defined by it. Finally, the consumer's outside option is normalized to $u_0 = 0$.

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

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

1. The seller selects $\tilde{N} \subseteq N$ products to make available and price vector $\mathbf{p}(\tilde{N})$.

2. The consumer observes $\tilde{N}$, $\mathbf{p}(\tilde{N})$, chooses between searching and her outside option, and, if she searches, what to inspect.

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

3. After each inspection, the consumer chooses between stopping and keeping searching (and what to inspect next) until she either purchases an inspected product or leaves without making a purchase.

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

3 A Simple Model of Multi-Attribute Search

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

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

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

Updating expected utilities. Suppose that the consumer already inspected one of the products. Because all products are *ex ante* identical, inspecting $(1,1)$ first is without loss of generality.[11] Whenever the product to inspect can be chosen randomly without loss of generality, I assume that products are inspected in increasing order of their indices. After the first inspection, the consumer has learned realizations A_1 and B_1. Which of the remaining products should be inspected next, if any?

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

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

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

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

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

- if $A_1 = B_1 = 0$, $(2,2)$ is searched next; no other search can take place because $A_2 \geq A_1$ and $B_2 \geq B_1$.

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

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

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

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

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

4 Seller's Optimal Strategy

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

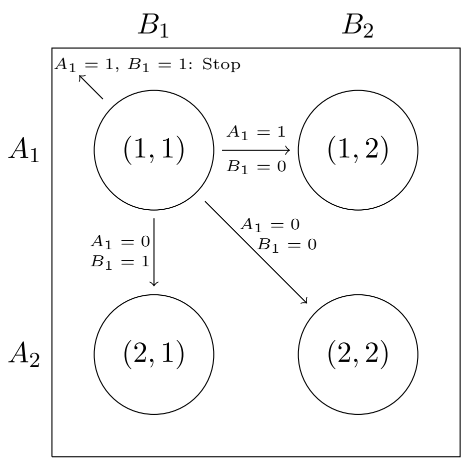

Figure 2: *Optimal search with binomial distribution and all products available, starting from* $(1,1)$.

The two decisions - menu selection and pricing - interact in non-obvious ways. Uniform prices, for example, cannot induce an order of search different from the one characterized above.

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

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

4.1 Uniform Prices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

where D stands for "discourage".

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

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

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

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

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

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

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

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

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

which maximizes probability of trade and is always valid,

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

4.2 Differential Prices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- the seller can let the consumer search $(2,2)$ first, and further increase $p_{1,2}$ and $p_{2,1}$ to $(p_{1,1} + 1 - s)$, or

- the seller can reduce prices $p_{1,2}$ and $p_{2,1}$ to encourage his preferred order of search to arise, or

- the seller can remove $(2,2)$ to induce his preferred order of search and keep the same prices for all other products.

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

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

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

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

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

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

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

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

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

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

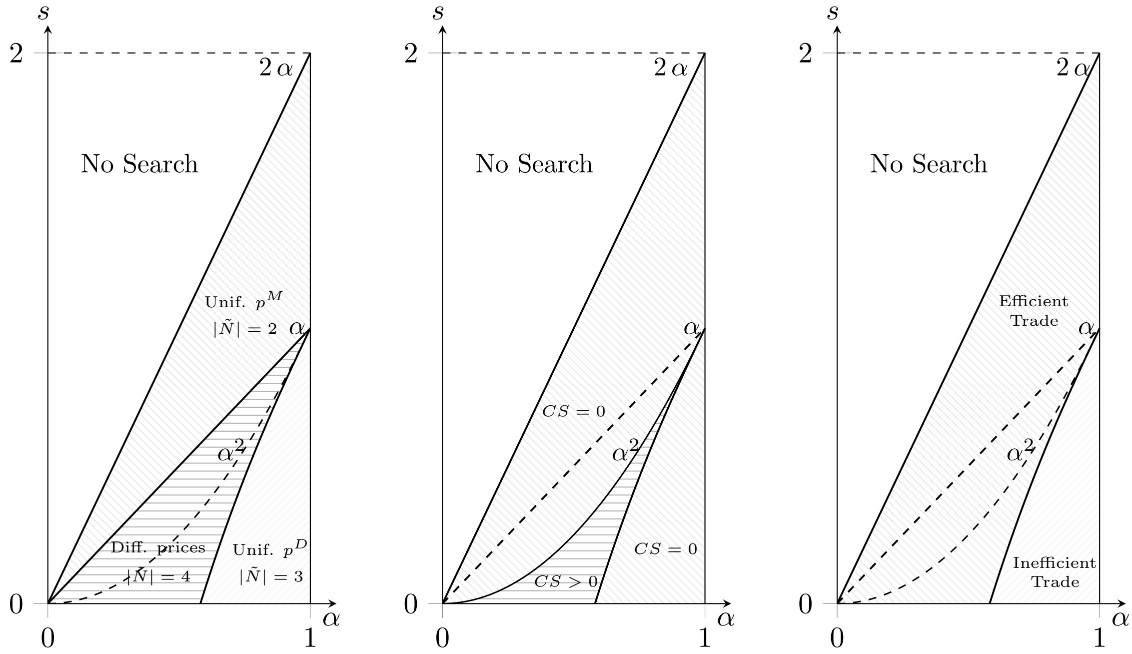

Figure 3: *Equilibrium monopoly menu selection and pricing (leftmost), expected consumer surplus (center), and trade efficiency (rightmost) for all feasible combinations of* $\alpha \in (0, 1)$ *and search costs* $s \in (0, 2\alpha)$

prices are set up to generate strictly higher profits while maintaining the same probability of trade, it is clearly an improvement to set differential prices.

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

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

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

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

4.3 Discussion of the Results

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

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

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

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

Consumers self-selecting based on taste also creates the incentive to condition pricing on their search history. Algorithmic pricing, the practice of pricing items automatically to adapt to the state of the market, are more and more commonly used in the digital world.[12] The model highlights the role that a product's position in the attribute space plays in their pricing.

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

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.

5 Optimal Search with Multi-Attribute Products

The simplified framework analyzed above hints at the mechanics underlying optimal search in this environment. A consumer inspects 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.

I now generalize this intuition. I rethink the problem in a way that allows to apply standard search 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 environment due to correlation: because products share attributes, it is not possible to assign to each one a score that only depends on the product itself. I show how this can be achieved by building "nests" of products to be scored as a single "box" and letting their score update following certain realizations to account for changes in the optimal search that would follow.[13]

The nests relevant for the search process can be considered effectively independent from each other at the moment of search. By leveraging this structure, static scores can be constructed following the same steps highlighted in seminal work by McCall (1970) and Kohn and Shavell (1974). These static scores can then be combined to account for possible changes in the structure of the search process in response to certain realizations, which generates the appropriate reservation values instructing the search process.

Adapting the framework I assume now A and B to come in infinite variants so that the number of products available for purchase is infinite as well. Further, I assume each attribute to be an i.i.d random variable distributed according to a cumulative distribution function F: given a generic attribute $y \in A \cup B$, $F(y)$ is assumed to have support $[0, \hat{y}]$ for some positive $\hat{y}$, and to be twice-differentiable everywhere on it. The cost of inspecting a product is still indexed by the constant s, and the consumer's outside option is still normalized to $u_0 = 0$. The timing of the interaction is unchanged.

[13]This approach was inspired by the work contained in Anderson et al. (2021); I thank Daniel Savelle for his many helpful comments.

5.1 A General Search Process

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).[14] 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:

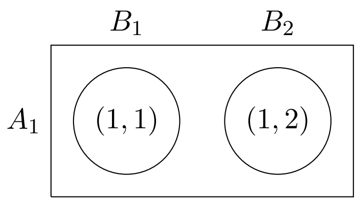

Figure 4: *Two products available*

$$s = \int_{z}^{\hat{y}} (B_2 - z)\, dF(B_2).$$

Then, the reservation value of inspecting $(1, 2)$ when A_1 is known is simply:[15]

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

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 $(1, 1)$: because the reservation value of each individual product depends on the other, we cannot apply Pandora's search algorithm.

Suppose however that the products were in a bigger box, and that the consumer had to decide whether to open one box containing $(1, 1)$ and a nested box containing $(1, 2)$, or nothing at all. The action of opening this "compound" box, that I refer to as $X_{1,1}$, can be scored.

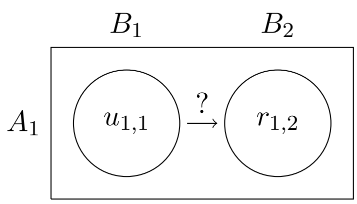

Figure 5: *Myopic search after inspecting* $(1, 1)$

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

The consumer could also want to inspect $(1, 2)$ first. We can imagine another compound box, $X_{1,2}$, containing $(1, 2)$ and a nested box containing $(1, 1)$. The two are *ex ante* identical

[14] This intuition can be found, for example, in Ke and Lin (2022).

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

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.

The value of a compound box. First, define $X_{i,j}$ a box containing products (i,j), immediately available upon opening the box, and all products $(i,j' \neq j)$, $(i' \neq i,j)$, inside smaller boxes that must be opened by paying an additional search cost s. In words, $X_{i,j}$ contains all products defined by A_i, B_j, or both. Notice that all products are contained in multiple boxes. Through this feature, correlation can be handily accounted for.

Consider any compound box $X_{i,j}$ in isolation. To define optimal search inside of it (that is, assuming no other compound box is available), we keep the assumption that, whenever indifferent, the consumer inspects closed boxes in ascending order of their index. Suppose the consumer is about to open $X_{i,i}$ and pay the relative search cost s. Let A^H, B^H be the highest past realization of the previously inspected $A_{j<i}$, $B_{j<i}$. If $\max\{A^H, B^H\}$ is low enough, the choice of searching $X_{i,i}$ is unaffected by all realizations that took place before. On the other hand, if either or both A^H and B^H are high enough, the choice is predictably affected by said realization.

Consider box $X_{i,i}$. The consumer is aware that inside she will find product (i,i) and will have the option to stop or inspect products $(i,j \neq i)$, $(j \neq i,i)$. How would she do so? All attributes $A_{j<i}$, $B_{j<i}$ have already been inspected and are known. Suppose the consumer already opened the box. The consumer can choose between

- stopping at (i,i), generating utility $u_{i,i} = A_i + B_i$, or
- searching again keeping A_i and
  - inspect a product defined by $B_{j<i}$, whose realization is already known, after paying cost s: $u_{i,j<i} = A_i + B_j - s$,
  - search a product defined by $B_{j>i}$, whose realization is unknown after paying cost s: $E[u_{i,j>i}] = A_i + E[B_j] - s$,
- searching again keeping B_i and:
  - search a product defined by $A_{j<i}$, whose realization is already known, after paying cost s: $u_{i,j<i} = A_j + B_i - s$,
  - search a product defined by $A_{j>i}$, whose realization is unknown after paying cost s: $E[u_{j>i,i}] = E[A_j] + B_i - s$.

After opening $X_{i,i}$, these choices can be ranked according to the classic result of Weitzman (1979). In particular, after $X_{i,i}$ has been opened, the remaining options are independent of each other because all attributes are assumed to be i.i.d. Therefore, we can assign a score to all of the options above by finding the certain equivalent of each. Stopping and inspecting a product whose realization is fully known trivially has certain equivalent matching the known ex post utility: $r_{i,i} = u_{i,i}$, $r_{i,j<i} = u_{i,j<i} - s$, $r_{j<i,i} = u_{j<i,i} - s$. Keeping the classic nomenclature, I refer to this as "reservation values" of these options.

Notice that these closed boxes, that I henceforth refer to as "nested" as they are accessible in this form only inside a compound box, are only unknown in one attribute after $X_{i,i}$ has been opened. This is the same object whose reservation value is provided in Choi et al. (2018).[16] In particular, because all attributes $y \sim F(y)$ with support $[0, \hat{y}]$, the certain equivalent of spending a search cost s to discover the realization of any unknown attribute y is z that solves:

$$s = \int_z^{\hat{y}} (y - z)\, dF(y) \tag{3}$$

and, therefore: $r_{i,j>i} = A_i + z$, $r_{j>i,i} = z + B_j$.

Notice that the choice between moving forward towards $(i, j > i)$ or $(j > i, i)$ and going backward to any known $(i, i' < i)$, $(i' < i, i)$ is resolved again simply by applying Weitzman (1979)'s optimal search policy: if there is at least one product $(i, j < i)$ (or $(j < i, i)$) such that $u_{i,j<i} - s > A_i + z$ (or $u_{j<i,i} - s > z + B_i$), no nested box with score $r_{i,j>i} = A_i + z$ (or $r_{j>i,i} = B_i + z$) would be opened, and the product generating the highest $u_{i,j<i} - s$ (or $u_{j<i,i} - s$) would be inspected and selected. This happens if $A^H > z + s$ (or $B^H > z + s$). Otherwise, all products $(i, j < i)$ (or $(j < i, i)$) would be ignored. Because A^H, B^H are the highest past realizations, they are known before $X_{i,i}$ is opened. Therefore, the consumer opens $X_{i,i}$ knowing already whether she would go forwards (that is, open nested boxes $(i, j > i)$ or $(j > i, i)$) or backwards (that is, inspecting a product $(i, j < i)$ or $(j < i, i)$) if she decides to search again).

This observation implies that unopened compound boxes that are constructed around products not sharing attributes are *de facto* independent for all values of A^H, B^H at the moment of making the choice of opening a new compound box. The possible configurations in which compound boxes can be found is illustrated in Figure 6. Depending on A^H and B^H, the effective path inside each unrelated box does not cross. We can then compute the expected value of searching box $X_{i,i}$ in isolation by tracing the optimal search therein for different values of A^H, B^H. This, in turn, means that each configuration as in Figure 6 can be solved independently. The value of opening these compound boxes in each configuration, then, can be combined to obtain the actual, history dependent reservation value governing optimal search.

5.2 Four Independent Configurations

Configuration 1. If $\max\{A^H, B^H\} < z + s$, the consumer will either stop at (i, j) or open nested boxes with unknown content. In this configuration the consumer searching inside a compound box always keeps the highest between A_i, B_j and either stops if the lowest is above z or opens nested boxes paying search cost s. In the latter case, the consumer will keep doing so until she finds something that beats z paying a search cost for each inspection.

To leverage the independence of compound boxes locked in a given configuration, it is necessary to obtain the distribution of values the consumer expect to find inside of it. Let w be

[16] In Appendix III I show that once a nested boxes is optimally opened, the consumer either stops or opens more nested boxes depending on the current realized payoff, making this branch of the optimal search policy myopic in nature.

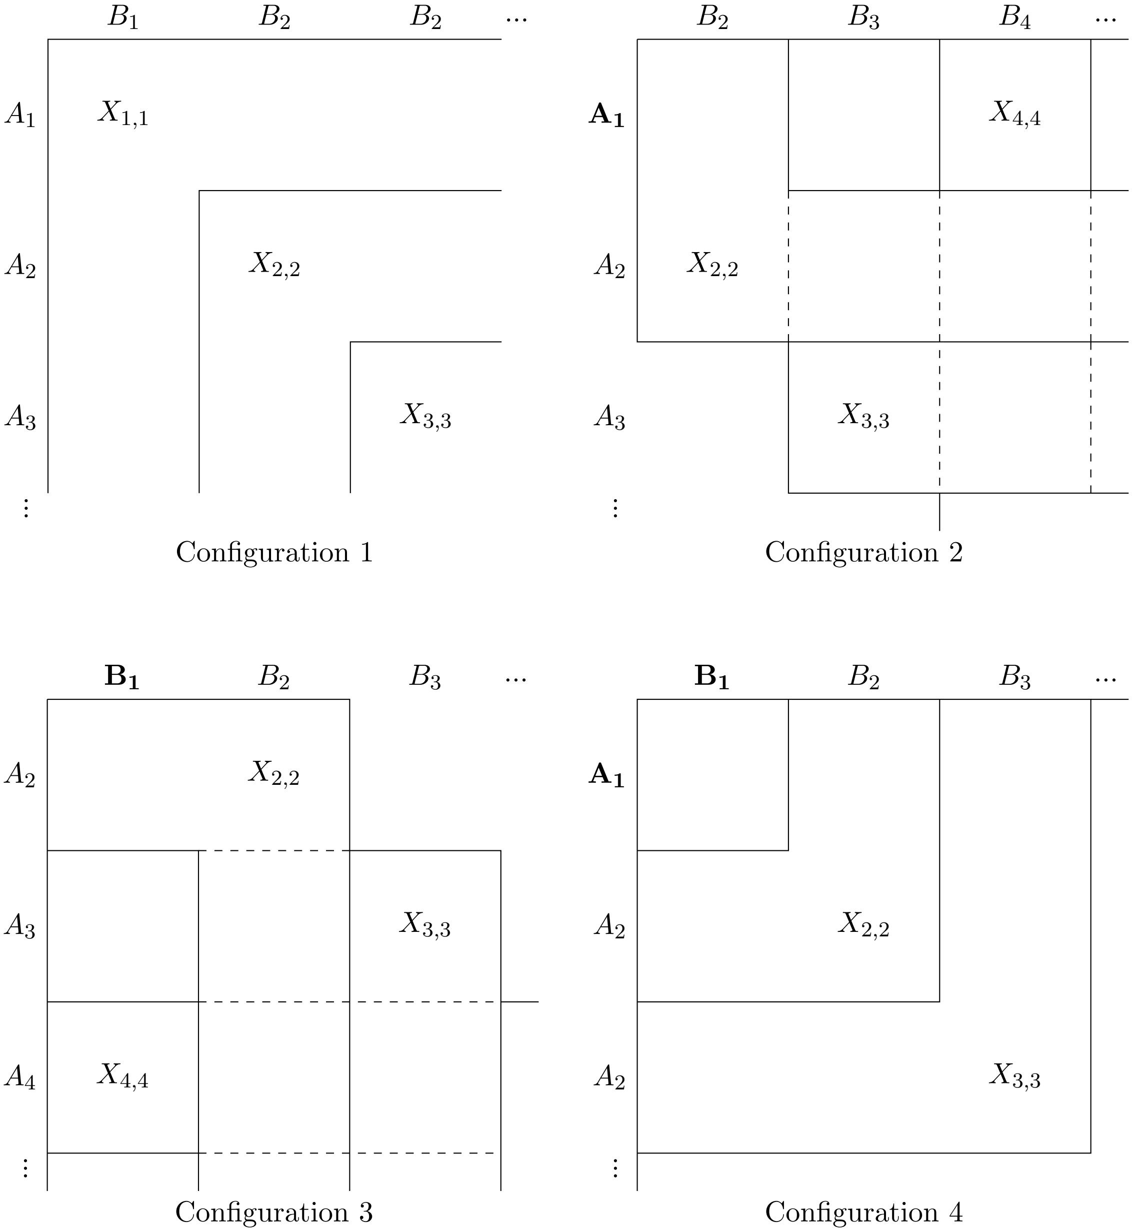

Figure 6: *Possible configurations of compound boxes. Conf.* 1 *represents boxes when* $\max\{A^H, B^H\} < z+s$*; Conf.* 2 *and* 3 *represent boxes when* $A_1 = A^H > z+s > B^H$ *and* $B_1 = B^H > z+s > A^H$ *respectively and, therefore, reroute search towards themselves; Conf.* 4 *represents boxes in which* $A_1 = A^H > z + s$ *and* $B_1 = B^H > z + s$.

the expected payoff of a consumer opening a compound box in this configuration and $H(w)$ be its CDF. Given optimal search inside the compound box, $w_{i,i}$ of opening box $X_{i,i}$ is:

$$w_{i,i} = \max\{A_i, B_i\} + \max\{z, \min\{A_i, B_i\}\}.$$

To compute $H(w)$, we must consider how different realizations for A_i and B_i interact. Fix a generic value B_i. If $B_i < z$, it is kept if and only if $B_i > A_i$. Otherwise, A_i is kept. In this case, $w = \max\{A_i, B_i\} + z$. If $B_i > z$, it is always kept over nested boxes $(i, j > i)$; A_i is also kept if its realization is above z, or: $w = B_i + \max\{A_i, z\}$.

Because $F_a(A) \equiv F_b(B) \equiv F(y)$ have support $[0, \hat{y}]$, $H(w)$ has support $[z, 2\hat{y}]$ and can be expressed as:

$$\begin{aligned} H(w) = & \int_0^z F_a \left(\int_0^B F_b\,(w-z)\; dF_a(A) + \int_B^{\hat{y}} F_b\,(w-z)\; dF_a(A) \right) dF_b(B) + \\ & + \int_z^{\hat{y}} F_a \left(\int_0^z F_b\,(w-z)\; dF_a(A) + \int_z^{\hat{y}} F_b\,(w-A)\; dF_a(A) \right) dF_b(B). \end{aligned}$$

Suppose that the all compound boxes are "locked" in this configuration as in Figure 6 (top left corner).[17] The optimal search in this simplified case can be obtained through definition of a value function as shown in McCall (1970) and Kohn and Shavell (1974). In particular, we want to find $\underline{W}$ that solves the dynamic programming problem:

$$\underline{W} = -s + \max\{w, E[\underline{W}]\}, \tag{4}$$

where w follows the cumulative distribution function $H(w)$, and $\underline{W}$ is the maximum return the consumer would obtain after opening a compound box (and searching optimally therein if she stopped there). In this case, the optimal process sees the consumer stopping and keeping $w \geq E[\underline{W}]$ and searching if $w < E[\underline{W}]$. Because compound boxes locked in a configuration are effectively independent objects, this problem bears the same solution as Weitzman (1979). In particular, the relevant threshold value above which a box is kept is $\underline{W}$ that solves:

$$s = \int_{\underline{W}}^{2\hat{y}} (w - \underline{W})\, dH(w), \tag{5}$$

Configurations 2, 3, and 4. The same procedure allows to obtain static reservation value associated with boxes locked in different configurations. Consider first configuration 2: if $\max\{A^H, B^H\} > z + s > \min\{A^H, B^H\}$, the consumer will not open nested boxes along one attribute but would do so along the other. W.L.O.G., assume $A^H > z + s > B^H$ so that after opening $X_{i,i}$, the consumer would always go back to a product $(j < i, i)$ rather than opening nested boxes $(j > i, i)$ (but could still open nested boxes $(i, j > i)$). In particular:

- if $B_i > z$, the consumer chooses between keeping (i, i), $u_{i,i} = A_i + B_i$, and returning to $(j < i, i)$, $u_{j<i,i} = A^H + B_i - s$

[17] That is, imagine boxes to be unchangeable and such that the value of its content always follows w without possibility of being updated.

- if $B_i < z$, instead, the consumer chooses between $(j < i, i)$ and inspecting nested boxes $(i, j > i)$, $r_{i,j>i} = A_i + z$.

Let $w_a(A^H)$ be the expected payoff of a consumer opening a compound box in this configuration and $H_a(w_a(A^H))$ be its CDF.[18] Assuming again that unopened compound boxes are locked in this configuration (illustrated in the top right corner of Figure 6), their static reservation value is $\underline{W}_a(A^H)$ that solves:

$$s = \int_{\underline{W}_a(A^H)}^{\hat{y}} (w^a - \underline{W}_a(A^H))\, dH_a(w^a). \tag{6}$$

The same exact exercise leads to $\underline{W}_b(B^H)$ (bottom left corner of Figure 6), relevant when $A^H < z + s < B^H$.

Consider now configuration 4 (bottom right corner of Figure 6): if $\min\{A^H, B^H\} > z + s$, the consumer will not open any nested box. In particular:

- if $B_i > B^H - s$ and $A_i > A^H - s$, the consumer stops,
- if $B_i > B^H - s$ and $A_i < A^H - s$, the consumer inspects and keeps (i', i), $u_{i',i} = A^H + B_i - s$,
- if $B_i < B^H - s$ and $A_i > A^H - s$, the consumer inspects and keeps (i, i'), $u_{i,i'} = A_i + B^H - s$,
- if $B_i < B^H - s$ and $A_i < A^H - s$, instead, the consumer chooses between (i', i) and (i, i'), depending on which has the highest utility.

Labeling $w_{a,b}(A^H, B^H)$ and $H_{a,b}(w_{a,b}(A^H, B^H))$ the expected payoff and CDF of boxes locked in this configuration, their reservation value is $\underline{W}_{a,b}(A^H, B^H)$ that solves:

$$s = \int_{\underline{W}_{a,b}(A^H,B^H)}^{\hat{y}} (w^a - \underline{W}_{a,b}(A^H, B^H))\, dH_{a,b}(w^{a,b}). \tag{7}$$

The final step requires to combine these thresholds to account for the fact that compound boxes are not locked in any given configuration but, rather, can move from one configuration to the next depending on the realizations A^H and B^H found along the search process.

5.3 Optimal Search Process

The values $\underline{W}$, $\underline{W}_a(A^H)$, $\underline{W}_b(B^H)$, and $\underline{W}_{a,b}(A^H, B^H)$ can be appropriately combined to obtain the reservation values of unopened compound boxes when they are not locked in any given configuration. Once again, which of this values is relevant depends on past realizations: if some $A_{j<i} > z+s$ and/or some $B_{j<i} > z+s$ is found, this affects the value of all future boxes because by construction all compound boxes contain at least one product defined by all attributes.

The relevant value of the unopened compound boxes can evolve only in one direction, from configuration 1 to 4, and never backwards. Indeed, once $A_{j<i} > z + s$ is found, it can never be forgotten: once the relevant reservation value of the current configuration of $X_{i,i}$ changes from $\underline{W}$ to $\underline{W}_a(A^H)$, it can never revert to $\underline{W}$ or change to $\underline{W}_b(B^H)$. From this point onward, it can

[18] A closed form expression for this and all subsequent CDFs can be found in Appendix III.

only stay at $\underline{W}_a(A^H)$ or change to $\underline{W}_{a,b}(A^H, B^H)$. Moreover, once configuration 4 is reached, all unopened compound boxes will keep this configuration.

Suppose all closed boxes reached configuration 4. This implies that $\min\{A^H, B^H\} > z + s$. Suppose the consumer has observed these A^H and B^H and must choose whether to open the next box. If boxes were to be locked, with any future A and B realization not being able to affect the next, the value of all closed boxes would be $\underline{W}_{a,b}(A^H, B^H)$. However, this does not capture the search dynamics appropriately.

Suppose the next box were to be opened and that $A_i > A^H$ was be found. The next compound box would have a different reservation value, $\underline{W}_{a,b}(A_i, B^H)$. The expected value of future boxes given the current values A^H, B^H can be obtained recursively. Let $\underline{W}^*_{a,b}(A^H, B^H)$ be the expected equivalent of costly opening the next box on the search path. This can be rewritten as a linear combination of expected $\underline{W}_{a,b}$ values:

$$\begin{aligned}\underline{W}^*_{a,b}(A^H, B^H) =& \underline{W}_{a,b}(A^H, B^H) \int_0^{B^H} \int_0^{A^H} dF_a(A) dF_b(B) + \\ &+ \int_0^{B^H} \int_{A^H}^{\hat{y}} \underline{W}_{a,b}(A, B^H)\, dF_a(A) dF_b(B) + \\ &+ \int_{B^H}^{\hat{y}} \int_0^{A^H} \underline{W}_{a,b}(A^H, B)\, dF_a(A) dF_b(B) + \\ &+ \int_{B^H}^{\hat{y}} \int_{A^H}^{\hat{y}} \underline{W}_{a,b}(A, B)\, dF_a(A) dF_b(B).\end{aligned}$$

Consider the choice of the consumer. If she chooses to open the next compound box, $X_{i+1,i+1}$, she knows that she will stop only if $w_{i+1,i+1}$ is higher than the new value $\underline{W}_{a,b}(A^H, B^H)$, which in expectation is equal to $\underline{W}^*_{a,b}(A^H, B^H)$ before $X_{i+1,i+1}$ is opened.

Notice that $\underline{W}^*_{a,b}(A^H, B^H)$ is strictly higher than $\underline{W}_{a,b}(A^H, B^H)$ because $\underline{W}_{a,b}(A^H, B^H)$ is increasing in A^H and B^H. This threshold captures not only the value of inspecting the next box, that by itself would have had reservation value $\underline{W}_{a,b}(A^H, B^H)$, but also that of the updating that opening the box might lead to. In words, the value of opening the next box is the expected "certain equivalent" of opening the next, which in itself depends on the outcome of the inspection.

We can repeat the same exact exercise with the other configurations. For configuration 2, we write:

$$\begin{aligned}\underline{W}^*_a(A^H) =& \underline{W}_a(A^H) \int_0^{z+s} \int_0^{A^H} dF_a(A) dF_b(B) + \\ &+ \int_0^{z+s} \int_{A^H}^{\hat{y}} \underline{W}_a(A)\, dF_a(A) dF_b(B) + \\ &+ \int_{z+s}^{\hat{y}} \int_0^{A^H} \underline{W}_{a,b}(A^H, B)\, dF_a(A) dF_b(B) + \\ &+ \int_{z+s}^{\hat{y}} \int_{A^H}^{\hat{y}} \underline{W}_{a,b}(A, B)\, dF_a(A) dF_b(B).\end{aligned}$$

which is both higher than $\underline{W}_a(A^H)$ and the hypothetical $\underline{W}_a(\tilde{A}^H)$ one would compute ignoring the possibility that the next box could change in value. An equivalent formulation can be found for configuration 3.

Finally, for configuration 1, we can write:

$$
\begin{aligned}
\underline{W}^* = &\underline{W} \int_0^{z+s} \int_0^{z+s} dF_a(A) dF_b(B) + \\
&+ \int_0^{z+s} \int_{z+s}^{\hat{y}} \underline{W}_a(A)\, dF_a(A) dF_b(B) + \\
&+ \int_{z+s}^{\hat{y}} \int_0^{z+s} \underline{W}_b(B)\, dF_a(A) dF_b(B) + \\
&+ \int_{z+s}^{\hat{y}} \int_{z+s}^{\hat{y}} \underline{W}_{a,b}(A,B)\, dF_a(A) dF_b(B).
\end{aligned}
$$

By taking into account all possible configurations, and all the ways in which this configurations can evolve into one another, we can then write the reservation value of all unopened boxes as:

$$
\mathcal{W}(A^H, B^H) = \begin{cases} \underline{W}^* & \text{if} \quad \max\{A^H, B^H\} < z+s, \\ \underline{W}^*_a(A^H) & \text{if} \quad A^H > z+s > B^H, \\ \underline{W}^*_B(B^H) & \text{if} \quad B^H > z+s > A^H, \\ \underline{W}^*_{a,b}(A^H, B^H) & \text{if} \quad \min\{A^H, B^H\} > z+s. \end{cases}
$$

The values above reflect the value of inspecting any given compound box given all information learned so far and anticipating how the game could change given future realizations: $\mathcal{W}(A^H, B^H)$ incorporates the value of searching along all possible paths, defined by the number of attributes found above $z+s$. Once a path is taken, that path can never be left. Each path is built through the optimal search process when boxes are in the appropriate state through $\underline{W}_\cdot$, which is defined based on the optimal search process inside the compound box as per its current configuration, captured by $w_\cdot(i,j)$ for each $X_{i,j}$. Because each branch is optimized, the whole process is too.

While the optimal search order cannot be determined *ex ante* because of the learning component, whenever the consumer must choose what to do there is no ambiguity regarding the value of her possible options:

Proposition 3. *Let $A^H = \max\{y \in A \cap I\}$, $B^H = \max\{y \in B \cap I\}$ be the highest discovered realization for A and B. Optimal search is characterized as follow:*

- ***Compound box selection:*** *compound boxes are opened until the expected payoff according to the optimal search policy inside of it, $w_{i,j}$, is higher than the reservation value of all unopened compound boxes, $\mathcal{W}(A^H, B^H)$.*

- ***Search inside the selected compound box:*** *Given selection of compound box $X_{i,j}$, (i,j) is kept if $u_{i,j} > \max\{r_{i,k\neq j}, r_{k\neq i,j}\}$; otherwise, the next box opened is the nested box with the highest $r_{i,k}$ or $r_{k,i}$.*

- ***Stopping rule:*** *Boxes (compound or nested) are opened until all unopened (compound or nested) boxes have updated reservation value below the highest realized payoff.*

Proof. All calculations and closed form equations for $H(w.)$ can be found in Appendix III. ■

The multi-attribute structure proposed here allows one to score search options appropriately by leveraging the fact that at any given point compound boxes can be thought of as effectively independent object along the search path. This, in turn, generates an environment in which the standard intuition behind optimal search can be adapted. This process can be thought of as a consumer sampling unrelated products until at least an attribute worth keeping is found. When this happens, the consumer ignores all remaining compound boxes and searches inside the one that let her find that first attribute to keep in order to find an appropriate other one to pair with it.

Notice that the structure of the compound boxes reflect the internal consistency of the search process: opening a compound box always carries more information than a nested box inside a previously opened one. Therefore, if a compound box is selected, the attribute that is kept when searching inside of it must have had a realization high enough to compensate for the lower informational value of not inspecting two new attributes.

5.4 Optimal Pricing with Infinite Products

In order to solve for the optimal pricing scheme in this more complex environment, we leverage once again the structure of compound boxes. This 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 can be used to score compound boxes appropriately, accounting for the price of all products on the relevant search paths. In other words, prices affect the values $w.$ of opening any compound box; the effect cascades to the reservation values $\mathcal{W}$, which allows to solve for optimal pricing.

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 search in it keeping attribute A_1. Then, she would inspect next the product $(1,j)$ that satisfies:

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

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:

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

The same structure governs inspection of products $(i,1)$.

In principle, all products $(1,j)$ could be priced differently. Suppose that prices where 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:

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

To see the effect of prices, it is useful to compute the expected value of a compound box when the products therein have prices posted. Consider the generic compound box $X_{i,j}$ in the first configuration for simplicity and for illustrative purposes. Let $\Delta_{i,k} \equiv p_{i,k+1} - p_{i,k}$ and $\Delta_{k,j} \equiv p_{k+1,j} - p_{k,j}$; further, let:

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

Then:

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

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 a 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 $E[w(\mathbf{p})]$, which translates to a lower reservation value. None of the boxes associated with this product, then, would ever be inspected as there are infinite better alternative for the consumer.

Another difficulty relates to the updating process 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, because the relationship between the different possible scores $\mathcal{W}$. 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 4. *Consider a multiproduct seller pricing infinite products defined by two infinite sets of i.i.d. attributes. There exist a unique, uniform equilibrium pricing vector such that* $p_{i,j} = p^* = \underline{W}^*$, $\forall (i,j)$.

Proof. All calculations can be found in Appendix III. ∎

Proposition 4 states that the only possible equilibrium features uniform pricing. In the simplified framework of Section 4, different prices could be optimal because the second product searched was generally the last one: if $(2, 2)$ was inspected after $(1, 1)$, the consumer would either purchase it or stop searching because no other products were available to inspect. To make it worth searching, it had to be priced accordingly. In the infinite attributes case, instead, 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.[19]

To further highlight the role of menu size, notice that $p^* = \underline{W}^*$ is at once the "encouraging" and "discouraging" price in the words of the simplified framework. It encourages search because it does not prevent search after a bad realization because compound boxes do not shrink in this framework. At the same, it matches the value of opening the compound box exactly, just like the discouraging price did in the simplified framework. This suggests that the relationship between optimal pricing and menu selection depends not only on the menu composition, but also on its size.

The fact that compound boxes do not shrink does not necessarily imply that products cannot be priced differently. In principle, given the reservation value of a compound box, 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 III, I show that this cannot be optimal. The intuition is as follows: suppose that 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 $\delta > 0$.[20] Plugging in these prices in the score of the compound box, one finds:

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

where $\underline{y}_\delta$ is the expected value of the highest of two realizations below $z - \delta$.

Studying $E[w_{1,1}(\mathbf{p_{1,1}})]$ 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 because $F(z-\delta)$ 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 expected 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.

That $p^* = \underline{W}^*$, the initial reservation value of any compound boxes, follows from the updating dynamic detailed in the previous subsection. In particular, it follows from the fact that all updates increase the value of subsequent boxes rather than shrink it. In particular, the lowest value of a compound box after any updating can be shown to be $\underline{W}^*_{a,b}(A^H, B^H) \geq \underline{W}^*_{a,b}(z+s, z+s) = \underline{W}^*$. Therefore, the highest prices that the monopolist can set is the highest price that does not prevent search from taking place and, in particular, $p^* = \underline{W}^*$.

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

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

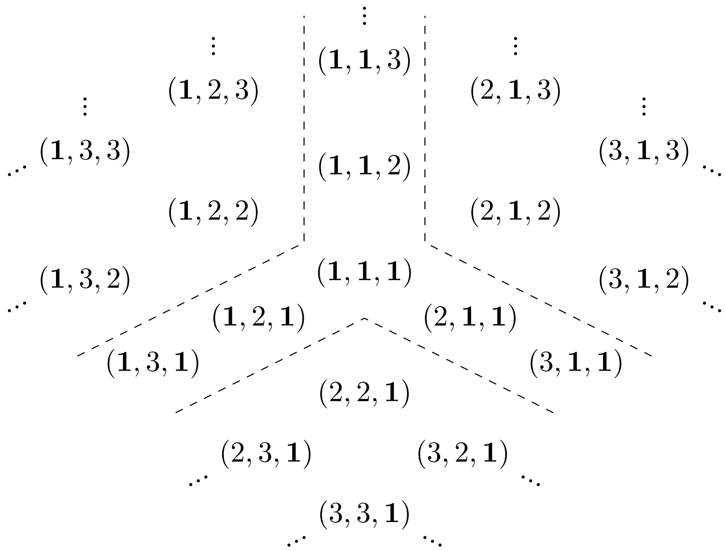

Figure 7: *Graphical representation of a three attribute compound box centered around* $(1,1,1)$*. Products that share two attributes with* $(1,1,1)$ *can be displayed along the edges of a cube (north for products sharing* A_1*,* B_1*, south-west for products sharing* A_1*,* C_1*, south east for products 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*).*

6 Extensions

6.1 More than Two Attributes

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. With three attributes, two kinds of closed boxes must be included in the compound box with the product it represents. On one hand, all products sharing exactly two attributes with the central product can be represented as small nested boxes equivalent to the ones contained in the two attributes case.

On the other, products that share only one attribute with the central one are unknown in two dimensions, and must be placed in two-dimensional boxes equivalent to the compound boxes of the baseline model. These "intermediate" boxes themselves contain infinite small nested boxes as well. 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 as in Figure 7.

We can conceptualize the same process to find the optimal search path for a consumer searching in this environment. First, it is necessary to rethink the structure of a generic nested box $X_{i,i,i}$. If $X_{i,i,i}$ is built around product (i,i,i), it contains all products that share at least one attribute with it. Therefore, it can be represented as the three edges of a cube and the sides delimited by them. Each side can be thought of as a two-dimensional grid in which "intermediate" compound boxes akin to the ones defined in the main model can be found. The choice of opening these boxes is governed by the same $\mathcal{W}.$ functions defined above. The choice of opening a different three-dimensional box, instead, requires computing the reservation value of the possible "locked" configurations this box can come out of. All configurations will always be made of three edges and three sides. Whether the edges stretch forward, towards undiscovered attributes, or backward, to known past realizations, depends once again on whether single attributes are found above or below $z + s$, or combinations of two attributes above or below $\mathcal{W}$.

6.2 Purchase Without Inspection

It is assumed throughout the paper that consumers must expend a search cost to inspect any product. Because products in this environment share attributes some uninspected products could be fully revealed without being inspected. If search is understood as the physical action of finding a product, this distinction is immaterial. 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.

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. This can be accomplished by modifying the way reservation values update after each observation. 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 because all products sharing an attribute with an inspected product would be contained, at zero additional cost, in all unopened compound boxes, and affects them all through the same updating detailed above.

The pricing game in Section 4 would be affected by this change. Recall that the price the multiproduct seller 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 seller 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. The change in interpretation does not affect the result qualitatively, but the mechanical change to the search process suggests that prices would be more dispersed in equilibrium under this alternative interpretation of search costs.

6.3 Limitations and Directions Forward

Finite number of products In the baseline general model I consider an infinite number of variants for each attribute: 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. Restricting the environment to finite sets of variants beyond the simple example reported in Section 3 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 8. 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 $\max\{A_1, B_1\} > z > \min\{A_1, B_1\}$ or $z > \max\{A_1, B_1\}$, the highest would be kept.

Differently from before, it is possible now for the consumer to search keeping one attribute fixed and, after exhausting all products sharing that attribute, switching to products characterized by the other. Consider again Figure 8, and suppose A_1 and B_1 both had very low realizations such that $z > A_1 > B_1$. The consumer will inspect $(1, 2)$ next. In the infinite attributes case, the consumer would never run out of $(1, j)$ products to inspect.

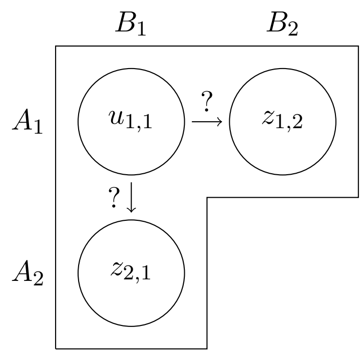

Figure 8: *Three products available*

The consumer would optimally inspect $(2, 1)$ next if it holds:

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

If $B_1 > B_2$, this is trivially true because $z > A_1 > B_1$. 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.

Finding one or more attributes above $z + s$ affects 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 for the same updating detailed in the baseline model to take place.

The final difference with the infinitely large boxes of the baseline model also makes the problem likely intractable and follows from the fact that when compound boxes are opened and discarded, the 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 index, 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 implication of this last remark is that while thinking about boxes as locked in some configuration achieves the same conceptual independence between objects, now every subsequent choice is "discounted" by the value associated with one more variant for each attribute. Effectively, this means that the choice of searching now and searching again later can never be the same. While in principle this could be accounted for, as the structure resembles of that of Weitzman (1979), combining the resulting locked reservation values to generate adaptive ones to take the place of $\mathcal{W}$ quickly leads to a computationally intractable problem.

Different distributions. In principle, removing the assumption of attributes following the same distribution can be accommodated. One can 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. 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. A general solution of this more complex problem, however, becomes quickly intractable, and is therefore left for future research.

Competition. It is natural to ask what would be the features of an hypothetical equilibrium with competing firms. The results above rely strongly on the seller'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 seller 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. This incentive to undercut, however, makes pinning down an equilibrium with competing firms exceedingly complex and well beyond the scope of this paper. Nonetheless, if such an equilibrium exists it should feature lower, uniform prices when consumers have the same prior considered here.

7 Conclusion

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

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

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

Chapter II

Not as Good as it Used to be: Do Streaming Platforms Penalize Quality?

Joint with Luca Sandrini

1 Introduction

Streaming platforms found great success in the digital era thanks in part to personalized features such as the well-known Spotify "Discover Weekly" playlist, automatically generated for each user every week.[21] Such features are extremely popular and have a real impact on consumption patterns. For example, Aguiar et al. (2021) show in their empirical investigation that inclusion in automatically generated playlists such as Spotify's "New Music Friday" boosts future popularity compared to similar songs not included. The two observations strongly point to the ability of such platforms to affect individual consumers' effective consumption bundle.

In this study, we focus our attention on the ability such a platform has to bias consumption through algorithmic recommendation: we study the role of a streaming platform in enabling consumers to "mix" different products in individually optimal proportions and its ability to profitably introduce bias. We consider a framework in which two horizontally differentiated sellers compete not only with each other but also with a monopoly platform offering a streaming service to all the consumers in the market. The platform attracts consumers who want to enjoy the content produced by both sellers, cashes in a subscription fee, and pays royalties to the sellers based on the effective consumption of their product. The set-up represents economic agents such as Spotify and different music labels who sell a product that is available both on the platform and outside of it, for example in the form of CDs or audio files up to purchase.

While most of the current literature on the topic addresses concerns regarding anti-competitive steering practices,[22] we adapt the framework to investigate the effect of subscription-based business models on the incentive of content providers to innovate. As shown in recent work by Bourreau and Gaudin (2022), streaming platforms have strong incentives to bias recommendations to reduce consumption of products that carry higher royalties.[23] The result is intuitive: since the platform charges a fixed fee to all consumers, she has the incentive to minimize costs once their participation is ensured. The finding speaks in favor of royalties being strategically used to gain prominence on this kind of platform (Bourreau et al., 2021). It is however unclear how this dynamic is affected by vertical differentiation in the products offered. On the one hand, a higher quality product is more desirable and, therefore, allows the platform to charge more to access it. On the other hand, as the seller with the superior product would demand to be

[21] Popper (2016) reports that 40 million out of Spotify's, at the time, 100 million users used it in 2016. More recently, Spotify reported that in the five years since its launch, Discover Weekly streamed 2.3 billion hours of music. See https://newsroom.spotify.com/2020-07-09/spotify-users-have-spent-over-2-3-billion-hours-streaming-discover-weekly-playlists-since-2015/.

[22] See, for example, Peitz and Sobolev (2022).

[23] In recent work, Aguiar et al. (2023) provide empirical support to the theoretical results. Using data from playlists on Spotify, the authors show that the share of songs owned by major music labels in prominent playlists offered by Spotify decreases over time. The interpretation of the finding is that Spotify is leveraging its power in the playlist market to obtain better deals in royalty payments.

paid a higher royalty rate, the incentive to bias away from his product would be higher as well. Which effect dominates and the overall impact on the platform, the sellers, and the buyers are at the core of our analysis.

More precisely, we propose a framework in which two horizontally differentiated content providers (henceforth labeled "a" and "b") sell their bundle good to a unit mass of consumers uniformly distributed on the $[0, 1]$ line for a price p_j, $j \in \{a, b\}$. Each seller offers a bundled good that consists of only the content they produce (i.e., bundle good a (respectively b) is entirely made of content a (b)). The sellers are assumed to be located at the extremes of the Hotelling line. Along the entire line, a platform for streamed content (labeled p) offers a subscription-based service: upon paying a uniform fee p_p, a consumer can access a mix of content from a and b.[24]

The platform remunerates the sellers and pays them a royalty rate per share of the content shown to each consumer. Consumers, therefore, can choose between three bundle goods: sellers a and b offer pure bundles, whereas the platform offers mixed bundle goods. If a consumer buys either of the pure bundle goods, she consumes only the content owned by one seller. Instead, by subscribing to the platform service, the consumers are offered a mix of content based on their preferences (Anderson and Neven, 1989; Hoernig and Valletti, 2007, 2011) and the platform's recommendation system (Bourreau and Gaudin, 2022). The platform can be understood as an intermediary that smooths consumption for those who value a balanced mix of content.

We show that when products are vertically homogeneous, the existence of such an intermediary represents a Pareto improvement compared to the alternative competitive outcome. The platform attracts consumers located in the middle of the Hotelling line: these are the consumers with the highest willingness to pay for the possibility of mixing products. Hence, the platform can charge a price higher than the sellers' and still make positive profits after paying royalties. This result holds under the extreme assumption that sellers have full bargaining power when setting the royalty rate.

In the baseline specification with uniform products, the platform has no incentives to introduce a bias for its users. When products are not vertically differentiated, the two sellers optimally select the same price in equilibrium anticipating the consumption taking place both in and out of the platform. Biasing consumption, in this case, would skew the demand in favor of one of the two sellers, inducing him to raise his price and monetize from it. The rival would, instead, choose to reduce the price of his product to induce consumers closer to him to leave the platform. Such a strategy cannot be optimal for the platform, as it would effectively bias consumption in favor of the most expensive option rather than the cheaper one.

The model's predictions change drastically when products are vertically differentiated. Since consumers value high quality, without platform intervention the equilibrium outcome features a higher price and larger share of consumption for the high-quality product. While the platform can raise its price to monetize the higher average quality of her bundles, her ability to do so is limited since the rival is forced to offer a lower price than under no vertical differentiation. When consumers are offered their optimal consumption bundle, the platform is hurt by the

[24]Such fees are the most common when it comes to music streaming platforms. Besides Spotify, other notable examples are Deezer and Pandora.

quality differential.

In this specification, we show that the platform always has the incentive to bias consumption away from the better, more expensive product. Since biasing consumption affects the sellers' equilibrium prices, the platform trades off consumption bias and the ability to monetize efficiently on the consumer side. When the quality difference is substantial, the platform is limited by consumers' participation constraints and offers a biased bundle that makes consumers indifferent between joining and leaving the platform. The result follows from the assumption that the platform can discriminate consumers and offer targeted recommendation bias.[25]

When the platform biases consumption away from the better and more expensive product, the respective seller is penalized. If this penalty is severe enough, the seller could choose not to make his product available on the platform and compete with the other seller directly instead. Whenever this happens, it is clear that the platform cannot be active. Consumers join the platform to mix content produced by both music labels: if the platform cannot attract both sellers, no consumer is interested in joining. Streaming platforms, however, have been known to popularize less-known artists and, therefore, generate demand. To capture this additional dimension, we split the unit mass of consumers in two. Some consumers are assumed to be ex-ante aware of the artists represented on the platform, while others are not. The latter group only learns about the artists and consumes their product if the platform manages to attract both music labels. When deciding to join, the sellers' outside option is worse if the group of consumers that only consumes if the platform is active is larger. It follows that the ability of the platform to bias consumption depends on the additional consumption she generates.

The findings have relevant implications both in the context of consumption steering in digital markets and in regard to the effect of subscription-based business models on the incentives of sellers to innovate. First, steering emerges in equilibrium, not because of sellers competing for prominence but rather as a response of the platform to soften competition: the platform has the incentive to contain the price effect generated by the difference in quality and the stronger market presence of the better product. The overall effect hurts consumers and the seller with the high-quality product and benefits the runner-up by skewing consumption towards him.

Incentives to innovate and produce higher-quality goods are weakened when a platform that can bias consumption is present in the market for a given demand. In particular, the platform always selects a positive level of bias when products are vertically differentiated. The bias is constrained by sellers' and consumers' participation decisions. In equilibrium, consumers who join the platform are exposed to more of the cheaper, low-quality content that they would optimally select. Furthermore, we show that the bias introduced by the platform distorts incentives to innovate and can lead to inefficient allocation of R&D efforts, with the more (resp., less) efficient seller investing less (resp., more) than it would have if the platform had been inactive.

The rest of the paper is structured as follows: after a review of the relevant literature, we introduce the model and solve the baseline specification with homogeneous quality products in

[25] The assumption is strong but realistic. It is well known that platforms such as Spotify offer personalized content in the form of playlists based on past consumption. The assumption, then, is simply a reversal of what is already known: the platform being aware of a consumer's taste instructs how much bias she would be willing to tolerate.

Sections 2 and 3, respectively. In Section 4 we introduce vertical differentiation by allowing one of the products to provide additional fixed, stand-alone utility to all consumers. After solving and discussing the seller's participation decision as a function of the demand generated by the platform, we endogenize the decision to invest in quality by the sellers (Subsection 4.2). We complement the main analysis by considering several extensions in Section 5. Section 6 concludes.

Related Literature Recommendation systems represent a core feature of digital platforms, and streaming platforms like the ones we study in this paper are no exception. The impact of these systems on consumer choice has been the focus of many empirical investigations. Among these, the aforementioned Aguiar et al. (2021) and companion paper Aguiar and Waldfogel (2021) speak of the impact that inclusion in automatically generated playlists has on the popularity of new songs on Spotify. Generally, recommendation systems have been shown to greatly widen the range of consumed products, a phenomenon generally referred to as the "long-tail effect" (Fleder and Hosanagar, 2009; Brynjolfsson et al., 2011; Oestreicher-Singer and Sundararajan, 2012; Datta et al., 2018). More recently, literature concerned with the incentives of intermediaries to strategically skew recommendations in a way that systematically harms consumers (see, for example, recent work by Peitz and Sobolev, 2022) has been on the rise.

It seems clear that the impact that these systems have on consumption makes them an obvious candidate for strategic manipulation. In this spirit, Bourreau et al. (2021) studies competition for prominence on digital platforms, comparing bias generated when prominence is gained via monetary or data-based compensation. We distance ourselves from this setting in various ways: first, we capture bias not through manipulation of the search query but through manipulation of the composition of available bundles. Second, we assume the platform already has relevant information on the buyers' side by building competition on the Hotelling line. From this perspective, the paper more closely resembles Bourreau and Gaudin (2022), who consider a market where the platform does not directly compete against the sellers. Instead, we explicitly model competition through the ability for consumers to purchase directly from the sellers active in the market. Furthermore, we allow the platform to condition the bias imposed on consumers based on their location while Bourreau and Gaudin (2022) focuses on uniform biases.

Novel to the literature is also the fact that we introduce vertical differentiation between the sellers: in our model, we assume sellers offer goods that are differentiated both horizontally and vertically. Moreover, we assume consumers are only sensible to horizontal variations of the good, whereas they do not differ in their willingness to pay for quality. In that sense, we depart from Mussa and Rosen (1978), where consumers' income is taken into consideration (see also Cremer and Thisse (1991) and Sutton (1986)). Here, we adopt a model of spatial competition à la Hotelling, in which we assume that one seller offers a good with higher intrinsic value than his rival.[26] We do so to better relate to the existing literature on innovation, which often includes both vertical and horizontal dimensions of differentiation (see Chen and Schwartz, 2013).

Our findings that the platform has the means and the incentive to bias consumption in

[26] In other words, we follow the textbook definition of vertical differentiation in Pepall et al. (2014). Two goods are vertically differentiated if, when offered at the same price, all consumers strictly prefer consuming more of one than of the other.

favor of cheaper content echo those detailed in Freimane (2022). The author examines the impact of a regulatory change affecting the bargaining process behind content provision on the Google News platform. In the paper, it is shown that the change, aimed at granting higher bargaining power to publishers vis-à-vis Google News led the platform to change the composition of articles shown to readers, substituting content provided by larger publishers towards cheaper alternatives. Even though the channel through which the asymmetry arises is different (we hold bargaining power fixed, and focus on vertical differentiation instead), the outcome is well-aligned with our equilibrium predictions.

For the timing of our model, we follow Fletcher et al. (2023): the platform commits to a recommendation system before prices are set, and all agents are aware of the implied potential bias in equilibrium. The choice is motivated both by technical sensibility and by the inner working of the music streaming industry. On one hand, it is sensible to assume recommendation systems to be implemented not in response to a specific interaction with a specific seller, but through an algorithm the platform commits to. Further, there is anecdotal evidence that streaming platforms and sellers therein trade off royalties for exposure. In 2014, the online radio company Pandora admitted to agreeing with indie-label coalition Merlin to exchange lower royalty rates for an increase in exposure.[27] The testimony highlights that these agreements are common practice, and justify a timing such that sellers set royalties aware of the bias the platform might generate in response. In passing, the testimony suggests that the platform might want to commit to its recommendation system regardless for fear of legal repercussions: if they were to condition the recommendation system on royalties, they would realistically be more vulnerable to legal action being taken against them.

Widening the scope of the discussion, the paper relates to the evolving literature on the economics of media markets. While most past contributions focused on the mix of content and advertising in media (Anderson and Coate, 2005; Anderson and Gabszewicz, 2006; Peitz and Valletti, 2008; Thomes, 2013; Halmenschlager and Mantovani, 2017), we ignore this dimension altogether. We do so for two reasons: first, while it is true that many streaming platforms offer free subscriptions with ads in alternative to the ad-free "Premium" ones, the latter in itself represents an enormous and still growing market.[28]

Second, while the literature on advertisement in media contraposes content and ads, bringing positive and negative utility to consumers respectively, we focus on the content bias because of the inherent alignment of interests it breaks. Consumers value good content and are willing to pay more for it. While the trade-off between content and ads is intuitive, the fact that the platform would have the incentive to penalize high-quality products she is not competing with is not. To our knowledge, this is the first paper to explore this dimension of the problem. The emerging result is in conflict with recent work by De Cornière and Taylor (2019): despite the alignment of interest of sellers and buyers to produce and consume better quality products, the "congruence" case studied by the authors, bias in our case can never be consumer surplus improving as they suggest. The reason follows from the discussion in the introduction. The

[27] See the testimony of Stephen McBride, Docket No. 14-CRB-0001-WR for a full recounting

[28] According to Spotify's earning report to investors, the number of premium subscribers in Q3 of 2022 was 195 million. Available at: https://s29.q4cdn.com/175625835/files/doc_financials/2022/q3/Q3-2022-Shareholder-Deck-FINAL-LOCKED.pdf

platform uses bias to strategically reduce cost rather than inflate revenue, which reverts the incentives and the direction of the bias that congruence would suggest.

More in general, the paper relates to the growing literature on platforms initiated by Armstrong (2006a). Streaming platforms find their footing and generate network effects by facilitating mixing in addition to facilitating contact between different sides – be these sides buyers and sellers, or users and advertisers. This distinction separates our work from other models studying platform steering. In particular, Teh and Wright (2022) show that steering can benefit consumers when searching for a product that represents a good enough match is very costly. In our context, instead, the platform profits by offering a service, that is, by allowing consumers to reach mixed bundles and represents a net welfare gain when she cannot, or chooses not to, bias consumption. Whenever she intervenes, however, she does so to the detriment of consumers.

To model our environment, we build on early work by Adams and Yellen (1976) and, more closely, by Anderson and Neven (1989): we consider a Hotelling framework in which location on the unit line uniquely determine the optimal mix of consumption. Intuitively, consumers closer to a (respectively b) want to purchase a higher share of the product produced by a (b). Consumers equidistant from the two find it optimal to consume the two in equal proportions. The framework has been used in the past to study advertisements when consumers mix their consumption (Gal-Or and Dukes, 2003) and, more recently, to study welfare implications of different pricing structures (Hoernig and Valletti, 2007, 2011; Döpper and Rasch, 2024).

The paper also relates to the literature on vertical relations and, in particular, to the coexistence of retailers and direct sale channels available to manufacturers, as well as the strategic interaction of a platform when competing against its suppliers. The recent work by Aguiar et al. (2023) relates very closely to our paper. They analyze the incentives of the platform to include certain types of artists and songs in its playlist to leverage its market power and obtain better licensing deals with major music labels. We differentiate from this work in two ways. First, we propose a theoretical investigation of the platform's incentives to bias its recommendation system to minimize costs. Second, we focus on individual recommendations.

Most of the literature considers consumers more or less sensible to prices and distribution channels.[29] In contrast, Tsay and Agrawal (2004) studies the manufacturer's optimal choice of distribution channels between direct, retail-based, and hybrid, with a focus on market penetration. While retailers extract part of the rent generated by the sale, the additional market penetration they lead to can be worth pursuing. We consider a similar dynamic: music labels, the manufacturers in our setting, can choose to sell through the platform on top of directly because the platform generates additional demand that would remain inactive otherwise.

More closely related to this paper is recent work by Ronayne and Taylor (2022). The paper studies the role of a competitive channel, like an online e-commerce platform, as an alternative distribution channel available to sellers. The authors focus their attention on the optimal governance structure of the competitive channel assuming both this channel and the sellers have some captive consumers to extract rent from. In contrast, our market shares emerge endogenously in equilibrium. The presence of captive consumers in our setting would allow the platform to bias more aggressively in equilibrium, a result that is proxied by the demand

[29] See, for example, Rhee and Park (2000), Chiang et al. (2003), and Kumar and Ruan (2006)

expansion we feature in our model: the more consumers would stay inactive if not for the platform intermediation, the less constrained the platform is in designing the recommendation system.

Finally, this paper encompasses the literature on the effect of rent-sharing mechanisms on innovation incentives (see Berton et al., 2021, for a review of the literature). Most of the research on this topic focused on the effects of institutions such as unions on the incentives of firms to invest in R&D. Grout (1984) first analyzed this topic and concluded that unions might act as rent-seekers, thus lowering firms' incentives to innovate. By appropriating part of the innovation-generated revenues, the argument goes, unions exert negative pressure on firms' incentives and introduce the well-known hold-up problem. Furthermore, Haucap and Wey (2004) focus on unionization structure (i.e., the degree of wage centralization) and its effects on innovation incentives. The authors show that centralized wage-setting institutions are the most efficient in generating innovation incentives. Indeed, under some conditions, a centralized union could also outperform a market where wages are determined competitively. On the contrary, Mukherjee and Pennings (2011) find that union centralization increases the incentive for technology licensing, which, under some conditions, may boost the investments in innovation by firms. In the same spirit, Kline et al. (2019) find that firms obtaining patent protection observe a rise in workers' compensation and productivity.

Our paper contributes to the literature by analyzing the issue in a B2B setting. More specifically, we consider the strategic interaction between innovators and a platform that can steer consumers' demand toward the most convenient good. We show that the platform may severely hinder incentives to invest in innovation even if the innovators have full bargaining power in determining their royalty rates. Moreover, we argue that the platform can appropriate part of the innovation value by biasing its recommendation system and artificially raising competitive pressure on the innovator. Finally, we show that the intervention of the platform can lead to severe inefficiency of the equilibrium allocation of R&D effort.

2 Model setup

There are two firms (sellers or music labels), indexed by $j = a, b$, who are located at the left and right extremes of the $[0, 1]$ Hotelling line. In addition, there is a streaming platform (p) that knows the consumers' location and offers them a personalized bundle of content from the two music labels. By doing so, the platform can better match consumers' preferences (Anderson and Neven, 1989; Bourreau and Gaudin, 2022).

We consider two groups of consumers, informed and uninformed, each uniformly distributed on the Hotelling line in the market for streamed products. The group of informed consumers has mass $\alpha \in [0, 1]$, whereas the group of uninformed has mass $1 - \alpha$. The information they possess (or do not possess) refers to the existence and location of the firms operating in the market. Moreover, informed consumers know ex-ante their own location on the line, as well as the exact location of the sellers. On the contrary, uninformed consumers only know about the streaming service and discover the two music labels if they are available on the platform.

The two sellers produce one good each, a and b respectively. We refer to them as the *pure* *bundles*, which are entirely made of contents produced *in-house*. These can be thought of as the albums produced by the two music labels. Instead, we define *mixed bundles* as the personalized good that the consumers can access via the streaming service, like a playlist that contains content produced by both sellers.

We indicate the location of consumers on the unit line with x. Then, we use $\lambda(x) + \varepsilon(x) \in [0, 1]$ to identify the share of content a consumed by the consumer located at x if she joins the platform service. In particular, $\lambda(x)$ is the preferred share that the consumers would choose to maximize utility, whereas $\varepsilon(x)$ is the personalized bias on the recommendation system imposed by the platform. Put differently, $\varepsilon(x)$ is the extra share of content a offered to each consumer by the platform's algorithm. Conversely, $1-\lambda(x)-\varepsilon(x)$ represents the share of content b offered to the same consumer.

Consumers purchase exactly one unit of the final good — either a pure bundle or the recommended mixed bundles. We use p_a and p_b to define the price of the pure bundles paid directly to the music labels. We use p_p to identify the subscription fee paid by consumers to access the platform's service instead.

Finally, the platform pays royalties (r_j) to the music labels per share of their content offered to consumers. We assume that the music labels charge a royalty rate equal to the market price: $r_j = p_j$. The assumption allows us to ignore any direct bargaining between sellers and the platform and any effect of eventual differences in bargaining power. In a way, we assume that sellers have full bargaining power in the royalty setting stage and, therefore, always select the highest rate possible given their own price in the external market.

The utility function of consumer i located in x_i can be written as:

$$
\begin{aligned}
U_{i,a} =& V_a - p_a - tx_i^2 \\
U_{i,b} =& V_b - p_b - t(1 - x_i)^2 \\
U_{i,p} =& (\lambda(x_i) + \varepsilon(x_i))V_a + (1 - \lambda(x_i) - \varepsilon(x_i))V_b - p_p - t(x_i - (1 - \lambda(x_i) - \varepsilon(x_i)))^2
\end{aligned}
$$

where $V_j = v+v_j$ is the intrinsic quality of the pure bundles (which is common to all consumers) and is composed of a common parameter $v > 0$, that we assume to be high enough to guarantee full coverage in the market — i.e., $v > 3t/2$ — and a music label-specific parameter $v_j \geq 0$. In what follows, we analyze the benchmark case of $v_a = v_b = 0$ and the asymmetric scenario where $v_b > v_a = 0$. Finally, the parameter $t > 0$ represents the transportation costs that multiply the utility loss from taste mismatch. For tractability, we assume $v_b < 2t/3$ always holds. Figure 9 shows the diagram of the model.

Importantly, we allow the platform to bias the bundles offered to consumers as a means to influence music labels' price decisions. By doing so, the platform alters the shares of content in the personalized mixed bundles: we analyze the incentives of the platform to steer consumers away from high-quality, and expensive, content and offer them a mixed bundle that is disproportionately rich in low-quality, and cheap, content.

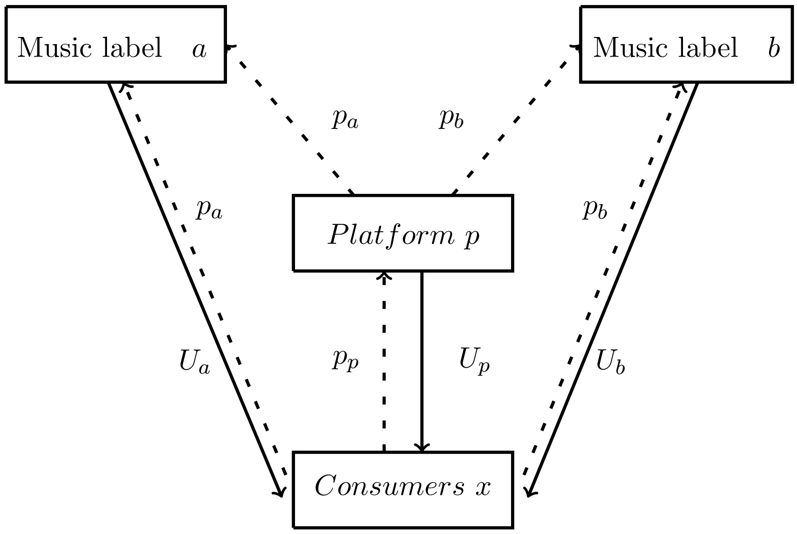

Figure 9: *The diagram of the model with payments and services when the platform is active, and both informed and uninformed consumers participate.*

Timing. The timing of the game is as follows: at stage 1, the platform chooses the level of bias of the recommendation system (ε) and commits to implementing it.[30] At stage 2, sellers observe the recommendation policy and decide whether to join the platform and serve both informed and uninformed consumers or to stay out and compete for informed consumers only in a standard Hotelling setting. Upon observing the entry decision and the quality attributes of the two contents, at stage 3, the two sellers and the platform set the prices for the pure bundles and the streaming service (p_a, p_b, and p_p). In Subsection 4.2, we augment the model with an additional stage 0 in which either one or both sellers costly invest in innovation to generate $v_j \geq 0$.

At stage 4, given the prices and the recommendation system, consumers make their consumption decisions and profits realize. One should remember that the share α of informed consumers know both their location on the Hotelling line as well as the locations of the sellers. Instead, the $1-\alpha$ uninformed consumers only know that a platform exists. We assume that all consumers can sample the platform for free before subscribing.[31] During the free sample period, uninformed consumers learn the location of the firms and their preferences. If the sellers decide not to join the platform at stage 1, it is clear then that uninformed consumers do not learn anything and make no purchase.

Our solution concept is Sub-game Perfect Nash-Equilibrium. We solve the game by backward induction.

[30] Once again, the assumption that the platform can commit to a certain level of bias is motivated by the evidence that platforms promise higher exposure to music labels in order to achieve better economic conditions, and by fear of possible legal repercussions if the platform conditioned bias on royalty rates.

[31] Many real-world streaming platforms, including Spotify, offer free trials to consumers. The assumption, therefore, well matches the kind of platform we aim to model.

3 Baseline: Homogeneous quality

We begin the analysis by focusing on the baseline case where the two music labels produce content of identical quality – i.e., $V_j = v \; \forall j$. We start from the demand faced by the three sellers (the two music labels and the platform) in the last stage of the game. The assumption $v_a = v_b = 0$ simplifies the utility functions to:

$$\begin{aligned} U_{i,a}^{bln} &= v - p_a - tx_i^2 \\ U_{i,b}^{bln} &= v - p_b - t(1 - x_i)^2 \\ U_{i,p}^{bln} &= v - p_p - t(x_i - (1 - \lambda(x_i) - \varepsilon(x_i)))^2 \end{aligned}$$

where the apex ^{bln} indicates we are in the baseline model specification.

As standard in these models, we derive the locations of indifferent consumers by equating the utility functions they obtain by choosing between the three options:

$$\begin{aligned} x_{ap}^{bln} &= \frac{p_p - p_a + t(1 - \lambda(x_{ap}))^2}{2t(1 - \lambda(x_{ap}))} &\implies U_{i,a}^{bln} = U_{i,p}^{bln} \\ x_{pb}^{bln} &= \frac{p_b - p_p + t(2 - \lambda(x_{pb}))\lambda(x_{pb})}{2t\,\lambda(x_{pb})} &\implies U_{i,p}^{bln} = U_{i,b}^{bln} \\ x_{ab}^{bln} &= \frac{p_b - p_a + t}{2t} &\implies U_{i,a}^{bln} = U_{i,b}^{bln} \end{aligned}$$

We adopt the following notation: x_{jk} indicates the indifferent consumer between buying from firm j and firm k, with $j, k = a, b, s$ and $k \neq j$. Notice that the location of the consumer who is indifferent between the two pure bundles a and b must lie between the other two. In the proceeding of the analysis, we use x_{ab} mainly as a reference point.[32]

Notice that $\varepsilon(x)$ does not enter the location of the indifferent consumers. This is simply because the platform knows consumers' locations and can offer them a personalized recommendation system. Indifferent consumers would change their consumption choice if subject to a bias that lowers their utility. Hence, the platform designs its algorithm to increase the bias in the distance between the consumers and their preferred music labels. This bias is personalized and it is bounded by the participation constraint of the consumers. More in detail, the personalized bias for each platform user is $\varepsilon(x_i) < \bar{\varepsilon}(x_i)$, where

$$\bar{\varepsilon}(x_i) = \{\varepsilon(x_i) \in [0, 1 - \lambda(x_i)) \text{ s.t. } U_{i,p}^{bln}|_{\lambda = \lambda(x_i) + \bar{\varepsilon}(x_i)} = \max\{U_{i,a}^{bln}, U_{i,b}^{bln}\} \; \forall x_i \in (x_{ap}^{bln}, x_{pb}^{bln})\}$$

is the maximum level of bias a consumer i located in x_i is willing to accept before leaving the platform and purchasing the pure bundle from her preferred music label. Hence, it is possible to see that the consumers in x_{ap}^{bln} and x_{pb}^{bln} would not accept any bias, as for them $U_a^{bln} = U_p^{bln}(\lambda(x_{ap}^{bln}))$ and $U_b^{bln} = U_p^{bln}(\lambda(x_{pb}^{bln}))$, respectively. Formally, $\varepsilon(x_{ap}^{bln}) = \varepsilon(x_{pb}^{bln}) = 0$.

With this in mind, we can now derive the personalized recommendation system set by the platform. Intuitively, the platform aims at maximizing consumption of the streaming service. To do so, it offers the *efficient bundle* to indifferent consumers. We define *efficient bundle* as

[32] The assumption $v > 3t/2$ is sufficient to ensure $U(x_{ab}^{bln}) > 0$.

the composite good that would be chosen by a consumer so that, for any prices p_a, p_b, and p_p, she would get the highest possible utility. By definition, the efficient bundle is not biased by the platform recommendation system ($\varepsilon = 0$). Formally:

$$\lambda^*(x_i) = \arg\max_{\lambda \in (0,1)} (U_{i,p}^{bln}) = 1 - x_i$$

Using this consumption choice, it is possible to update the location of the indifferent consumers as:

$$x_{ap}^{bln}|_{\lambda(x_{ap}^{bln})=\lambda^*(x_{ap}^{bln})} = \sqrt{\frac{p_p - p_a}{t}}; \qquad x_{pb}^{bln}|_{\lambda(x_{pb}^{bln})=\lambda^*(x_{pb}^{bln})} = 1 - \sqrt{\frac{p_p - p_b}{t}}. \tag{8}$$

This information allows us to compute the demand and expected profits of all agents that we use to derive in the Appendix the following Lemma:

Lemma 1. *The profits of the music labels and the platform for any given level of $\varepsilon(x_i)$ are*

$$\pi_a^{bln} = \frac{1}{18}t(3+2\epsilon)^2; \qquad \pi_b^{bln} = \frac{1}{18}t(3-2\epsilon)^2; \quad \pi_p^{bln} = \frac{t(1-39\varepsilon^2)}{27};$$

and the indifferent consumers are located in:

$$x_{ap}^{bln} = \frac{1}{3} - \varepsilon; \quad x_{pb}^{bln} = \frac{2}{3} - \varepsilon; \quad x_{ab}^{bln} = \frac{1}{2} - \varepsilon;$$

Proof. See the appendix. ■

Finally, we proceed backward to the first stage of the game, when the platform announces and commits to a level of bias $\varepsilon(x_i)$. It follows from the Lemma 1 above that:

Proposition 1. *The equilibrium recommendation system with homogeneous quality is the one that recommends the efficient bundle to all consumers.*

Proof. See the Appendix. ■

To understand Proposition 1, one should look at the prices set by the music labels given the level of bias — i.e., $p_a^{bln} = t + \frac{2\varepsilon t}{3}$ and $p_b^{bln} = t - \frac{2\varepsilon t}{3}$. Interestingly, the bias exerts a positive effect on the price of the seller that it favors. Hence, if the two music labels offer bundles of the same quality, the demand differential implied by a biased recommendation system would result in higher prices for the content that the platform offers more to consumers. It goes without saying that the platform would never engage in such behavior, as it would be detrimental to her profitability. It follows that the only optimal recommendation schedule possible is the unbiased one. Put differently: when goods are homogeneous the platform always offers the *efficient bundle* to all consumers.

Corollary 1. *In equilibrium, all firms make positive profits. In particular, music labels obtain the standard Hotelling outcomes. Formally:*

$$\pi_a^{*,bln} = \pi_b^{*,bln} = \frac{t}{2}; \quad \pi_p^{*,bln} = \frac{t}{27}.$$

Proof. The proof follows from plugging $\varepsilon = 0$ in the results derived in Lemma 1. ■

4 Heterogeneous quality

In this section, we relax the assumption that goods are homogeneous. Indeed, the music industry is both vertically and horizontally differentiated.[33] Music labels experiment and research new ways of expressing their art. In other words, they innovate. It is, therefore, credible to assume that music labels compete with products that embed different quality levels. Therefore, we repeat the analysis assuming that seller b produces content of better quality that ensures higher utility to all consumers irrespective of their preferences for the available varieties. We first assume such quality differential to be a primitive of the model; afterward, we endogenize the choice of costly investment in quality to study which distortions, if any, the intervention of the platform leads to.

4.1 Exogenous quality differential

First, we assume $V_b = v + v_b > v = V_a$ with v_b exogeneously given. The utility functions become:

$$\begin{aligned} U_{i,a}^{hq} &= v - p_a - t x_i^2 \\ U_{i,b}^{hq} &= v + v_b - p_b - t(1 - x_i)^2 \\ U_{i,p}^{hq} &= v + (1 - \lambda(x_i) - \varepsilon(x_i))v_b - p_p - t(x_i - (1 - \lambda(x_i) - \varepsilon(x_i)))^2 \end{aligned}$$

where the apex ^{hq} stands for "heterogeneous quality". Recall that $\lambda(x_i)$ indicates the share of content a in consumer i's individual mix, whereas $\varepsilon(x_i)$ is the personalized bias introduced by the platform. As before, we derive the locations of indifferent consumers by equating the utility functions they obtain by choosing between the three options. These locations do not depend on the bias, as the personalized bias optimally selected for the indifferent consumers is simply $\varepsilon(x_{ap}^{hq}) = \varepsilon(x_{pb}^{hq}) = 0$. We write:

$$\begin{aligned} x_{ap}^{hq} &= \frac{p_p - p_a + (t(1 - \lambda(x_{ap}^{hq}) - v_b)(1 - \lambda(x_{ap}^{hq}))}{2t(1 - \lambda(x_{ap}^{hq}))} && \Longrightarrow \quad U_{i,a}^{hq} = U_{i,p}^{hq} \\ x_{pb}^{hq} &= \frac{p_b - p_p + (t(2 - \lambda(x_{pb}^{hq}) - v_b))\lambda(x_{pb}^{hq})}{2t\,\lambda(x_{pb}^{hq})} && \Longrightarrow \quad U_{i,p}^{hq} = U_{i,b}^{hq} \\ x_{ab}^{hq} &= \frac{p_b - p_a + t - v_b}{2t} && \Longrightarrow \quad U_{i,a}^{hq} = U_{i,b}^{hq} \end{aligned}$$

absent the price effect the quality gap v_b moves the indifferent consumers towards the location of the music label a, thus shrinking her demand: quality shifts demand.[34]

We define the *efficient bundle* in this scenario as:

$$\lambda^{hq}(x_i) = \arg\max_{\lambda \in (0,1)} (U_{i,s}) = 1 - x_i - \frac{v_b}{2t}$$

[33]Music labels backed by major labels generally have more resources than comparable independent music labels to produce better products – e.g., in terms of sound quality or international collaborations.

[34]As we are interested in gradual innovations, we focus on values of v_b for which there is always at least a consumer that prefers the pure bundle b but would derive positive utility from mixing. Formally, $v_b < 2t/3$.

We use this information to update the location of the indifferent consumers:

$$x_{ap}^{hq} = \sqrt{\frac{p_p - p_a}{t}} - \frac{v_b}{2t}; \qquad x_{pb}^{hq} = 1 - \sqrt{\frac{p_p - p_b}{t}} - \frac{v_b}{2t}. \tag{9}$$

The location of the indifferent consumers helps us compute the demand faced by each agent of the model. In particular, the platform's demand is simply given by $D_p = x_{pb}^{hq} - x_{ap}^{hq}$. Instead, the demands of the two sellers change because of the different proportions of content in the new biased bundles. Formally:

$$D_a^{hq} = x_{ap}^{hq} + \int_{x_{ap}^{hq}}^{x_{pb}^{hq}} (\lambda^{hq}(x_i) + \varepsilon(x_i))\, dx \tag{10}$$

$$D_b^{hq} = 1 - x_{pb}^{hq} + \int_{x_{ap}^{hq}}^{x_{pb}^{hq}} (1 - \lambda^{hq}(x_i) - \varepsilon(x_i))\, dx \tag{11}$$

From the above demand functions, it is possible to anticipate that music label b faces a demand that is decreasing in the intensity of the bias ε; consequently, music label a faces an increased demand because of the favorable bias. This is analogous to what we observed in the baseline model, and it explains why in that case bias would not emerge in equilibrium. However, in this model specification, the variation in goods' quality makes the two demands asymmetric, to begin with, and the platform may have incentives to increase the demand for the music label that offers the cheapest good.

Because the variation in music labels' demands is known before prices are set, they affect the equilibrium prices of both the platform and the sellers. Music label b is expected to lower its price in response to the decreased demand, whereas music label a would likely do the opposite as a consequence of the increased demand. Possibly, the two prices would converge towards a common value if the bias is sufficiently intense. Because the consumers demand more content from the high-quality seller, reducing its price is indeed in the interest of the platform, absent any constraint on the sellers' participation.

Because consumers have different tastes, the level of bias is personalized. Hence, the platform cares that the participation constraint of each consumer is satisfied. In deciding the total mass of demand to shift from one seller to the other — i.e., the total bias — the platform pays attention that it does not exceed the sum of the participation constraint of all the consumers (see Figure 10). Formally:

Condition 1. *Given $v_b > 0$, the aggregate personalized bias imposed by the platform to users of the streaming service can be identified by a general mass of bias ε^p, such that*

$$\varepsilon^p \equiv \int_{x_{ap}^{hq}}^{x_{pb}^{hq}} \varepsilon^p(x_i) dx \leq \int_{x_{ap}^{hq}}^{x_{pb}^{hq}} \bar{\varepsilon}(x_i) dx \equiv \varepsilon^c$$

where the apexes ^{p},^{c} indicate the total bias selected by the platform and the maximum bias that satisfies consumers' participation constraint, respectively; $\varepsilon^p(x_i)$ represents the individual level of bias the platform designs.

In order to solve the problem, and since what matters to the platform and sellers is the

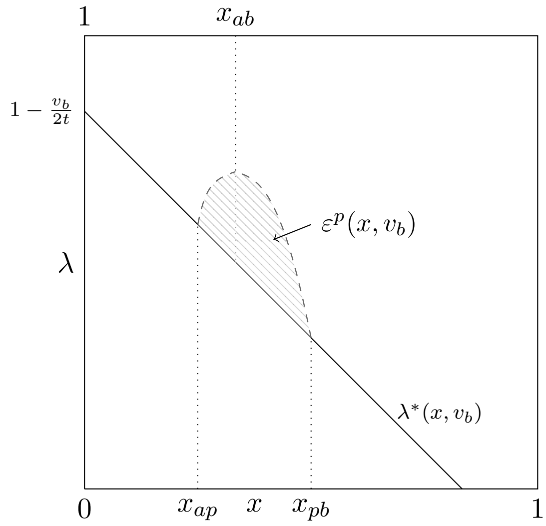

Figure 10: *Personalized biased bundle. The area in yellow is the total demand that the platform can shift from music label b to music label a without losing users.*

total mass of consumption shifted from one seller to the other, we forego solving the optimal individual level of bias for each consumer that joins the platform. Instead, we consider the total mass ε^p, noting that it must always be compatible with the total participation constraints of all buyers combined. This ensures consistency of the solution while maintaining the problem tractable.

The new recommendation system can be written as $\int_{x_{ap}^{hq}}^{x_{pb}^{hq}} \lambda^{hq}(x_i)\, dx + \varepsilon^p$. We adjust the profit functions accordingly:

$$\pi_p^{hq} = p_p\,(x_{pb}^{hq} - x_{ap}^{hq}) - p_a \left(\int_{x_{ap}^{hq}}^{x_{pb}^{hq}} \lambda^{hq}(x_i)\, dx + \varepsilon^p \right) - p_b \left(\int_{x_{ap}^{hq}}^{x_{pb}^{hq}} \Big(1 - \lambda^{hq}(x_i)\Big)\, dx - \varepsilon^p \right) \quad (12)$$

$$\pi_a^{hq} = p_a \left(x_{ap}^{hq} + \int_{x_{ap}^{hq}}^{x_{pb}^{hq}} \lambda^{hq}(x_i)\, dx + \varepsilon^p \right) \quad (13)$$

$$\pi_b^{hq} = p_b \left(1 - x_{pb}^{hq} + \int_{x_{ap}^{hq}}^{x_{pb}^{hq}} \Big(1 - \lambda^{hq}(x_i)\Big)\, dx - \varepsilon^p \right) \quad (14)$$

Then, from the system of first-order conditions, we derive the profit-maximizing prices:

$$p_a^{hq}(\varepsilon^p) = t - \frac{v_b}{3} + \frac{2\,\varepsilon^p\, t}{3}; \quad p_b^{hq}(\varepsilon^p) = t + \frac{v_b}{3} - \frac{2\,\varepsilon^p\, t}{3}; \quad p_p^{hq}(\varepsilon^p) = \frac{10t}{9} + \frac{v_b^2}{4} - \varepsilon^p(v_b - \varepsilon^p\, t) \quad (15)$$

The next Lemma follows directly:

Lemma 2. *Consider the case in which the platform offers a biased mix $\lambda^{hq}(x, v_b) + \varepsilon^p$ to the consumers. Then the stage 2 equilibrium prices are as derived in (15), the profits of the music labels and the platform are*

$$\pi_a^{hq}(\varepsilon^p) = \frac{(t(3 + 2\varepsilon^p) - v_b)^2}{18t}; \quad \pi_b^{hq}(\varepsilon^p) = \frac{(t(3 - 2\varepsilon^p) + v_b)^2}{18t}$$

$$\pi_p^{hq}(\varepsilon^p) = \frac{t + 3\varepsilon^p\,(7v_b - 13t\varepsilon^p)}{27} - \frac{v_b^2}{36t}$$

and the indifferent consumers are located in:

$$x_{ap}^{hq} = \frac{1}{3} - \varepsilon^p; \qquad x_{pb}^{hq} = \frac{2}{3} - \varepsilon^p; \qquad x_{ab}^{hq} = \frac{1}{2} - \frac{v_b}{6t} - \frac{2\varepsilon^p}{3}$$

Proof. See the Appendix. ■

Before proceeding backward to derive the sellers' entry decision and the equilibrium bias, let us discuss the results in Lemma 2. The bias affects the two sellers in opposite ways. Seller a benefits from employing a biased mix, as it allows it to sell more of its content to the platform subscribers, mitigating the quality gap. As seen in the benchmark case, the increase in the demand for the seller a's content generates a positive pressure on the price p_a. The indifferent consumer shifts to the left, but the price effect and the larger share of content a in the biased mix more than compensate for the reduction of demand on the direct channel.

Conversely, seller b suffers from the recommendation bias. Consumers are exposed to a lower-than-optimal level of content b on the platform. To compensate for this loss, seller b lowers the price p_b, inducing more consumers to purchase the pure bundle good b. However, the negative price effect and the reduced exposure of content b in the mixed bundle good dominate the demand expansion on the direct channel.

The platform does not lose demand but reshuffles its cost function more conveniently. It is worth mentioning that a positive bias $\varepsilon > 0$ makes sense provided that $p_b > p_a$, which in this case requires $v_b > 2\varepsilon^p t$. If that were not the case, the recommendation bias would backfire: if $\varepsilon^p > 0$ such that $p_a > p_b$, then the platform would find herself shifting demand toward the most expensive content, which is obviously non-optimal.

Recall that the bias is set before the game starts and that the platform commits to that level. Hence, once decided, it cannot be modified to adjust for the new ordering of the prices. Recall also that the bias cannot exceed the maximum one (Condition 1): the platform anticipates the effects of the bias on the entry decision of consumers and on pricing and sets it consistently with their participation constraints.

4.1.1 Sellers' participation decision

Let us now proceed backward and consider the sellers' participation decision given the bias ε^p. At stage 1, the sellers must decide whether to compete with each other in a market where $\alpha \in [0, 1]$ consumers are aware of their products and locations, or to join the platform and also reach the other $1 - \alpha$ consumers who are ex-ante unaware of the two sellers. As mentioned in Section 2, we refer to the former group as the "informed consumers". We refer to the latter group as the "uninformed consumers" instead. Uninformed consumers learn of the existence of the sellers or their relative position only if the platform is active, which can only happen if the platform manages to attract both sellers.[35]

[35]Consider the case in which only one seller $j = a, b$ joins the platform. Uninformed consumers learn about her and her position during the free trial of the streaming service. After the trial, they decide what to purchase (the subscription to the streaming service or the pure bundle). However, the platform operates as a retailer here

In all sub-games where at least one of the sellers decides not to join the platform, only the informed consumers are active. With no streaming service available, consumers cannot mix their consumption and are therefore limited to purchasing a pure bundle from either a or b. In these sub-games, sellers compete in a standard Hotelling setting. Given $v_b \geq 0$ and $\alpha \in [0, 1]$, equilibrium prices and profits when the platform is inactive are:

$$p_a^{out} = t - \frac{v_b}{3}; \quad p_b^{out} = t + \frac{v_b}{3}$$

$$\pi_a^{out} = \alpha\frac{(3t - v_b)^2}{18t}; \quad \pi_b^{out} = \alpha\frac{(3t + v_b)^2}{18t}$$

Where the apex ^{out} indicates the case where only consumption outside the platform is possible.

When seller $j = a, b$ decides whether to join the platform, he compares profit π_j^{out} and π_j anticipating equilibrium pricing and any consumption bias the platform might introduce. Notice that, compared to seller a, seller b has the better outside option if the platform is inactive. Moreover, b is the seller that would be penalized if the platform biased consumption. It follows that it is sufficient to consider the participation decision of b to determine whether the platform can be active or not in equilibrium. This decision depends on the share of informed consumers, α, and the quality difference, v_b.

In the baseline specification with homogeneous goods ($v_b = 0$), it is clear that the platform is always active: because firms make the Hotelling profits in equilibrium (Corollary 1) they are strictly better off if they are exposed to the uninformed consumers. In the limit case in which $\alpha = 1$ (that is, there are no uninformed consumers), moreover, sellers are indifferent between joining or not; in this case, we assume that the indifference is split in favor of the platform, which can then become active.

The prediction changes drastically if the products are vertically differentiated. In particular, the ability of the platform to bias consumption is limited in that it must induce both sellers and buyers to join. In other words, the equilibrium bias the platform can design is bound by two constraints: the consumers' participation constraints, addressed above, and the sellers' participation constraints.

Intuitively, the latter becomes stricter the higher α is. If there are many informed consumers, the high-quality seller has a stronger bargaining chip at the entry stage. Suppose that there are, in fact, no uninformed consumers — i.e., $\alpha = 1$. The platform's optimal bias policy would negatively affect music label b. Clearly, b anticipates that joining the platform does not expose his product to more consumers. Then, b would rationally choose not to join the platform if she commits to any positive level of bias. Therefore, the platform would reduce her optimal bias to zero to induce both sellers to join.

At the opposite limit, suppose that $\alpha = 0$: if the music label b does not join the platform, he cannot make any sale in the direct market as there are no consumers who are aware of her existence. Regardless of how biased the recommendation system is in favor of his rival, he would always optimally choose to join the platform. In turn, this implies that the platform is only

(it only offers the pure bundle of the music label, as there are no other goods to include in the mix). Because the royalty rate is $r_j = p_j$, and royalties enter the cost structure of the platform, it must be that the subscription fee $p_p \geq p_j$, which means all consumers weakly prefer purchasing the pure bundle j directly by the seller.

constrained by the consumer participation decision when setting up her bias policy.

Formally, we can distinguish a threshold for α as a function of v_b and the chosen bias ε^p that sorts between when it is profitable to join the platform and when it is profitable to operate only on the direct market.

$$\alpha^* = \{\alpha \in [0,1] \;\|\; \text{s.t. } \pi_b^{hq}(\varepsilon^p) = \alpha\, \pi_b^{out}\} \quad \Longleftrightarrow \quad \alpha^* = \frac{\pi_b^{hq}(\varepsilon^p)}{\pi_b^{out}} = \frac{((3-2\varepsilon)t + v_b)^2}{(3t+v_b)^2}$$

Because the platform is not active unless both sellers join the streaming service, she must then choose a bias such that:

$$\varepsilon \leq \varepsilon^s = \frac{(3t+v_b)(1-\sqrt{\alpha})}{2t} \tag{16}$$

where the apex ^{s} indicates it is the *sellers' constraint.*

4.1.2 Equilibrium Bias

We can finally proceed backward to stage 0 and determine the equilibrium level of bias that the platform includes in her recommendation system. From the analysis above, the problem of the platform can be written as

$$\begin{aligned} \max_{\varepsilon} \quad & \pi_p^{hq}(\varepsilon) = \frac{t + 3\varepsilon(7v_b - 13t\varepsilon)}{27} - \frac{v_b^2}{36t} \\ \text{subject to} \quad & \varepsilon < \min\{\varepsilon^c, \varepsilon^s\} \end{aligned}$$

It is easy to observe that the unconstrained maximization leads to $\varepsilon^p = \frac{7v_b}{26t}$. Using the prices in (15) to evaluate the maximum bias consumers are willing to accept before leaving the platform, we obtain:

$$\varepsilon^c = \int_{x_{ap}^{hq}}^{x_{ab}^{hq}} \frac{2t + 3v_b - \sqrt{72t^2x^2 - 4t^2 - 18tv_b^2 + 72tv_bx - 12tv_b + 27v_b^2}}{12t}\,dx + \int_{x_{ab}^{hq}}^{x_{bp}^{hq}} \frac{-2t + 3v_b - \sqrt{72t^2x^2 + t^2(68-144x) + 72tv_bx - 6tv_b(3v_b+10) + 27v_b^2}}{12t}\,dx$$

where the apex ^{c} indicates that it is the *consumers' participation constraint.* Using the location of indifferent consumers in Proposition 2, tedious calculations reveal that ε^c is decreasing in v_b.

Proposition 2. *In equilibrium, the platform sets a bias such that all sellers join the platform and the market is fully covered. Moreover, the equilibrium bias is:*

$$\varepsilon^* = \min\{\varepsilon^p, \varepsilon^c, \varepsilon^s\}$$

Proof. See the appendix. Moreover, see Figure 11 for an illustration of the result. ■

In choosing the bias intensity, the platform trades off cost minimization and the intensity of competition. Since biasing consumption away from b leads it to optimally reduce p_b, the plat-

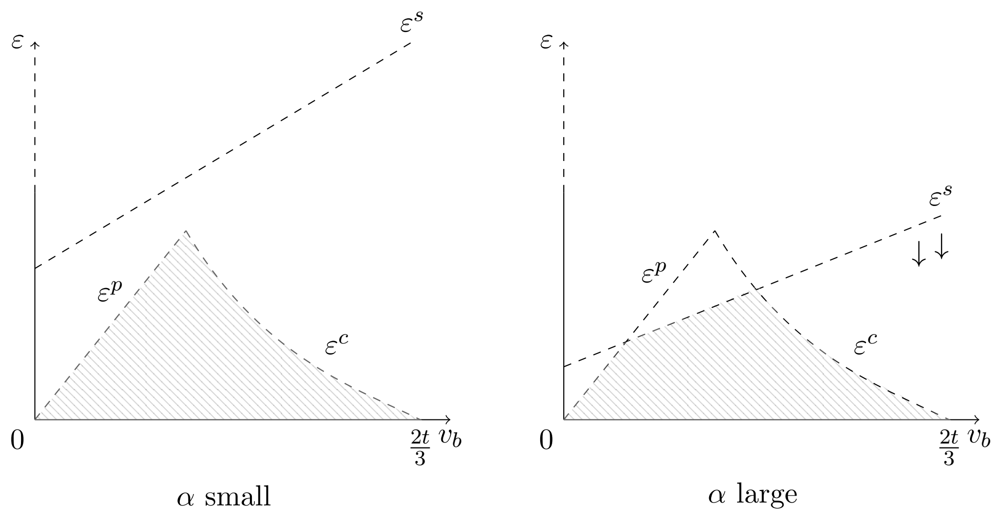

Figure 11: *The constrained equilibrium level of recommendation bias designed by the platform. When the share of informed consumers is sufficiently small (left panel), the optimal bias (ε^p) is bounded by the consumers' participation constraint (ε^c). When it kicks in, the platform needs to reduce the bias as quality increases because fewer users are willing to substitute content a with the high-quality good b. Otherwise (right panel), the platform must consider an additional constraint (ε^s) and needs to secure sellers' participation in the streaming service.*

form must update its own optimal price downwards because the better product has a stronger competitive effect than the worse one. The platform, therefore, wants to reduce the bias to set higher fees for consumers and raise the bias to minimize costs. The trade-off is optimally resolved at ε^p. The dynamic echoes the result found in Arya et al. (2007): since sellers compete with the platform on the margin, they set lower prices than they would if their only revenue stream was through the platform. These prices translate to lower operational costs for the streaming platform even without the adjustment that follows from the strategic use of the recommendation system discussed here.

Crucially, the platform is aware of several constraints. In particular, consumers and sellers must be offered conditions that make them want to join. Consumers' participation constraint (ε^c) arises because each biased bundle offered to each consumer must provide weakly more utility than what they could obtain outside of the platform.

Sellers face a different trade-off when choosing to join the platform or staying out: the more consumers are already aware of their existence and product, the less they benefit from exposure. Vice versa, if there are many uninformed consumers in the market, joining the platform boosts the visibility of the sellers and, consequently, their sales. This trade-off is incorporated in the constraint represented by ε^s. The higher α is, the harder it is to convince sellers to withstand the bias on the platform. This, in practice, limits the total bias that the platform can realistically design and commit to. Notably, this implies that the minimum level α that makes seller b strictly better off on the platform than outside of it is non-monotonic.

One can notice that the unconstrained bias is increasing in v_b: the larger the quality differential is, the stronger the effect it has on costs relative to the effect it has on prices, and the more the platform is interested in biasing consumption in equilibrium. Instead, the threshold ε^c decreases in v_b: the higher v_b is, the more consumers benefit from consuming the product offered by b, and the less flexible they are in accepting a biased bundle. To attract them, the platform must design its recommendation system such that the total bias never exceeds ε^c.

Finally, for the sellers to be better off on the platform, they need to be exposed to more new consumers to compensate for the lower individual exposure. When the consumer participation constraints start binding — that is, for v_b high enough — ε^* decreases. As a consequence, the high-quality seller's penalty generated by the platform gets smaller the better his product is since more consumers want to purchase more of its content.

Overall, the model highlights a subtle interaction between the business model of most streaming platforms and the incentive of creators to produce high-quality content. It is clear that absent the bias, sellers would have a strong incentive to innovate and compete with high-quality products: consumers value better content and are willing to pay more for it. The intermediation of the platform dilutes these incentives by introducing a bias that penalizes better content when it is more expensive to offer it. We assumed so far that V_j is exogenously given for both sellers. However, it is easy to extend the analysis to a hypothetical stage 0 in which the sellers had to costly invest in the creation of higher-quality content. In what follows, we show that unless the platform generates significant new demand (that is unless α is small), the incentive to costly invest in higher quality is lessened when the platform can bias consumption. Further, we show that the equilibrium allocation of R&D effort can be made inefficient by the platform's intervention.

4.2 Endogenous investment in quality

We now endogenize the choice of sellers to costly invest in quality. To maintain the direction of the asymmetry studied above, we assume seller b always to be more efficient than seller a. We consider two scenarios: in the first, we assume seller a to be incapable of investing in quality, while b is free to invest in quality at the beginning of the game. Afterward, we assume both sellers simultaneously select their quality investment from the same baseline value $V_j = V$. For tractability reasons, we ignore agents' participation constraints' effect on equilibrium quality levels when solving for the equilibrium investment in quality.[36]

Formally, the timing of the interaction between sellers and the platform is unchanged but augmented by an earlier stage (we refer to it as stage 0) in which sellers independently, and simultaneously, select v_a and v_b. To model this investment, we consider standard convex cost function $I(v_j) = \phi_j v_j^2$. Notice that in our simplified setting, a firm chooses how much to invest in R&D and these investments uniquely determine the quality of the good sold in the market, $v + v_j$. As a consequence, a firm maximizes profits by either choosing the amount of its investments, $I(v_j)$, or its degree of innovativeness, v_j; for this reason, with a slight abuse of terminology, in what follows, we often refer to v_j as the level of investment in R&D chosen by seller j.

[36]The aim of this exercise is to qualitatively evaluate the effect recommendation bias has on innovation investments. As discussed above, both consumer and seller participation constraint limits the role of the bias for high enough quality differential. We acknowledge this effect, but choose to ignore it for the sake of clarity of the exposition of the results.

The sellers' profits change to:

$$\pi_a^{eq} = D_a p_a - \phi_a \, v_a^2; \quad \pi_b^{eq} = D_b p_b - \phi_b \, v_b^2; \tag{17}$$

where apex ^{eq} stands for "endogenous quality", $\phi_a = \infty$ or $\phi_a = 1$ (depending on whether a is assumed to be able to invest in quality or not), and $\phi_b < 1$ captures the higher efficiency of seller b compared to a. We make the following technical assumption to ensure the existence of an interior solution:

Assumption 1. *The investment cost functions are sufficiently convex:* $\phi_b > \tilde{\phi} \approx \frac{0.1183}{t}$.

We update the formulas detailed in Lemma 2 by including the cost functions for the two firms and the cost differential $v_b - v_a$:

$$\pi_a^{eq}(\varepsilon^p) = \frac{(t(3+2\varepsilon^p) - (v_b - v_a))^2}{18t} - v_a^2; \quad \pi_b^{eq}(\varepsilon^p) = \frac{(t(3-2\varepsilon^p) + (v_b - v_a))^2}{18t} - \phi \, v_b^2$$

$$\pi_p^{eq}(\varepsilon^p) = \frac{t + 3\varepsilon^p \left(7(v_b - v_a) - 13t\varepsilon^p\right)}{27} - \frac{(v_b - v_a)^2}{36t}$$

As before, we obtain $\varepsilon^p = \frac{7(v_b - v_a)}{26t}$ by simple F.O.C. argument; we use the equilibrium platform bias to obtain equilibrium investment levels by plugging it in the equations for $\pi_a^{eq}(\varepsilon^p)$ and $\pi_b^{eq}(\varepsilon^p)$. Then, we proceed backward to stage 0, when investments in innovation are decided.

4.2.1 Only one seller invests

First, we directly extend the analysis brought forth, assuming v_b to be exogenously selected by endogenizing the choice of the more efficient seller with respect to the optimal quality differential on and off the platform. In this scenario, we impose $\phi_a = \infty$ to guarantee $v_a = 0$ in equilibrium. After plugging in $\varepsilon^p = \frac{7v_b}{26t}$ in the profit function of b, by F.O.C. we obtain the optimal investment in quality:

$$\left.\frac{\partial \pi_b^{eq}}{\partial v_b}\right|_{\varepsilon^p = \frac{7v_b}{26t}} = 0 \quad \Longleftrightarrow \quad v_b^{eq} = \frac{13t}{169t\phi_b - 2}.$$

Unsurprisingly, v_b^{eq} is decreasing in ϕ_b: the more efficient seller b is, the higher the optimal quality differential, all else being equal. To give meaning to this value, however, we must compare it with the relevant counterfactual. First, one should consider what the optimal investment level would be in the sub-game in which sellers do not join the platform. In this case, only a fraction $\alpha \in (0,1)$ of consumers is active. Thus, the rent that the sellers can generate depends on the size of this group. It is easy to show that the higher the number of informed consumers α is, the stronger the incentives to invest in quality outside the platform ecosystem. If we update the payoff of the innovative seller b when it decides not to join the platform ($\pi_b^{eq,out} = \pi_b^{out} - I(v_b)$), the standard maximization procedure yields the following value of investment:

$$\frac{\partial \pi_b^{eq,out}}{\partial v_b} = 0 \quad \Longleftrightarrow \quad v_b^{eq,out} = \frac{3\alpha t}{18t\phi_b - \alpha}.$$

As anticipated, for all (valid) values of ϕ_b, it holds:

$$\frac{\partial v_b^{*,out}}{\partial \alpha} > 0, \qquad \text{and} \qquad v_b\big|_{\alpha=0} = 0.$$

Second, we compare our equilibrium result with the profit-maximizing level of investments assuming that the platform is not able to bias the recommendation system ($\varepsilon^p = 0$). In this hypothetical scenario, both sellers have a clear incentive to join as argued in the previous subsection. Furthermore, as shown above, when the platform does not introduce a bias, the competitive stage is equivalent to a standard Hotelling framework (Corollary 1).

It follows that in this scenario firm b optimally invests:

$$v_b^{eq,nobias} = v_b^{eq,out}\big|_{\alpha=1} > v_b^{eq} \; \forall \phi_b,$$

where the last relationship is established by direct comparison.

In summary, firm b anticipates the effect of the bias and lowers its investment in quality v_b. The result above highlights that for some α joining the platform may reduce the incentives to invest in innovation. This happens if the share of informed consumers is sufficiently large. In this case, the positive shock in demand from the additional uninformed users who only listen to music via the platform is insufficient to compensate for the demand shift towards the low-quality rival entailed by the biased recommendation system. However, for low enough values of α, the opposite holds. The increased demand creates strong enough incentives to invest despite the distortion introduced by the platform.

4.2.2 Both sellers invest

Suppose now that both sellers are able to costly invest in quality, that is, $\phi_a = 1 > \phi_b$. After plugging in ε^p in the equations above, it is once again straightforward to obtain the equilibrium levels of investment by simple F.O.C. argument. It is also immediate to obtain the equilibrium level of investment in the sub-game in which the platform is inactive. Remember that, in this case, only a fraction α of consumers is active and sellers compete in standard Hotelling. Then:

Lemma 3. *Consider the case in which both sellers can invest in quality, and seller b is more efficient than seller a by a factor ϕ_b. Equilibrium levels of investment on the platform are:*

$$v_a^{eq} = \frac{169t\phi_b - 4}{(2197t - 26)\phi_b - 26}; \quad v_b^{eq} = \frac{169t - 4}{(2197t - 26)\phi_b - 26};$$

equilibrium levels of investment outside the platform is:

$$v_a^{eq,out} = \frac{\alpha(9t\phi_b - \alpha)}{3((18t - \alpha)\phi_b - \alpha)}; \quad v_b^{eq,out} = \frac{\alpha(9t - \alpha)}{3((18t - \alpha)\phi_b - \alpha)};$$

and equilibrium levels of investment, if the platform is assumed not to bias, is:

$$v_a^{eq,nobias} = v_a^{eq,out}|_{\alpha=1}; \quad v_b^{eq,nobias} = v_b^{eq,out}|_{\alpha=1}.$$

Proof. See the Appendix. ■

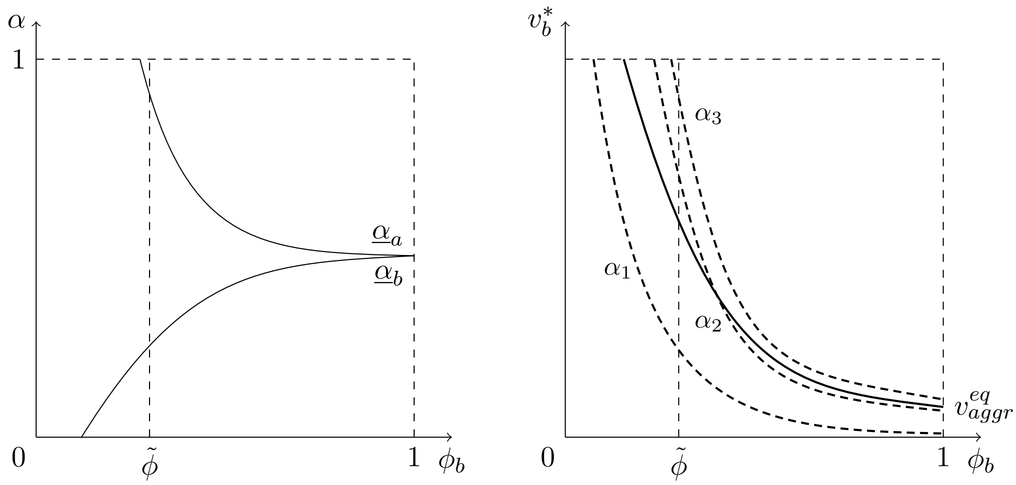

Figure 12: *On the left: the minimum value $\underline{\alpha}_a$ and $\underline{\alpha}_b$ such that, for $\alpha > \underline{\alpha}_j$, seller j would have invested more in quality outside of the platform; On the right: aggregate investment level for increasing values of α ($\alpha_1 = 0.1$ $\alpha_2 = 0.45$ $\alpha_3 = 1$). $t = 1$ in both graphs.*

The bias imposed by the platform distorts the incentives to innovate even if we allow both sellers to invest in innovation. In particular, for given demand, the intervention of the platform dilutes the incentives of seller b to innovate, which in turn leads to an equilibrium investment v_a lower than the one a would have selected in the absence of bias. This can be seen by direct comparison of v_a^{eq} and $v_a^{eq,nobias}$, and of v_b^{eq} and $v_b^{eq,nobias}$. The result is once tempered when we account for the additional demand generated by the platform: if α is low enough, joining the platform exposes the labels to enough new demand that, despite the bias, equilibrium investment in quality is higher on the platform than off of it.

The equations in Lemma 3 allow us to make some statements regarding the overall effect on aggregate investment of quality, which can be shown to crucially depend on the relative values of α and ϕ_b. In particular, the lowest value α such that equilibrium aggregate investments are higher off the platform than on it is increasing in ϕ_b: the more marked the difference in efficiency is, the stronger the distortion on equilibrium investment levels introduced by the platform is. Vice versa, for α high enough, aggregate investment would be higher outside of the platform than it is on the platform. Furthermore, For all values of ϕ_b, there are values α such that b would have invested more and a would have invested less outside of the platform than on it. The last remark implies that platform intervention can also lead to inefficient allocation of R&D effort by innovators, with efficient (resp. inefficient) sellers investing less (resp. more) than they would have without the bias in the recommendation system. The results are illustrated in Figure 12 and summarized as follows:

Proposition 3. *If the platform does not generate enough new demand, recommendation bias leads to lower aggregate investment in innovation; the distortion is stronger when the difference in efficiency between labels is not too high. Furthermore, for all values $\phi_b \in (\tilde{\phi}, 1)$, there exist values α such that equilibrium allocation of R&D effort is inefficient.*

Proof. See the Appendix. ■

5 Extensions

We extend the analysis in several directions. First, we consider the implications of letting consumers costly produce their own favorite mix of content on the platform. We assume that the associated cost is a design choice of the platform and, therefore, is selected strategically at the beginning of the game. We show that the platform has the incentive to increase the said cost to force consumers to use her recommendation system, which follows seamlessly from the analysis above.

Second, we consider a different timing of the interaction: following Bourreau and Gaudin (2022) we assume that the recommendation system is not set up at the beginning of the game but, rather, after the agents have selected prices. We take a reduced form approach and show that the platform has the incentive to bias more than in the baseline model since agents cannot condition prices on the equilibrium recommendation system.

Finally, we consider a different source of vertical differentiation in the form of asymmetric costs of production by the sellers. We show that since the most efficient seller selects lower prices in equilibrium, the platform optimally biases consumption towards him. On one hand, this suggests that streaming platforms reward efficiency. We argue however that this, too, indicates that streaming platforms penalize experimentation and, instead, create the incentive to produce commodified content to avoid the added penalty linked to inefficiency.

On-platform search. Our baseline model relies on the implicit assumption that consumers are passive agents when it comes to the design of the mix they consume. In other words, we assume that consumers join the streaming service and observe the bundle offered by the platform, with no opportunity to modify it.

In real-world examples, this is hardly the case. Consumers, once they join a streaming platform, have some freedom in choosing the music or digital content they want to consume.

It is however apparent that the extent to which it is easy to search for the desired digital content is part of the design of a platform, of its architecture. Hence, a natural question arises: what is the optimal "degree of freedom" that the platform should give to consumers to build their own consumption mix?

In this extension, we address this question by assuming that the consumer may (or may not) decide to pay a cost in order to get the efficient bundle and get rid of any recommendation bias imposed by the platform.

Formally, the utility of the consumers when they join the platform becomes:

$$U_{i,p} = \begin{cases} v + (1 - \lambda(x_i) - \varepsilon(x_i))v_b - p_p - t(x_i - (1 - \lambda(x_i) - \varepsilon(x_i)))^2 & \text{if she does not pay} \\ v + (1 - \lambda(x_i))v_b - p_p - k - t(x_i - (1 - \lambda(x_i)))^2 & \text{if she pays } k > 0 \end{cases}$$

where k is the search cost that consumers must incur in order to avoid recommendation bias. Keeping in mind that the efficient bundle is $\lambda^{hq}(x_i) = 1 - x_i - \frac{v_b}{2t}$, it is possible to derive the cost level that makes any consumers on the platform indifferent between building consuming the efficient mix and get the biased bundle:

$$\bar{k}_i = t\big(\varepsilon(x_i)\big)^2$$

If unconstrained, the platform can simply design the search process in a way such that $k = \bar{k}_i$ for the consumer i who is willing to accept the largest bias (who is located in x_{ab}). Thus, the consumers are induced to accept the bias set by the platform, and the analysis in the previous section goes through. Otherwise, if constrained in some ways, the platform has to set a smaller k. In such a case, the consumers located closer to x_{ab} will be willing to pay the search cost in order to get the efficient mix. Thus, the ability of the platform to bias the recommendation is limited to consumers at the margins of the platform's demand.

Under our assumption of a monopoly platform, we observe no incentives to design its search process in a way that limits its ability to bias the recommendation system and save on costs. However, how efficient the search process looks to consumers may be a relevant characteristic when competition between platforms kicks in. We leave this question open to further research.

Different timing. In our baseline model, we adopt a timing that implies sellers are ex-ante aware of the ex-post level of bias because the platform announces it in the first stage and commits to it. The question of what would occur were the sellers unaware of the recommendation bias naturally emerges.

In this section, we address this issue assuming that the platform does not commit to a bias level, but firms are aware of its incentives to steer demand toward the cheapest product. Hence, the new timing is the following: at stage 1, the sellers anticipate the bias level and, simultaneously with the platform, set the prices. Then, at stage 2, the platform observes prices and quality and adjusts its recommendation bias. Finally, consumers make their choices, and payoffs realize.

In the baseline specification with homogeneous goods, the platform offers an efficient mix to all consumers, and the game unfolds as in Proposition 1. However, Proposition 2 reveals that prices are not symmetric if there is a variation in the qualities of the goods offered. Hence, the platform finds biasing its recommendation system profitable, as there is a demand mass that can be steered toward the cheapest (inferior) good without impacting the number of consumers joining the streaming service. Moreover, as sellers are not able to react after the choice of the recommendation bias, the platform would always steer as much demand as possible, subject to the consumers' participation constraint (see the definition of $\bar{\varepsilon}(x_i)$ in Section 4).

The sellers anticipate this incentive and adjust their prices accordingly. On the one hand, the seller of the superior good reacts to the anticipated bias by lowering her price. On the other hand, the seller of the inferior good, anticipating the demand expansion following the recommendation bias, has the incentive to increase its price. Eventually, this adjustment occurs until profitable.

One should notice that an equilibrium exists only provided that the prices do not "cross" - i.e., provided that the most valuable good is also the most expensive one.

Otherwise, assume the sellers anticipate the bias and adjust the prices to such an extent that the inferior good is now as expensive as the superior one. The platform observes the prices and optimally reacts by recommending the cheapest high-quality content to more consumers. However, this reaction goes in the opposite direction of what was anticipated by the sellers, who would like to change their strategies ex-post.

In other words, consider the prices $p_a(\varepsilon)$ and $p_b(\varepsilon)$ chosen by the sellers in anticipation of the total recommendation bias. We can state the following:

Lemma 4. *Assume that* $p_a|_{\varepsilon=0} < p_b|_{\varepsilon=0}$. *Then if the maximum bias that the platform can impose* (ε^c) *is such that* $p_a(\varepsilon^c) < p_b(\varepsilon^c)$, *an equilibrium exists in which the platform sets the maximum bias and the price difference between the two sellers shrinks. Otherwise, if the bias is such that* $p_a(\varepsilon^c) \geq p_b(\varepsilon^c)$, *an equilibrium in pure strategy does not exist.*

One should notice that because ε^c is decreasing in v_b, we can say that an equilibrium in pure is more likely to emerge when v_b is sufficiently large.

Asymmetric costs. We now consider a different source of vertical differentiation, namely an asymmetry in the cost functions of sellers a and b. The quality of the content produced by the sellers is now constant and equal to v, assumed to be high enough to guarantee full coverage. Sellers maximize:

$$\pi_i = D_i(p_i - C_i), \qquad i \in \{a, b\}$$

where C_i is a measure of the marginal cost of producing the content sold by a and b. To preserve the direction of the asymmetry in the main model, we assume that $C_a = c_a > 0$, $C_b = 0$.

The framework differs from the main model in two substantial ways. First, since the asymmetry lies in the costs rather than the value consumers attach to the content, the optimal mix of consumers is not affected by the asymmetry, and $\lambda^*(x_i) = 1 - x_i$. Second, when the better seller is more efficient rather than offering higher quality content, the price he would optimally select in the absence of bias would be lower than the one offered by his competitor. The platform, therefore, would have an incentive to penalize the worst of the two sellers with her biased recommendations rather than the best one as it was in the main model.

The analytical steps to solve the model mirror the ones made explicit for the main model. As before, the difference in equilibrium prices is tempered by the intervention of the platform: ε^* is selected to reduce the distance between p_a and p_b. Unlike in the main model, however, since $C_a > C_b$, $\varepsilon^* \leq 0$: the platform introduces bias to boost consumption of seller b's content. In doing so, she induces higher p_b and lower p_a to emerge in equilibrium compared to the case in which bias was not introduced. The platform balances the incentive to increase the bias to reduce operational costs and reduce to bias to increase her own subscription fee p_p.

Finally, since both consumers and sellers must choose to join the platform, constraints ε^c and ε^s still have to be taken into account. Unlike in the main model, consumers do not want to consume the content of one of the two sellers disproportionately. Since the bias is introduced to penalize the content of the least efficient seller, and since this seller charges a higher price in equilibrium because of this inefficiency, consumers are less sensitive to the bias than before. It follows that the constraint represented by ε^c is still decreasing in the cost differential, but is less tight than in the main model.

The constraint relevant to induce sellers to join is also looser than the one considered in the main model. Unlike before, since the bias penalizes the inefficient seller, the highest possible bias the platform can introduce must make the worse seller indifferent between joining or not, rather than the better one as it was before. From standard Hotelling logic, the seller with higher marginal costs would see his profits decrease in the cost differential. Since the seller penalized by the bias is the one with the worst outside option, then, it is clear that the platform can ignore the constraint represented by ε^s for a wider range of values α.

We can state the following result:

Proposition 4. *When sellers a and b have different marginal costs of production, the platform introduces a positive bias in favor of the most efficient of the two and increases his equilibrium profits. In particular:*

$$\varepsilon^* = \min\{\varepsilon^p, \varepsilon^c, \varepsilon^s\}$$

where

$$\frac{\partial|\varepsilon^p|}{\partial\Delta_c} > 0, \qquad \frac{\partial|\varepsilon^c|}{\partial\Delta_c} < 0, \qquad \frac{\partial|\varepsilon^s|}{\partial\Delta_c} < 0,$$

and $\Delta_c = |c_a - c_b|$.

Proof. See the Appendix. ∎

This last exercise serves two purposes. On one hand, it highlights the difference between vertical differentiation driven by consumer taste and efficiency: in particular, the two approaches generate bias of opposite signs (favoring the least liked content and the most efficient seller respectively). Considering only efficiency as a driver of asymmetry might lead to the partial conclusion that streaming platforms' intervention might be socially desirable since it creates the incentive to reduce marginal costs or production and, with them, equilibrium prices.

On the other hand, however, we believe that modeling costs as we have proxies the choice of music labels to experiment with their content, which can be expected to increase costs of production, instead of optimizing the process of creating said content. From this perspective, the outcome of the exercise is well in line with the one presented in the main model: the platform discourages risk and, instead, rewards "assembly line" production of content. The overall takeaway, then, becomes straightforward: if we assume as true that novelty and experimentation are costlier than producing commodified content, the platform always has the incentive to penalize it through strategic manipulation of what consumers are exposed to.

6 Discussion and conclusion

In this paper, we study the incentives of a streaming platform to bias bundling in an effort to achieve optimal economic conditions. The platform has the potential to generate utility for consumers who value a balanced mix of content. When content is of equal quality, sellers select uniform prices, and the platform has no incentive to bias consumption. When sellers offer vertically differentiated products, instead, they have the incentive to set different royalties. In particular, the seller with the higher quality product wants to raise royalties since consumers value his product more. When this happens, the platform has an incentive to bias consumption towards the "cheaper", lower quality product to minimize costs. This comes at the detriment of consumers, who lose the additional utility generated through efficient content mixing, and the higher quality seller, who sees his demand artificially shrunk. In equilibrium, the latter would set a lower price than without intervention: the platform dampens the incentive to introduce higher quality products by punishing them with reduced exposure. Further, platform intervention can significantly distort equilibrium R&D efforts.

Based on several real-life examples, we assumed that the platform cannot price discriminate consumers. If she could, it is clear that she would have the incentive to offer different bundles at different prices in an effort to extract the rent she helps generate. The ability to price discriminate does not eliminate the incentive to bias. However, since consumers must be convinced to join the platform, personalized pricing would remove the ability to bias consumption. Price discrimination and consumption bias are substitute strategies. If personalized pricing were possible, the higher-quality seller would be better off in equilibrium. On the other hand, consumers would be as well off if products were vertically differentiated; they would also be strictly worse off if products were of the same quality. The reason is straightforward: the platform has no incentive to bias consumption under the baseline specification, but she would still have the incentive to price discriminate if it was possible.

More subtly, the result is carried forward by the assumption of sellers bargaining their royalty rate individually. The incentive to bias consumption follows directly from the difference in cost for the platform to stream the content of the sellers. Suppose, however, that the sellers were both represented by an intermediary, such as a copyright collecting agency, bargaining royalty rates for both. It is clear that such an agent would have the incentive to set equal royalties to reduce the incentive to bias consumption towards the cheaper product. It is less clear that this would not be to the detriment of the higher quality product's seller.

The model, as it stands, is limited by the assumption that sellers set prices and royalties equally. Separating the two, letting sellers set prices outside the platform and royalty rates inside of it, might lead to new insights.

On the one hand, it is natural to think that the platform would exploit this opportunity by inducing sellers to compete for its internal demand. The sellers would, accordingly, lower their royalties, as they would be trapped in a prisoner-dilemma-like situation. Possibly, the platform's subscription fee may approach the prices of the seller, forcing them to lower their prices also on the captive segments.

On the other hand, one can also envision that sellers may have the incentives to increase their prices in their *captive* segments to increase the surplus extraction from their loyal consumers. In order for this strategy to be profitable, the sellers have to increase their royalties, which enters the platform objective function as marginal costs. By doing so, the sellers ensure their price cannot be matched by the platform and operate as monopolists in their captive segments. The platform might be limited in its possibility to bias demand in this case, as sellers would be using royalties to increase their rival's costs rather than to extract surplus directly. A formal analysis of these scenarios is left as an open question for further research.

It is important to stress that the mechanism studied in this paper requires the platform's algorithmic component to be relevant. In a world in which consumers had no access to automatically generated content susceptible to manipulation but were always in perfect control of what they consume, the distortions predicted by the model would not bite. From the opposite perspective, if the algorithmic recommendation of streaming platforms were provably biased in a way that damages consumers, regulatory intervention would prevent such distortions from arising. The broadest takeaway from the paper follows in the footprints of many others, calling for inspection and direct regulation of the algorithms used by digital platforms to provide their service. While we acknowledge that this might disincentivize R&D expenditure and innovation to ameliorate these algorithms, perhaps the loss could be more than compensated by the stronger incentives to innovate on the content that algorithms would no longer be able to penalize.

It must be noticed that, in our model, the platform has to provide the sellers with at least the same payoffs they would gain if they were not joining the streaming service. Hence, profit-wise, sellers' participation in the platform implies higher incentives to invest in content quality. Therefore, our analysis examines the distortion of the investments from the *potential* level the seller would select absent the recommendation bias.

In doing so, we take an overly benevolent view towards the platform, which, by construction, never directly harms sellers. Consistently, the shares of informed and uninformed consumers are exogenously determined and, in our model, proxy the popularity of music labels. Debuting music labels are unknown by the large public and would likely be unable to strive outside the platform service. On the contrary, established music labels benefit from a large network of consumers interested in their content and willing to purchase it regardless of whether they join the streaming service or not.[37]

Evaluating the net effect on the music industry imposed by the advent of the platform is well beyond the scope of our model. Here, our goal is to understand whether the platform can distort sellers' incentives to invest in quality by changing the algorithmic design of the recommendation system.

Abstracting from the model, it is important to consider the implications of the ability of the platform to make content accessible to more consumers. It is difficult to argue that streaming platforms do not represent a significant portion of the market they host. It follows that many music labels, especially new ones, have little hope of reaching the public without being hosted on one of these platforms. Joining, however, requires coming to terms with the ability of the platform to act as a gatekeeper. If music labels and content creators need the platform to reach interested consumers, and the platform is designed to punish good content if it comes at a higher price, incentives to vertically differentiate weaken. The model, then, highlights a potential risk embodied in the platform ecosystem: if a significant share of the users is held "captive" by the platform, the content available for streaming may become more commodified, which occurs at a loss for society that is difficult to quantify.

[37]Notably, Joni Mitchell and Neil Young's collections are not available on Spotify since 2022. Their motivation to delist from Spotify was not driven by commercial disputes but rather by the debate of Covid-19 misinformation in Joe Rogan's controversial podcast streamed on the platform. Yet, their decision to abandon the streaming service reflects the larger freedom celebrities enjoy vis-a-vis beginners.

Chapter III

Platform-enabled information disclosure

Joint with Martin Peitz

1 Introduction

With the growth of e-commerce and advances in data collection and storage, information brokers are collecting detailed information about consumer behavior. Some of these information brokers are digital platforms that operate marketplaces for sellers and consumers to trade with each other. Consumers may decide to which extent information on them can be collected or passed on to sellers. A platform that exclusively makes revenues from sellers may be thought of acting in the interest of sellers and thereby adjusting information design accordingly. However, if consumers foresee that they will receive unattractive terms some of them will stay away from the marketplace forcing the platform to balance seller and consumer interests. How does a monopoly platform choose price and information design when sellers can use disclosed information to price discriminate between consumers? How does the platform's choice compare with what is best from the seller and/or the consumer perspective?

We embed Ali et al. (2023)'s model with monopoly or competing sellers into a platform setting with heterogeneous consumers and a continuum of product categories and derive the platform-preferred information design. We contrast this design with the ones preferred by consumers and sellers. In our setting, the platform monetizes on the seller side; that is, consumers do not pay to visit the marketplace. The platform decides how much freedom to give to consumers to voluntarily disclose relevant information about their valuations to interested sellers. Within the disclosure regime chosen by the platform, consumers have control over their data – this is motivated by the recent activities of regulators aimed at empowering consumers in digital markets. With some information disclosure and seller prices that condition on this information, there is third-degree price discrimination. The platform's information design affects the sellers' pricing strategy and, thus, determines how attractive trade on the platform is for consumers and sellers. From the viewpoint of the platform, the chosen information design determines the strength of the cross-side network effect exerted from consumers on sellers and vice versa.

Following Hagenbach and Koessler (2017) and Ali et al. (2023), we distinguish between two disclosure regimes: when consumers have the binary choice of fully disclosing or not disclosing information at all, we refer to "simple evidence"; when consumers have more control over how much information to share and can disclose that they belong to a group of consumers with certain characteristics, instead, we refer to "rich evidence". Simple evidence can be exemplified by the practice on online marketplaces to collect tracking cookies from various websites. Then, consumers can choose to be tracked (full disclosure) or not be tracked at all (no disclosure). Rich evidence can be thought of as the act of selectively deleting cookies from some websites that reveal a particular information: for example, eliminating cookies collected by some air travel companies before looking to purchase a ticket can blur a higher valuation for a flight.

The effect of voluntary disclosure on consumer surplus has been analyzed by Ali et al.

(2023) in a single product market in which consumer participation is exogenous. Translated into a platform setting, when all consumers join the platform, sellers and the platform (that absorbs part of the sellers' profits) are best off if consumes can not disclose any information. By contrast, with endogenous consumer participation, a platform that monetizes on the seller side must take two effects into account. The direct effect by which information disclosure is price-reducing is counterproductive from the viewpoint of the platform. This tends to make "no evidence" the platform's preferred choice. However, in the presence of mutual cross-side network effects (that is, sellers benefit from more consumers participating and consumers benefit from a larger set of available product categories), enabling price-reducing information disclosure tends to attract a larger number of consumers. This indirect effect may be countervailing and dominate the direct effect. Then, the platform will enable voluntary information disclosure.

We explicitly model the consumer-seller interaction when sellers are either monopolists or duopolists. A seller's product belongs to a product category that is drawn from a continuum and each consumer is interested in one of them. We assume that, when joining the platform, consumers do not yet know their product category of interest. Consumers have heterogeneous valuations for the product(s) in their product category of interest, but discover their valuation only upon joining the platform. They make their participation decision based on the number of product categories carried by the platform, as this determines their expected valuation.The information design affects the incentives to join for consumers and sellers directly, through changes in the expected gains of joining, and indirectly, through the effect on the participation decisions of the other side. Our functional form assumptions allow us to obtain explicit solutions.

Under monopoly, different disclosure regimes may generate different expected total gains from trade of a transaction, and all economic actors may benefit from consumers being allowed to disclose information. Under seller competition in the classic Hotelling duopoly model with linear transportation costs, all disclosure regimes generate the same total gains from trade and only affect how these gains are shared between the different economic actors. We spell out the conditions under which the incentives of consumers, sellers, and platform are misaligned. Consumers tend to be the ones who benefit from a disclosure regime that gives more possibilities to them to reveal some information on their valuations to sellers. When buyers and sellers coincide in benefiting from a richer disclosure regime, it may still not be in the interest of the platform to enable such voluntary information disclosure.

However, consumers' and seller's preferences about the information design are sometimes aligned. In the monopoly seller case, when the expected gains from trade of a transaction are high, consumer participation is particularly valuable, and sellers (and the platform) benefit from richer voluntary disclosure through the subsequent increase in trade volume, as do consumers. In the duopoly case, when gains from trade on the platform are relatively low, sellers benefit from voluntary disclosure despite the harsher competition that follows, as the increase in consumer participation more than compensates the reduction of profit margins. Alignment on no disclosure is also possible: when gains from trade are relatively high, consumers would be better off without the ability to disclose their valuation, as this would lead to more seller participation and, thus, an increase in the number of product categories available on the platform. Then, consumer, seller, and platform interests are aligned.

Information design falls within a broad set of platform design decisions. For instance, Crémer et al. (2019, p. 60) in their report to the European Commission note that "platforms impose rules and institutions that reach beyond the pure matching service [...], e.g., [...] by regulating access to information that is generated on the platform, imposing minimum standards [...] Such rule setting and 'market design' determine the way in which competition takes place." We show that the platform chooses a disclosure regime that may be different from the one preferred by consumers and/or sellers (in a second-best sense according to which the platform always sets the platform fee). We find that a regulator maximizing consumer or seller surplus (or a convex combination thereof) may want to force the platform to enable information disclosure. More broadly, our results speak to whether and which public interventions may increase consumer welfare relative to the platform's self-regulation of information disclosure.

Literature review. This paper relates to two strands of literature: the economics of two-sided platforms and the literature on information design.

In this paper the platform plays the role of an information designer: as emphasized in Hagiu and Wright (2015), a multi-sided platform is special in its ability to shape interaction and communication between consumers and sellers. The approach we take is related to Teh (2022) in which a platform chooses a governance design that affects a consumer's transaction benefit and a seller's markup. It is also related to Choi and Jeon (2023) in which a platform makes investment decisions which affect consumers and sellers differently.[38] Our setting differs in two ways: we look at a particular, discrete platform governance decision different from the ones considered by Teh (2022) and Choi and Jeon (2023) and focus on mutual cross-side network effects thereby making consumers and sellers taking into account each others' participation decision.[39]

We consider a platform that manages the interaction between sellers and consumers. When sellers offer substitutes (as in our version with duopoly sellers), seller competition affects the expected surplus of consumers and sellers per transaction, and several papers have looked at the platform managing seller competition by limiting seller access to the platform (e.g., Nocke et al., 2007; Hagiu, 2009; Belleflamme and Peitz, 2019; Karle et al., 2020; Teh, 2022).[40] Intra-platform competition is also incorporated in recent work on hybrid platforms that allows the platform to be vertically integrated with one of the seller (see e.g. Anderson and Bedre-Defolie, ming).

According to earlier work outside a platform context, consumers may benefit from information disclosure as this may increase the competitive pressure among competing sellers under spatial price discrimination (e.g. Thisse and Vives, 1988) and behavior-based price discrimination (e.g. Fudenberg and Tirole, 2000). Competitive pressure is also affected by voluntary information disclosure in the model by Ali et al. (2023) with competing sellers, which is limited

[38] See Belleflamme and Peitz (2021, chap. 6) for a discussion of platform design decisions.

[39] Our focus is on regulating the platform's information design. In a different vein, Jeon et al. (2022) consider regulatory intervention regarding platform liability; they do so in a setting in which a copy-cat seller may free-ride on the investment by a brand manufacturer.

[40] Another way in which the platform may shape seller competition is to steer consumers to particular products or reduce the visibility of others. Johnson et al. (2020) consider a platform's demand-steering policy when sellers use Q-learning algorithms for their pricing strategy and show that such policies can undermine seller collusion. Casner (2020) shows that a platform with an exogenous proportional fee on seller revenues benefits from obfuscating search as this leads to higher seller mark-ups.

by the disclosure technology. We take a platform perspective according to which the platform not only charges sellers but is also an information designer who may limit the extent to which consumers can disclose information. With exogenous consumer participation, the platform would choose the disclosure regime that maximizes seller profits. If consumer participation is endogenous, as in our model, this is no longer necessarily the case and leads to richer results on the platform's information design. Thus, our paper contributes to the growing literature studying the effect of information and privacy on market outcomes.[41]

Voluntary disclosure can be placed between cheap-talk (studied in a monopoly platform setting in Hidir and Vellodi, 2021) and information-based mechanisms (such as the segmentation strategy an intermediary can commit to as in Bergemann et al., 2015, or the recommendation system studied in Lefez, 2022). Voluntary disclosure as in Ali et al. (2023) and, in their footsteps, our paper differs from other works in which the action is purely on the seller side (Armstrong, 2006b; Liu and Serfes, 2004; and Thisse and Vives, 1988). Such settings are related but ultimately fundamentally different. Our contribution also differs from Bounie et al. (2021) in which a data broker decides which consumer information to provide to sellers. In their setting, the data broker withholds some information on consumers to soften seller competition. Finally, the framework differs substantially from the one studied by Armstrong and Zhou (2022). The authors compare optimal signal structures from the perspective of consumers and sellers when consumers are uncertain about their valuation, but do not allow for voluntary communication between consumers and sellers. In contrast, following Ali et al. (2023), we assume that consumers decide whether and which information to disclose about their valuations.

It matters whether consumers disclose information or sellers acquire information about consumer tastes because the two sides' incentives may be opposed: Ali et al. (2023), Armstrong and Zhou (2022), Bergemann et al. (2015), and Elliott and Galeotti (2019) tell us that, with horizontal differentiation between sellers, consumers have an incentive to share information if located relatively far from both sellers, as this leads to lower prices. Sellers, on the other hand, have an incentive to learn the location of consumers located at the extremes (either very close or very far from their location) to better extract surplus from those consumers.

The rest of the paper proceeds as follows: in Section 2, we present the platform model. In Section 3 we consider the version with monopoly sellers and characterize equilibrium outcome and compare it with the equilibrium choices made by a regulator who can select the disclosure regime but not interfere with the selection of the entry fee (second best outcome). In Section 4, we analyze the version in which product categories are served by duopoly sellers. In each of the two sections, we define the admissible parameter space and discuss the relationship between our framework and that of Ali et al. (2023). Section 5 concludes. Appendix A collects calculations

[41] See Acquisti et al. (2016) and Bergemann and Bonatti (2019) for comprehensive reviews. In particular, Montes et al. (2019) study duopoly competition in a setting in which sellers can use information on consumer tastes to price discriminate, but in which consumers can prevent this use by opting out at a cost; this corresponds to the setting with simple evidence if that cost is zero. They investigate a data broker's incentives to exclusively sell the information it collected from consumers to one of the two sellers. Ichihashi (2020) considers a multi-product monopoly sellers who can price discriminate conditional on the disclosure decision by a consumer. The consumer chooses upfront a disclosure rule. In line with Ali et al. (2023), by withholding information about which product is most valuable, the consumer can induce the seller to set lower prices. If the seller has the option to commit to not use the consumer's information for pricing it prefers to use this option.

on the different disclosure regimes in the version with monopoly sellers and Appendix III in the version with duopoly sellers. Appendix III contains supplementary material on the consumer-seller interaction when sellers are duopolists.

2 Model

We consider a monopoly platform that facilitates trade between consumers and sellers. The platform operator manages the platform by setting a uniform seller fee, which is a proportional tax α_s on a seller's revenue, and choosing an information disclosure regime.[42] The prevailing disclosure regime (together with the platform fee) determines how large is the realized surplus in the consumer-seller interaction and how it is split between consumers, sellers, and the platform.

There is a unit mass of product categories and a unit mass of consumers. Sellers are single-product firms that have zero marginal costs of production. Each product category is associated with an opportunity cost f_s drawn from the uniform distribution $U[0, \overline{f}_s]$. Sellers decide whether or not to enter after observing the opportunity cost in their product category, the uniform fee, and the disclosure regime. In Section 3, we analyze the monopoly version of the model – that is, only one seller per product category can join the platform; in Section 4, we analyze the duopoly version of the model – that is, two sellers can enter per product category and, if they do, they compete with horizontally differentiated goods located at the endpoints of the unit interval.

Consumers are heterogeneous in three dimensions: each consumer is interested in exactly one product category (which is drawn with equal likelihood from the continuum of product categories), each consumer has a particular taste regarding the product or products in the product category that is of interest to this consumer, and each consumer has an opportunity cost f_b of joining the platform.

Consumers learn only f_b before they decide whether to join. We consider two groups of consumers. The fraction β of consumers are eager to join and, thus, have zero opportunity cost of joining ($f_b = 0$); thus, they are always present on the platform and β constitutes the minimum network size from the sellers' perspective. The fraction $(1-\beta)$ of consumers are hesitant to join; they draw their opportunity cost from the uniform distribution $U[0, \overline{f}_b]$ independently across consumers. Opportunity cost f_b is private information of the consumer when deciding whether to join.

Consumers do not know whether the product category that they are interested in is available on the platform. Two interpretations are compatible with this setting: (i) Consumers need to search on the platform to figure out which product category is the one they like or (ii) consumers know the product category of interest, but do not know whether it is carried by the plaform. In either case, consumers form their expectation on how likely it is that they will find a match based on the number of product categories available on the platform, which is assumed to be observable prior to joining. After joining the platform, these consumers learn whether their product category of interest is available and their willingness to pay or location: if their product category of interest is not represented on the platform, they do not purchase anything. If it is, in the monopoly version, a consumer has valuation v for the available product

[42] In reality, we observe very limited fee discrimination by platforms. This may partly be seen as a commitment device by the platform to protect inframarginal seller rents.

in the product category that they are interested in; the valuation is drawn from the uniform distribution $U[0, \overline{v}]$; in the duopoly version, we model taste heterogeneity by assuming that each consumer has location x drawn from the uniform distribution on $[0, 1]$. Consumers observe this realization only after they have joined the platform. Under monopoly, consumers buy in their product category of interest if the price they are asked to pay is weakly lower than their valuation. Under duopoly, we restrict attention to sufficiently attractive products such that, in equilibrium, all consumers buy from their product category of interest; that is, there is full coverage.

Consumers have access to a disclosure technology that is provided by the platform. This disclosure technology allows them to communicate some information on their valuation to the relevant sellers. We follow Ali et al. (2023) and define three disclosure regimes (further details are provided in Appendix III):

- No Evidence (NE): consumers cannot disclose any information regarding their location to sellers;
- Simple Evidence (SE): consumers decide whether or not to disclose their willingness to pay v in the monopoly version or location x to seller i, $i \in \{1, 2\}$ in the duopoly version (disclosure can be seller-specific);
- Rich Evidence (RE): consumers decide whether to disclose (partial) information regarding their willingness to pay v in the monopoly version or location x to seller i, $i \in \{1, 2\}$, in the duopoly version (again, disclosure can be seller-specific); this information can be any convex set of values v or locations x such that the true value is contained in this set.

Intra-platform interactions take the following form: first, consumers make disclosure decisions, then sellers set retail prices simultaneously. Prices can be conditioned on the information received from consumers. Both "simple evidence" and "rich evidence" regimes allow consumers to send messages to the seller(s). In most of the exposition, we take a reduced-form approach to quantify the impact of different disclosure regimes on shares of the realized gains from trade. In particular, following Ali et al. (2023), we assume that the disclosure regime $z \in \{NE, SE, RE\}$ is associated with total gains from trade w^z that are split according to shares λ^z and $(1 - \lambda^z)$ between consumers and sellers, gross of the payment sellers make to the platform.

We consider the following timing:

1. The platform chooses a disclosure regime and an ad valorem fee on seller revenues (that is, the percentage of the seller profit that goes to the platform); consumers are not charged.[43]
2. Sellers learn their opportunity cost of joining and decide whether or not to join.
3. Consumers learn their opportunity cost of joining and choose whether or not to join the platform.

[43]Regarding the former, in practice, platforms often ask for a fraction of seller revenues. However, when variable costs are negligible (as is typically the case with digital products) this is indistinguishable from "taxing" profits. Regarding the latter, in many real-world examples of e-commerce platforms, consumers do not pay the platform to be able to join.

4. Consumers learn their product category of interest and their valuations and make their disclosure decision given the disclosure regime.

5. Given the consumers' information disclosure decision, sellers set prices for each identifiable consumer group.

6. Consumers make purchase decisions and payoffs are realized.

We solve for Perfect Bayesian Equilibria, where, for every subgame starting in stage 4, we select the equilibrium that is most favorable for consumers (as derived in Ali et al., 2023). This allows for a clear distinction between the "simple evidence" and the "rich evidence" regimes.

To avoid uninteresting corner solutions when $\beta < 1$, we assume that there are always some sellers and some consumers who find it too costly to join the platform no matter the specification in place, which holds if $\overline{f}_b$ and $\overline{f}_s$ are sufficiently large. Furthermore, in the duopoly model, we will assume that when a seller has opportunity costs such that this seller would find it profitable to join the platform, another seller always joins as well so that every product category is served by a duopoly; this assumption is reminiscent of the one in Jeon et al. (2022).

Discussion. We assume that consumers do not observe whether "their" product category is available prior to joining. This generates a cross-side network effect from sellers to consumers since these consumers will be more likely to join the more product categories are available on the platform and, thus, leads to a "true" platform problem. Under this assumption, it is immaterial whether or not consumers observe how many product categories are available on the platform and the predictions would be the same in an alternative model in which sellers and consumers simultaneously make their participation decision.[44]

If we were to assume that all consumers observe which product categories are available on the platform prior to making their participation decision and consumers know in which one they are interested, only consumers with an interest in one of the available categories would consider joining. In such a world, the platform becomes fully segmented, meaning that the number of consumers for a given product category does not depend on the availability of other product categories on the platform and, thus, no network effects would be at play.

3 Disclosure with monopoly sellers

3.1 The consumer-seller interaction with monopoly sellers

At stage 4, the consumer and seller participation decisions have already been made, and, given the disclosure regime, consumers decide which information if any to disclose and sellers then set prices conditional on the information about consumer valuations that is available to them. Here, we reproduce the findings by Ali et al. (2023).

Under no evidence, consumers cannot share any information about their valuation v. Then, each seller sets the same price p^M for all consumers. Since $v \sim U[0, \overline{v}]$, it is immediate that the seller sets monopoly price $p^M = \frac{\overline{v}}{2}$ and sells to all consumers with willingness to pay $v \geq p^M$.

[44] With simultaneous participation decisions, consumers would need to observe and "interpret" the seller fee charged by the platform, since they would need to infer the equilibrium fraction of available product categories.

The expected surplus for consumers and sellers (gross of the ad valorem fee) under this regime are $u^{NE} = \frac{1}{2}\left(\frac{3}{4}\overline{v} - p^M\right) = \frac{1}{8}\overline{v}$ and $\pi^{NE} = \frac{1}{2}p^M = \frac{1}{4}\overline{v}$ and the total gain from trade is $w^{NE} = \frac{3}{8}\overline{v}$.

Under simple evidence, consumers can decide to disclose for free their exact valuation v to the seller in their product category. If they do, they receive a personalized price offer. As shown in Ali et al. (2023), with monopoly sellers consumers are at best not better off than under no disclosure if the platform were to select simple evidence. In particular, if a consumer discloses v, the monopoly seller extracts the full surplus from the interaction. Thus, consumers such that $v < p^M$ are indifferent between not disclosing and disclosing information and receive personalized price $p^{SE} = v$. In equilibrium, consumers buy at this price. Expected gains from trade for consumers, sellers, and in total under this regime are then $u^{SE} = \frac{1}{8}\overline{v}$, $\pi^{SE} = \frac{1}{2}p^M + \int_0^{\frac{\overline{v}}{2}} v dv = \frac{3}{8}\overline{v}$, and $w^{SE} = \frac{1}{2}\overline{v}$.

Consider now rich evidence. A consumer with willingness to pay v can now send any message $m = [a, b]$ such that $0 \leq a \leq v \leq b \leq \overline{v}$. Since consumers are not restricted to revealing their exact willingness to pay as they would be under simple evidence, there exists an equilibrium disclosure strategy that leads to a strictly higher consumer surplus. The consumer-preferred equilibrium disclosure studied in Ali et al. (2023) that we select here generates a partition that allows consumers to retain additional utility by inducing the monopoly to offer different prices to different segments of the resulting truncated distribution. Formally, Ali et al. (2023) propose an equilibrium in which the interval $[0, \overline{v}]$ is segmented by threshold values $\left(2^{-k}\right)\overline{v}$, $k \in \mathbb{N}_0 \cup \{\infty\}$ so that consumers in each segment pool their messages:

$$m(v) = \begin{cases} m_k = \left(\left(2^{-(k+1)}\right)\overline{v}, \left(2^{-k}\right)\overline{v}\right] & \text{if } v \in \left(\left(2^{-(k+1)}\right)\overline{v}, \left(2^{-k}\right)\overline{v}\right], \\ m_\infty = \{0\} & \text{if } v = 0. \end{cases}$$

Each seller then sets $p_k = \left(2^{-(k+1)}\right)\overline{v}$ to consumers with message m_k. For example, consumers with $v \in (\frac{v}{4}, \frac{v}{2}]$ have an incentive to disclose $m_1 = (\frac{v}{4}, \frac{v}{2}]$ to induce the profit-maximizing price $p_1^{RE} = \frac{v}{4}$.

This equilibrium generates expected gains from trade equal to: $u^{RE} = \frac{v}{2}\sum_{k=1}^{\infty}\left(\frac{1}{4}\right)^k = \frac{1}{6}\overline{v}$, $\pi^{RE} = v\sum_{k=1}^{\infty}\left(\frac{1}{4}\right)^k = \frac{1}{3}\overline{v}$, and $w^{RE} = \frac{1}{2}\overline{v}$.

The result at the interaction stage in the three disclosure regimes is illustrated in Figure 13. In the left panel, under "no evidence", sellers set a uniform price p; in the center, under "simple evidence", they set the same price to consumers who do not disclose ($v \in [1/2, 1]$) and a price equal to v to all other consumers; in the right panel, under "rich evidence", sellers set price p_k for each segment $\left(\left(2^{-(k+1)}\right)\overline{v}, \left(2^{-k}\right)\overline{v}\right]$. The figure also reports the associated consumer surplus for each possible realitation of v.

Expected gains from trade and shares λ are reported in Lemma 5; correspondingly, Table 1 reports gains from trade $w = u + \pi$ and how much of them go to consumers u and how much to sellers (π).

Lemma 5. *(Propositions 1 and 2 in Ali et al., 2023) Suppose that consumers have privately known willingness to pay v extracted from a uniform distribution $U[0, \overline{v}]$. Then, when the sellers are monopolists in their product category, gains from trade w and their share obtained*

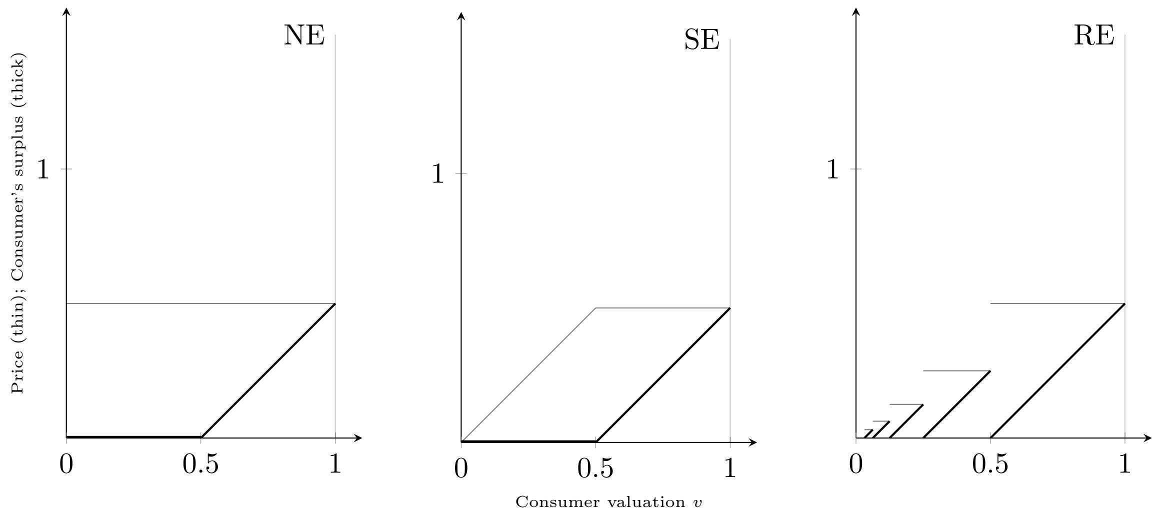

Figure 13: *Equilibrium prices and consumer surplus with monopoly sellers for* $z \in \{NE, SE, RE\}$ *(with* $\overline{v} = 1$*).*

NE	SE	RE
$u^{NE} = \frac{1}{8}\overline{v}$	$u^{SE} = \frac{1}{8}\overline{v}$	$u^{RE} = \frac{1}{6}\overline{v}$
$\pi^{NE} = \frac{1}{4}\overline{v}$	$\pi^{SE} = \frac{3}{8}\overline{v}$	$\pi^{RE} = \frac{1}{3}\overline{v}$
$w^{NE} = \frac{3}{8}\overline{v}$	$w^{SE} = \frac{1}{2}\overline{v}$	$w^{RE} = \frac{1}{2}\overline{v}$

Table 1: *Expected gains from trade for consumers and monopoly sellers under NE, SE, RE.*

by consumers λ *in the consumer-preferred equilibrium are:*

$$w^{NE} = \frac{3}{8}\overline{v}, \quad w^{SE} = \frac{1}{2}\overline{v}, \quad w^{RE} = \frac{1}{2}\overline{v},$$

$$\lambda^{NE} = \frac{1}{3}, \quad \lambda^{SE} = \frac{1}{4}, \quad \lambda^{RE} = \frac{1}{3}.$$

We return to these expressions when evaluating the different disclosure regimes.

3.2 Seller and consumer participation and the profit-maximizing platform fee with monopoly sellers

To shorten notation, we omit the superscript when this does not create ambiguities. Given the disclosure regime and the ad valorem fee, at stages 2 and 3, first sellers and then consumers decide whether to join the platform.

To determine n_b, consider the hesitant consumers' participation decision. Since consumers do not know their product category of interest prior to joining the platform, their participation decision is based on how many product categories are available on the platform. The expected utility of consumers with entry cost f_b at the participation stage is $n_s \lambda w - f_b$. Since f_b is uniformly distributed on $[0, \overline{f}_b]$, we can write the share of hesitant consumers joining the platform as

$$n_b = \frac{n_s \lambda w}{\overline{f}_b}. \tag{18}$$

Notice that that not all of them will make a purchase: since only a fraction of product categories are available, some of the joining consumers end up not purchasing as the product category of interest is not available. The equilibrium volume of trade, then, is $n_s(\beta + (1-\beta)n_b)$.

The share n_s is obtained from the marginal sellers' participation constraint:

$$(1-\alpha_s)(1-\lambda)w(\beta + (1-\beta)n_b(n_s)) \geq f_s.$$

Since f_s is uniformly distributed on $[0, \overline{f}_s]$, the fraction of available product categories n_s solves:

$$n_s = \frac{(1-\alpha_s)(\beta + (1-\beta)n_b(n_s))(1-\lambda)w}{\overline{f}_s}. \tag{19}$$

Substituting for $n_b(n_s)$, we obtain the fraction of active product categories

$$n_s = \frac{(1-\alpha_s)\beta(1-\lambda)w\overline{f}_b}{\overline{f}_s\overline{f}_b - (1-\alpha_s)(1-\beta)(1-\lambda)\lambda w^2}. \tag{20}$$

Hence, using equation (18), the share of hesitant consumers joining the platform is

$$n_b = \frac{(1-\alpha_s)\beta\lambda(1-\lambda)w^2}{\overline{f}_s\overline{f}_b - (1-\alpha_s)(1-\beta)\lambda(1-\lambda)w^2} \tag{21}$$

and the overall share of active consumers is

$$\beta + (1-\beta)n_b = \frac{\beta\overline{f}_s\overline{f}_b}{\overline{f}_s\overline{f}_b - (1-\alpha_s)(1-\beta)(1-\lambda)\lambda w^2}. \tag{22}$$

The expression for active product categories and consumers joining the platform and sellers' profit is positive under the parameter assumption that $\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda)\lambda w^2 > 0$ in all three disclosure regimes (i.e. the denominator of all three expressions is positive for any value of α_s and any admissible z). Equilibrium participation levels are decreasing in α_s: if the platform increases the seller fee, seller profit decreases and, thus, fewer sellers join the platform. This, in turn, reduces the number of consumers joining the platform and, consequently, demand for every product category, which suppresses seller profits even further.

The monopoly platform does two things: it chooses upfront the disclosure regime and the revenue share sellers have to pay to the platform. In this subsection, we solve for the profit-maximizing platform fee for any given disclosure regime. The platform sets a uniform ad valorem fee to all sellers; sellers know their entry cost and choose whether to join the platform or not after observing the entry fee. Note that the decisions of the platform affect the entry decision of consumers and sellers because of the cross-side network effects.

We solve the platform's problem for a given disclosure regime z. Since there is a monopoly seller in each active product category, the platform sets α_s to maximize

$$\Pi^z(\alpha_s) = \alpha_s \left[(1-\lambda)w(\beta + (1-\beta)n_b)\right] n_s, \tag{23}$$

where n_s represents the share of product categories available on the platform, $\beta + (1-\beta)n_b$ the share of consumers who join. Both n_b and n_s depend on α_s.

Plugging in the share of active sellers and consumers from equations (20) and (22) into expression (23), we obtain

$$\Pi^z(\alpha_s) = [(1-\lambda^z)w^z]\frac{\alpha_s(1-\alpha_s)\beta^2(1-\lambda^z)w^z\overline{f}_s\overline{f}_b^2}{(\overline{f}_s\overline{f}_b - (1-\alpha_s)(1-\beta)(1-\lambda^z)\lambda^z(w^z)^2)^2}.$$

The derivative with respect to α_s can be written as

$$\frac{\partial\Pi(r)}{\partial\alpha_s} = \frac{\beta^2\overline{f}_s\overline{f}_b^2(1-\lambda^z)^2(w^z)^2\left[(1-2\alpha)\overline{f}_s\overline{f}_b - (1-\alpha)(1-\beta)(1-\lambda^z)\lambda^z(w^z)^2\right]}{\left(\overline{f}_s\overline{f}_b - (1-\alpha)(1-\beta)(1-\lambda^z)\lambda^z(w^z)^2\right)^3}.$$

The equation is equal to zero if the term in square brackets is zero. The profit-maximizing fee is

$$\alpha_s^* = \frac{\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda)\lambda w^2}{2\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda)\lambda w^2}. \tag{24}$$

We note that α_s^* depends on the disclosure regime through λ as long as $\beta < 1$: if only eager consumers are present in the market there are no cross-group network effect exerted by consumers on sellers, and the platform sets $\alpha_s^* = \frac{1}{2}$ in all disclosure regimes.

3.3 The optimal disclosure regime with monopoly sellers

Given the profit-maximizing fee α_s^*, we characterize the optimal disclosure regime with monopoly sellers from the perspective of the platform, the consumers, and the sellers. We obtain conditions for each in order, we recount the equilibrium trade results as per Ali et al. (2023) for this specification and, finally, we characterize the equilibrium disclosure regime under laissez-faire and under regulation.

Platform-optimal disclosure regime. By plugging in the equilibrium fee from equation (24), the number of sellers joining the platform n_s^* is readily obtained:

$$n_s^* = \frac{\beta(1-\lambda)w\overline{f}_b}{2[\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda)\lambda w^2]}. \tag{25}$$

We rewrite the platform's profit as a function of λ and w through n_s^* and α_s^* by rewriting sellers' profit as a function of n_s^* for a generic disclosure regime and obtain threshold values around which the platform strictly prefers one disclosure regime over the others.

Using equation (23) and noting that $n_s = \frac{(1-\alpha_s)(1-\lambda)w(\beta+(1-\beta)n_b)}{\overline{f}_s}$, we have:

$$\Pi(\alpha_s^*) = \overline{f}_s\frac{\alpha_s^*}{1-\alpha_s^*}(n_s^*)^2 = \frac{\beta^2(1-\lambda)^2w^2\overline{f}_b}{4[\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda)\lambda w^2]}.$$

Consider two disclosure regimes z_1 and z_2. As shown in Ali et al. (2023), when the seller representing a product category is a monopolist, different disclosure regimes can lead to different gains from trade being generated. As such, we want to compare the expected platform profit when w^z is different for different $z \in \{NE, SE, RE\}$. Since $\overline{f}_s$, $\overline{f}_b$, and β are the same in all regimes, the platform prefers z_1 over z_2 if:

$$\frac{[(1-\lambda^{z_1})w^{z_1}]^2}{\overline{f}_s\overline{f}_b-(1-\beta)(1-\lambda^{z_1})\lambda^{z_1}(w^{z_1})^2} > \frac{[(1-\lambda^{z_2})w^{z_2}]^2}{\overline{f}_s\overline{f}_b-(1-\beta)(1-\lambda^{z_2})\lambda^{z_2}(w^{z_2})^2},$$

which can be rewritten as:

$$\begin{aligned}&\overline{f}_s\overline{f}_b[[(1-\lambda^{z_1})w^{z_1}]^2-[(1-\lambda^{z_2})w^{z_2}]^2] > \\ &\qquad (1-\beta)(w^{z_1})^2(w^{z_2})^2(1-\lambda^{z_1})(1-\lambda^{z_2})[(1-\lambda^{z_1})-(1-\lambda^{z_2})].\end{aligned} \tag{26}$$

The expression above captures the trade off of the platform choosing between different disclosure regimes. The left-hand side is larger the greater the gains from trade appropriated by the seller; the right-hand side, instead, reflects the intensity of the network effects, and how valuable it is to attract consumers rather than accommodate sellers. Direct comparison of the two sides requires equilibrium values of w and λ from the trading phase. A few observations, however, can be made directly: first, if two regimes z_1 and z_2 split gains from trade in the same proportion, the one generating the higher gains from trade w would be selected by the platform. Second, if two regimes generate the same gains from trade, the choice of the platforms depends on how costly it is to incentivize the two sides to join. Finally, and intuitively, if consumers generate no indirect network effects (i.e. for $\beta = 1$), the regime that generates higher gains on the seller side maximizes platform profits.

We compare the disclosure regime that maximizes platform profits to the one that maximizes some form of social welfare, while the platform continues to set the fee. In particular, we solve the regulator's problem when the regulator selects the disclosure regime to understand society's incentives for non-price regulation.[45] The market participants are the same as in the model under laissez-faire and have the same choice variables as before, except for the platform that now only sets the fee, while the regulator commits at an earlier stage to a disclosure regime. The regulator's choice of regime then accounts for the profit-maximizing fee derived above. We study the surplus generated for consumers and sellers separately in order to answer the possible misalignment of private and social incentives.

Consumer-optimal disclosure regime. Suppose first that the regulator is interested in maximizing consumer surplus. Consumer surplus CS under disclosure regime z is equal to:

$$\begin{aligned}CS^z &= (1-\beta)n_b^z\left(n_s^z\left(\lambda^z w^z\left(1-\frac{n_s^z}{2}\right)\right)+(1-n_s^z)\left(-\frac{\lambda^z w^z n_s^z}{2}\right)\right)+\beta(\lambda^z w^z n_s^z)\\ &= (1-\beta)\frac{(n_s^z n_b^z\lambda^z w^z)}{2}+\beta(\lambda^z w^z n_s^z).\end{aligned} \tag{27}$$

Where the first component captures the gains from trade that hesitant consumers would gain net of their opportunity cost of joining, and the second captures the gains from trade of the

[45] Recent efforts in the European Union with its GDPR can be seen in this light: the regulation emphasizes individual consent and affects information disclosure, but it does not directly intervene in the platform's pricing decision.

eager consumers. Since it holds that $n_b^z = \frac{\lambda^z w^z n_s^z}{\overline{f}_b}$, the expression can be rewritten as:

$$CS^z = (1-\beta)\frac{(\lambda^z w^z n_s^z)^2}{2\overline{f}_b} + \beta(\lambda^z w^z n_s^z).$$

From the above, it is clear that a regulator interested in maximizing consumer surplus would aim at maximizing $\lambda^z w^z n_s^z$, which in turns maximizes consumer participation. Given the expression for equilibrium seller participation as per Equation 25, and since β and $\overline{f}_b$ are common across regimes, then, a regulator interested in consumer surplus would select disclosure regime z_1 over z_2 as long as:

$$\frac{\lambda^{z_1}(1-\lambda^{z_1})(w^{z_1})^2}{\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_1})\lambda^{z_1}(w^{z_1})^2} > \frac{\lambda^{z_2}(1-\lambda^{z_2})(w^{z_2})^2}{\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_2})\lambda^{z_2}(w^{z_2})^2},$$

which, after rearranging, can be rewritten as:

$$\overline{f}_s\overline{f}_b\left[\lambda^{z_1}(1-\lambda^{z_1})(w^{z_1})^2 - \lambda^{z_2}(1-\lambda^{z_2})(w^{z_2})^2\right] > 0.$$

As it was the case for the platform, if two regimes split the gains from trade in the same proportion, the regulator strictly prefers the one that generates most gains from trade overall. Other considerations are less straightforward to make: the regulator wants to balance consumer and seller participation and compares how dispersed the shares of gains from trade are under different disclosure regimes. Notice, in particular, that if two regimes generate the same gains from trade w, the condition above is equivalent to z_1 being selected over z_2 if:

$$\lambda^{z_1}(1-\lambda^{z_1}) > \lambda^{z_2}(1-\lambda^{z_2}).$$

Seller-optimal disclosure regime. We recall that sellers belong to product categories that differ by their privately known “cost of entry”. Thanks to the uniform distribution of sellers' cost of entry with the lower bound of zero, sellers make on average a profit equal to half the threshold cost of entry. Producer surplus (PS) under disclosure regime z is:

$$PS^z = n_s^z\left(\frac{1}{2}(1-\alpha_s^*)(1-\lambda^z)w^z(\beta + (1-\beta)n_b)\right) = (n_s^z)^2\frac{\overline{f}_s}{2}. \tag{28}$$

To study the seller-optimal disclosure regime, it is sufficient to compare n_s across the various regimes: n_s increases in the expected profit of sellers given market conditions and, therefore, reflects the changes in share of trade, consumer participation and equilibrium platform fee brought forth by different disclosure regimes. As before, we derive the condition such that the regulator maximizing PS prefers regime z_1 to z_2:

$$\frac{(1-\lambda^{z_1})w^{z_1}}{\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_1})\lambda^{z_1}(w^{z_1})^2} > \frac{(1-\lambda^{z_2})w^{z_2}}{\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_2})\lambda^{z_2}(w^{z_2})^2},$$

which can be rewritten as

$$\overline{f}_s\overline{f}_b[(1-\lambda^{z_1})w^{z_1}-(1-\lambda^{z_2})w^{z_2}]$$
$$>(1-\beta)w^{z_1}w^{z_2}(1-\lambda^{z_1})(1-\lambda^{z_2})[\lambda^{z_2}w^{z_2}-\lambda^{z_1}w^{z_1}]. \tag{29}$$

As for the platform, the choice of disclosure regime requires comparing which agents are more costly to encourage joining the platform, which in turns depends on how dispersed entry costs of consumers and sellers are, and how strong the network effects are. It holds for sellers as well as for the platform that if λ^z is the same for two different regimes, the one generating the most gains from trade is preferred; moreover, if consumers exert no cross-group network effects, the regime generating more gains for the sellers is obviously preferred by sellers.

Equilibrium disclosure regime selection In order to precisely pin down the platform's and regulator's optimal disclosure regime we embed the equilibrium results in the consumer-seller interaction as obtained in Section 3.1. As follows from Lemma 5, allowing consumers to share information about their preferences is always strictly better than selecting NE from the perspective of the platform. When disclosure is allowed, sellers are able to generate higher gains from trade by conditioning prices on the messages optimally sent by consumers. Selection between SE and RE, instead, depends on the relative ease with which the platform is able to attract buyers and sellers.

More generally, disclosure is preferable to NE for all agents. Sellers are obviously better off when they can reach consumers that would not purchase anything if they could not disclose their willingness to pay. Consumers do not benefit directly from SE as the additional gains from trade are fully captured by sellers, but they benefit indirectly from the larger number of sellers joining the platform. While platform's and sellers' optimal disclosure regime depends on the parameters, consumers always strictly prefer RE to be selected:

Proposition 5. *Suppose that each product category on the platform is served by a monopoly seller. Given $F=\overline{f}_b\overline{f}_s$, and given the set of disclosure regimes $\{NE, Se, RE\}$, it holds:*

- *if $w<\sqrt{\frac{2}{(1-\beta)}F}$, SE maximizes platform profits and producer surplus,*
- *if $\sqrt{\frac{2}{(1-\beta)}F}<w<\sqrt{\frac{17}{6(1-\beta)}F}$, SE maximizes platform profits and RE maximizes producer surplus,*
- *if $\sqrt{\frac{17}{6(1-\beta)}F}<w$, RE maximizes platform profits and producer surplus,*

RE always maximizes consumer surplus

Proof. A formal proof is relegated to Appendix III. ■

The result in the monopoly case highlights the distortions that a platform can introduce when selecting her preferred disclosure regime. As illustrated in Figure 14, both the platform and the seller generally prefer SE over RE: under SE, sellers can capture a higher share of the same gains from trade $w=\frac{1}{2}\overline{v}$; as w grows, however, attracting consumers becomes relatively more valuable and, for high enough gains from trade, both prefer to switch to RE. The weaker

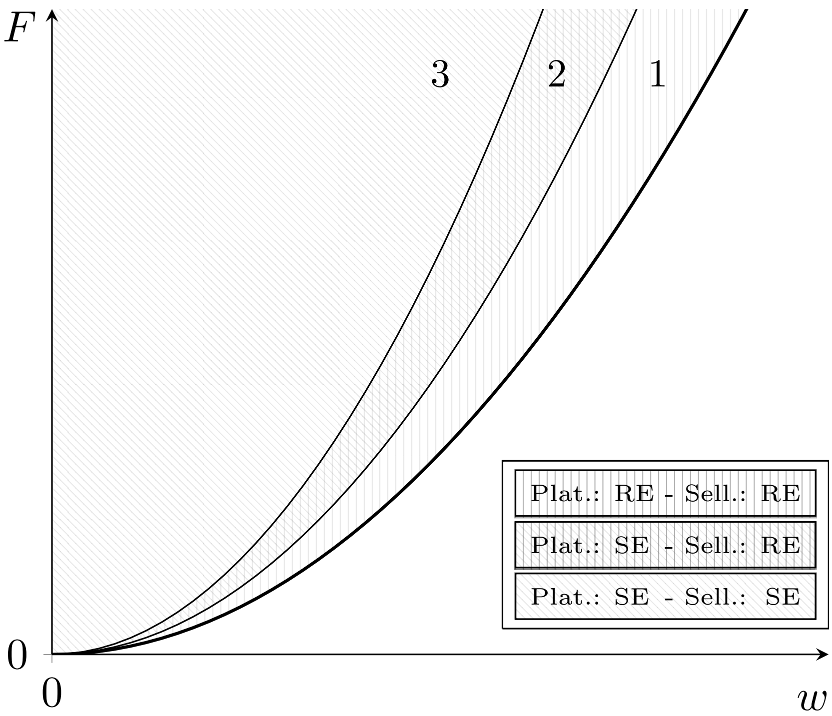

Figure 14: *Preferred disclosure regimes with monopoly sellers.*

the network effects are (that is, the higher β is), the larger w needs to be for the switch to happen. For all values of $\beta < 1$, as w grows, the sellers are interested in switching to the less restrictive regime "earlier" than the platform does.

Consumers, instead, always prefer more disclosure to be available to them: when SE is selected, they obtain, in expectation, as much gains from trade as they would without disclosure in absolute values. They still value having the chance to disclose, since this encourages more sellers to join and, therefore, more consumers have higher expected utility from trade and can, therefore, join the platform. The ability to curate the information they can disclose, however, allows consumers to retain a higher share of w, making it their preferred regime. This holds for eager and hesitant consumers and for all values of β. Sellers and platforms, instead, strictly prefer SE over RE when $\beta = 1$, as in this case there are no cross-group network effects exerted by consumers on sellers: there is no need to allow consumers to strategically disclose information about themselves to encourage them to join the platform in this case.

The analysis so far suggests that the equilibrium disclosure regime favored by consumers would never be selected by a profit-maximizing platform unless consumers exerted strong enough network effects. Note that $\beta = 1$ corresponds to the model proposed by Ali et al. (2023) embedded in a platform environment that, however, does not feature cross-group network effects exerted by consumers on sellers. Then, regulatory intervention with the mandate that consumers must be allowed to disclose freely would be beneficial to consumers (under our equilibrium selection).

	Scenario 1	Scenario 2	Scenario 3
Platform	RE	SE	SE
Sellers	RE	RE	SE
Consumers	RE	RE	RE

Table 2: *Scenarios with monopoly sellers*

	Scenario 1 ($F = 0.071$)			Scenario 2 ($F = 0.095$)			Scenario 3 ($F = 0.13$)		
	NE	SE	RE	NE	SE	RE	NE	SE	RE
α_s	0.497	0.289	0.228	0.498	0.357	0.321	0.498	0.4	0.381
n_s	0.031	0.12	0.14	0.024	0.067	0.07	0.019	0.0427	0.0417
n_b	≈ 0	0.081	0.132	≈ 0	0.044	0.062	≈ 0	0.0267	0.0347
π	≈ 0	0.0021	**0.0023**	≈ 0	**0.0013**	0.0012	≈ 0	**0.0008**	0.0007
PS	≈ 0	0.0027	**0.004**	≈ 0	0.001	**0.0015**	≈ 0	**0.00059**	0.00056
CS	≈ 0	0.002	**0.0037**	≈ 0	0.0027	**0.004**	≈ 0	0.0006	**0.0008**

Table 3: *Numerical results across disclosure regimes with monopoly sellers ($v = 1$, $\beta = 0.1$)*

Corollary 2. *Suppose $\beta = 1$; then:*

- *SE always maximizes platform profits,*
- *SE always maximizes producer surplus,*
- *RE always maximizes consumer surplus.*

Overall, when product categories are served by monopolists, all ecnomic actors prefer to give consumers some ability to disclose their preferences as it generates more trade. If cross-group network effects from consumers to seller are sufficiently strong, and enough gains from trade are generated, all agents align in their interest to allow consumers to curate what kind of information they would like to share. If w is relatively low, instead, both platform and sellers prefer to restrict the ability of buyers to disclose information to a simple evidence regime, while consumers would still prefer to retain more freedom. For intermediate values of w, the platform deviates from both consumers' and sellers' preferred regime in the same direction, restricting disclosure to SE while both consumers and sellers would want RE to be in place. A regulator interested in consumer surplus and, in some cases, producer surplus, then, would optimally intervene to allow consumers to curate the information they share with sellers. Figure 14 and Table 3 illustrate the outcome for a given value of β and for feasible combinations of w and $F = \overline{f}_b \overline{f}_s$, and summarize the possible misalignments in the agents' interests. A detailed discussion of the feasibility constraints and how they interact with β can be found at the end of this section.

The misalignment between sellers' and platform interests may be surprising. To shed some light on it, we report the results of a simple numerical exercise aimed at decomposing this misalignment in Table 3. Given our modeling assumptions, the fee α_s chosen by the platform

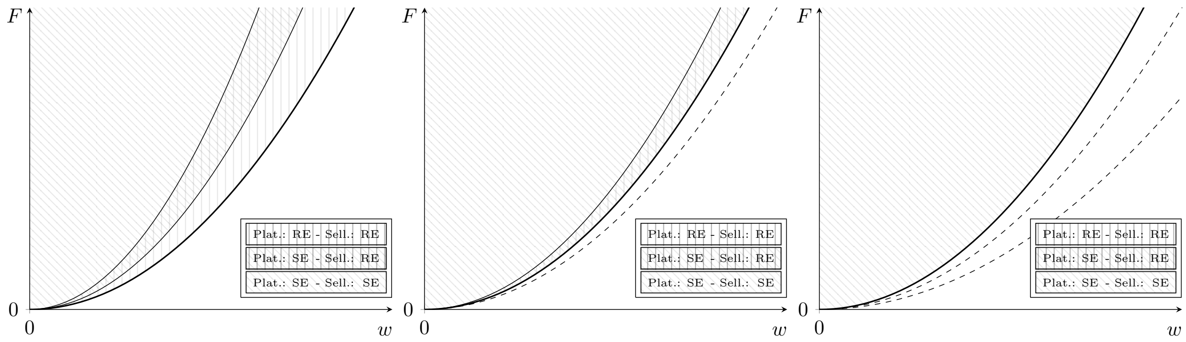

Figure 15: *Equilibrium choices of disclosure regimes for different agents for β small (left), intermediate (center), and large (right) under monopoly sellers.*

is lower the richer the disclosure regime. When a more restrictive disclosure regime is in place, fewer consumers join the platform all else equal. For higher F, allowing consumers to disclose information becomes less effective in inducing them to join the platform at the margin. The tension between platform's and sellers' interests arise for intermediate values of F (for a given value w), that is, when consumers' become more costly to attract. When this is the case, the platform prefers to restrict disclosure and sacrifices consumer and seller participation.

Discussion of admissible parameter constellations. We provided conditions on w and $\overline{v}$ that determine the optimal disclosure regime. However, not all parameter combinations are feasible. Opportunity costs of joining the platform must be such that the model has an interior solution.

$$\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^z)\lambda(w^z)^2 > 0 \quad \forall z \in \{NE, SE, RE\},$$

Thus, it must hold that $\overline{f}_s$, $\overline{f}_b$ is high enough such that some consumers and sellers always find it too costly to join the platform under all disclosure regimes. As lower bounds, we set

$$\overline{f}_s = (1-\lambda^{SE})w^{SE}, \quad \overline{f}_b = \lambda^{RE}w^{RE};$$

that is, the highest absolute gains from trade either side can obtain if the other side joined in full, under their most-preferred disclosure regime, net of the effect of β. We take the above as lower bound values defining the distributions of opportunity costs for sellers and buyers respectively.

The constraint depends on β, which reflects the role of the network effects: for β low enough, there exist parameter constellations such that all agents are aligned in their interest of selecting RE. As β grows, network effects experienced by sellers become weaker and, for β high enough, the platform stops finding RE to be profit-maximizing. For β even higher, the same is true for the sellers. It follows that the platform's interests align with those of consumers and sellers (and, therefore, the social optimum outcome) only if gains from trade are high enough and β sufficiently small. Otherwise, she always chooses a regime different from the one preferred by consumers and, in some cases, also different from the one preferred by sellers. Figure 15 illustrates the possible outcomes for low, intermediate, and high values of β.

4 Disclosure with duopoly sellers

In this section, we allow for seller competition in each product category and consider a particular duopoly model with consumer information disclosure. We recall the setup. We model the duopoly case assuming all product categories to be represented by a $[0, 1]$ interval with two sellers located exogenously at its end points. Consumers' preferences depend on their location x on the line representing their product category of choice. As for the monopoly case, we assume that a fraction β of consumers are eager and that the remaining fraction $(1-\beta)$ are hesitant. We continue to assume that eager consumers have opportunity cost of joining equal at zero, and that hesitant consumers have an opportunity cost that is independently drawn from the uniform distribution $U[0, \overline{f}_b]$. After joining the platform consumers learn their preferred product category and their location on the Hotelling line, which is independently drawn from $U[0, 1]$: a consumer located at x obtains utility $V - tx - p_1$ if buying from seller 1 at price p_1 and $V - t(1-x) - p_2$ if buying from seller 2 at price p_2, where t measures the degree of product differentiation and V is a stand-alone utility of the product, which is assumed to be sufficiently large such that all consumers buy in every admissible disclosure regime. On the seller side, we assume that when a seller has opportunity cost such that it would want to join the platform, a second one always joins as well. Thus, every product category is served by a duopoly. Below in this section, we provide conditions on the parameters such that this assumption is satisfied. Sellers' opportunity costs depend on their product category as in the version with monopoly sellers.

4.1 The consumer-seller interaction with duopoly sellers

From Ali et al. (2023)'s results, reproduced in detail in Appendix III, we note that in our setting with competing sellers, in contrast to our previous analysis with monopoly sellers, there is no deadweight loss in the consumer-seller interaction under any disclosure regime because the market is fully covered and each consumer buys from the seller that is closest in the product space. Therefore, the disclosure regime has no impact on the overall gains from trade for given participation levels. Moreover, the three regimes can be ordered based on the share λ of gains from trade w obtained by consumers; for all values w it holds that

$$\lambda^{RE} > \lambda^{SE} > \lambda^{NE}.$$

Equilibrium shares are reported in the following lemma.

Lemma 6. *(Propositions 6, 7, and 8 in Ali et al., 2023) Suppose consumers have privately known location x extracted from a uniform distribution $U[0, 1]$ and linear transportation costs t, and that the market is fully covered by competing sellers located at the extremes of the unit interval. Then, gains from trade (w) are $w^{NE} = w^{SE} = w^{RE} = V - \frac{1}{4}t$ and the share received by consumers (λ) in the consumer-preferred equilibrium are:*

$$\lambda^{NE} = 1 - \frac{t}{w^{NE}} \qquad \lambda^{SE} = 1 - \frac{3t}{8(w^{SE})} \qquad \lambda^{RE} = 1 - \frac{t}{3(w^{RE})}$$

Proof. See Appendix III. ∎

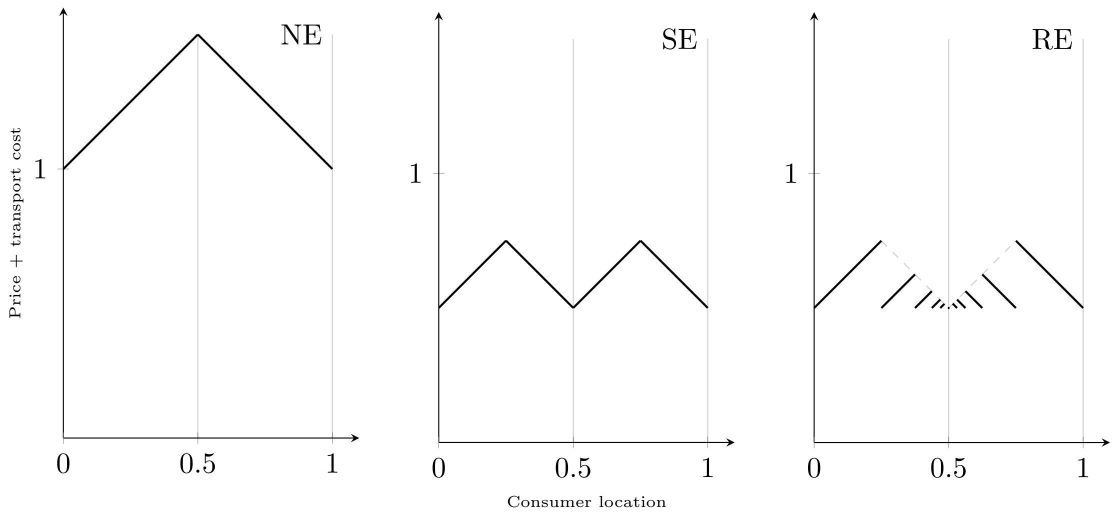

Figure 16: *Price plus transport cost paid by consumer x with duopoly sellers for $z \in \{NE, SE, RE\}$ (with $t = 1$).*

NE	SE	RE
$u^{NE} = V - \frac{5}{4}t$	$u^{SE} = V - \frac{5}{8}t$	$u^{RE} = V - \frac{7}{12}t$
$\pi^{NE} = \frac{1}{2}t$	$\pi^{SE} = \frac{3}{16}t$	$\pi^{RE} = \frac{1}{6}t$
$w^{NE} = V - \frac{1}{4}t$	$w^{SE} = V - \frac{1}{4}t$	$w^{RE} = V - \frac{1}{4}t$

Table 4: *Expected gains from trade for consumers, competing sellers, and in total under NE, SE, RE.*

Two differences from the monopoly setting can be noted. First, as mentioned above, there is no obviously dominated disclosure regime: since all regimes generate the same gains from trade w, NE is now a viable choice for both the platform and a regulator maximizing consumer and/or seller surplus. Additionally, it should be noted that under duopoly the share of gains from trade received by the consumers depends on the level w. This follows from the Hotelling framework: sellers are limited in their price setting for all w. For example, under NE, the standard Hotelling model applies and equilibrium prices are equal to t. As w grows, a higher share of the overall surplus is retained by the consumers since the uniform price does not change.

Price plus transport cost paid by a consumer located at $x \in [0, 1]$ are plotted in Figure 16 for the three disclosure regimes.

Table 4 reports the net expected surplus for any successful interaction between a consumer and a seller in the three disclosure regimes. The three regimes then can be order based on the consumers' and sellers' shares of gains from trade: the more flexibly consumers can disclose their location, the higher their share of the gains from trade. This implies that seller profits are largest under no disclosure, lowest under rich evidence, and at an intermediate level under simple evidence. As for the monopoly case, we use the above values to construct gains from trade w and shares λ and $(1 - \lambda)$ as given in Lemma 6.

4.2 Consumer and seller participation and platform fee setting

The analysis follows the same steps as for the version with monopoly sellers. Participation of consumers and sellers is according to

$$\begin{aligned} n_s &= \frac{(1-\alpha_s)(\beta+(1-\beta)n_b)(1-\lambda^z)w^z}{2\overline{f}_s}, \\ n_b &= \frac{n_s\lambda^z w^z}{\overline{f}_b}. \end{aligned}$$

Thus, equilibrium participation levels are

$$\begin{aligned} n_s &= \frac{(1-\alpha_s)\beta(1-\lambda^z)w^z\overline{f}_b}{2\overline{f}_s\overline{f}_b-(1-\alpha_s)(1-\beta)(1-\lambda^z)\lambda^z(w^z)^2}, \\ n_b &= \frac{(1-\alpha_s)\beta\lambda^z(1-\lambda^z)(w^z)^2}{2\overline{f}_s\overline{f}_b-(1-\alpha_s)(1-\beta)\lambda^z(1-\lambda^z)(w^z)^2}. \end{aligned}$$

The platform's profit function is

$$\Pi^z(\alpha_s) = \left[2\alpha_s(1-\lambda^z)\frac{w}{2}(\beta+(1-\beta)n_b)\right]n_s.$$

The first-order condition of profit maximization can be rewritten as

$$(1-2\alpha)2\overline{f}_s\overline{f}_b-(1-\alpha)(1-\beta)(1-\lambda^z)\lambda^z(w^z)^2=0$$

which leads to the profit-maximizing fee:

$$\alpha_s^* = \frac{2\overline{f}_s\overline{f}_b-(1-\beta)(1-\lambda^z)\lambda^z w^2}{4\overline{f}_s\overline{f}_b-(1-\beta)(1-\lambda^z)\lambda^z w^2}.$$

Platform profit under regime z is

$$\Pi^z(\alpha_s^z) = 2\overline{f}_s\frac{\alpha_s^*}{1-\alpha_s^*}(n_s^*)^2.$$

Expressions for surpluses CS and PS are

$$CS^z = (1-\beta)\frac{(\lambda^z w^z n_s^z)^2}{2\overline{f}_b}+\beta(\lambda^z w^z n_s^z), \qquad PS^z = (n_s^z)^2\frac{\overline{f}_s}{2}.$$

To compare the different disclosure regimes, we make use of the equilibrium outcomes in the consumer-seller interaction as characterized in Lemma 6. Plugging in each value of λ^z in the above expressions allows us to compare expected platform profits, CS and PS across the three disclosure regimes. The main result of the section follows:

Proposition 6. *Suppose that each product category on the platform is served by duopoly sellers. Given $F = 2\overline{f}_s\overline{f}_b$, disclosure regimes from the set $\{NE, SE, RE\}$ perform as follows:*

- *if $w < \frac{11F}{(1-\beta)3t} = w_1^P$, NE maximizes platform profits; if $w_1^P < w < \frac{17F}{(1-\beta)3t} = w_2^P$, SE does; otherwise RE does.*
- *if $F < \frac{1-\beta}{8}t^2 = F_1^S$, RE maximizes producer surplus; if $F_1^S < F < \frac{3(1-\beta)}{8}t^2 = F_2^S$, SE does; otherwise NE does*
- *if $w < \frac{17}{24}t = w_1^C$, RE maximizes consumer surplus; if $w_1^C < w < \frac{11}{8}t = w_2^C$, SE does; otherwise NE does.*

Proof. See Appendix III. ■

The platform's incentives to set different disclosure regimes mirror the ones that we obtained when product categories are served by monopoly sellers. In general, the platform has a strong incentive to select the disclosure regime that generates higher gains from trade on the seller side, since that is the side the platform monetizes. The network effects in place, however, create some incentive to allow for disclosure to take place to attract more consumers and generate more trade. As w grows, since the per trade profit of sellers is fixed to the equilibrium prices, this incentive grows as well. Switching from a less-flexible to a more-flexible disclosure regime implies that, under seller duopoly, a lower fraction of the gains from trade can be extracted by the sellers. However, this leads to higher consumer participation, which may mean more trade and higher expected profits for sellers and the platform. For w high enough, the latter effects dominates. The threshold at which the platform is indifferent between two regimes increases in the dispersion of buyers and sellers, captured by F, which stands for how difficult it is to attract both sides of the trade. The threshold also increases in β: the weaker the network effects (i.e. the larger is β), the lower the incentives to allow for disclosure. It follows that NE is the regime selected by the platform if $\beta = 1$.

A regulator interested in maximizing producer surplus, on the other hand, is primarily interested in the level of dispersion of the two sides of the transaction. The reason lies in the interplay between network effects and per-category competition. Since the absolute value of gains from trade retained by sellers is fixed for all levels of w, sellers' incentives to allow for consumer disclosure come into play when F, the level of dispersion of the two sides, is particularly low. When this happens, network effects are relatively stronger at the margin, and the lower share of gains from trade obtained under a more-flexible disclosure regime is more than compensated by the additional trade it generates. The threshold at which the two effects cancel each other out increases with the strength of the network effects, that is, it decreases in β. As it was the case for the platform, NE is best from the sellers' perspective if $\beta = 1$.

Finally, a regulator interested in maximizing consumer surplus balances seller participation and the share of gains from trade retained by consumers. Unlike the sellers, consumers benefit from higher levels of w since seller competition implies that consumers to receive higher shares of the gains from trade under all disclosuer regimes as w increases. When w is low, consumers benefit relatively more from more-flexible regimes since disclosure allows them to obtain larger shares. As w grows, however, this effects becomes weaker, and it becomes relatively more

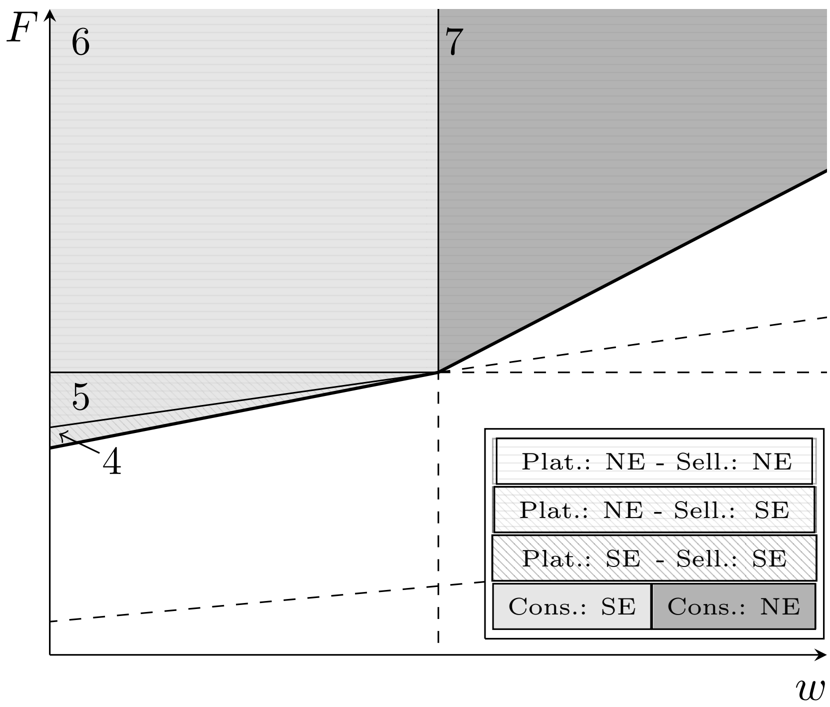

Figure 17: *Equilibrium disclosure regime as selected by a profit maximizing platform, a CS-maximizing regulator, and a PS-maximizing regulator for $t = 1$ when sellers compete in duopoly.*

important to create incentives for sellers to join and restricting disclosure as this allows sellers benefit from, on average, higher prices and thus obtain a higher share of w. For w high enough, the latter effect dominates and consumers benefit from restricting disclosure.

Corollary 3. *Suppose that each product category on the platform is served by duopoly sellers. If $\beta = 1$, that is, consumers do not exert a cross-group network effect on sellers, the following holds:*

- *NE always maximizes platform profits,*
- *NE always maximizes producer surplus,*
- *consumer surplus is maximized by the disclosure regime reported in Proposition 6.*

The results above must be qualified because ofthe distributional assumptions required for consumer and seller participation to lead to an internal solution and for all categories to be fully covered in the trading phase. Taking these constraints into account, only regimes NE and SE are viable under laissez-faire and regulation, as shown in Figure 17, given the parameter restriction of our model (more on those restrictions at the end of the section).

Table 5 reports the possible scenarios regarding the preferred disclosure regime by the platform, consumers, and sellers. When the platform's, the sellers' and the consumers' interests are not full aligned, the following two outcomes are possible: (i) the platform selects no disclosure, while sellers and consumers would prefer the simple evidence regime; (ii) the platform selects no disclosure, which is also in the interest of sellers, whereas consumers would prefer the simple evidence regime. The latter case is intuitive, since both platform and sellers benefit from

	Scenario 4	Scenario 5	Scenario 6	Scenario 7
Platform	SE	NE	NE	NE
Sellers	SE	SE	NE	NE
Consumers	SE	SE	SE	NE

Table 5: *Scenarios with duopoly sellers*

impeding disclosure when network effects are relatively weak. The former, instead, sheds light on the subtler interaction between disclosure regime and the optimal platform fee. When w is particularly low and F is relatively high, the platform has an incentive to select NE and the associated higher equilibrium fee, which leads to higher platform profits but lower seller participation, where sellers would prefer SE. This misalignment emerges for parameter constellations such that CS-maximization would also lead to SE rather than NE: when the platform deviates from the regime maximizing producer surplus, it deviates from the one maximizing CS as well, and in the same direction.

The numerical examples reported in Table 6 shed some light on the factors driving the misalignment in platform's and sellers' selection of disclosure regime. In particular, the platform's trade off (higher fee and less trade versus lower fee and higher trade) depends on the ease with which buyers and sellers can be encouraged to join. Scenario 1 to 3 have the same overall gains from trade being shared, but buyers' and sellers' side become less responsive to the attraction spiral generated by network effects. When F is low, the platform optimally chooses a lower α_s to maximize participation, as the network effects are stronger all else equal. As their strength dwindles, the platform restricts disclosure opportunities and selects the regime associated with the highest fee. Sellers would prefer less restrictions when F is not too high: while they lose in terms of the per transaction share of gains from trade, they gain indirectly because of the lower fee they need to pay, and because network effects are still relatively strong. When F becomes too high, however, this is not the case anymore: retaining a higher share of gains from trade is better even at a higher fee because the attraction spiral does not generate enough additional demand. Finally, in Scenario 4, the platform selects NE even if such a regime is associated with a lower fee: since w is higher than before, consumers value seller participation over disclosure. The platform, then, selects a low α_s and no disclosure to maximize sellers' expected profit, which in turn increases equilibrium number of product categories, consumer participation, and trade. Disclosure, in this case, is detrimental to all agents since it does not generate enough additional value to compensate for the lower seller participation given the high value of F.

The analysis has implications for the direction and intensity of intervention that regulators should aim at in the context of data sharing between platform and sellers: in our model, data sharing should be encouraged, especially when the focus is on consumer surplus. Indeed, in our setting the platform will never select a more permissive disclosure regime than what is in the interest of consumers and sellers; consumers have the strongest interest in being able to disclose.

Second, data sharing should require consumers' consent in the spirit of the "privacy by default" in the EU's General Data Protection Regulation (GDPR). This would also apply to

	Scenario 4 ($w = \frac{21}{16}$, $F = 0.353$)			Scenario 5 ($w = \frac{21}{16}$, $F = 0.36$)			Scenario 6 ($w = \frac{21}{16}$, $F = 0.45$)			Scenario 7 ($w = \frac{23}{16}$, $F = 0.45$)		
	NE	SE	RE	NE	SE	RE	NE	SE	RE	NE	SE	RE
α_s	0.11	0.014	0.078	0.123	0.032	0.093	0.238	0.185	0.219	0.035	0.11	0.16
n_s	0.052	0.358	0.107	0.094	0.148	0.043	0.039	0.02	0.014	0.328	0.037	0.021
n_b	0.054	0.357	0.36	0.031	0.146	0.044	0.011	0.017	0.013	0.13	0.036	0.02
π	0.00009	**0.0007**	0.0005	**0.00047**	0.00028	0.00007	**0.00019**	0.00003	0.00002	**0.0016**	0.00007	0.00004
PS	0.0002	**0.012**	0.0011	0.0008	**0.0021**	0.0001	**0.00016**	0.00004	0.00002	**0.011**	0.0001	0.00005
CS	0.002	**0.063**	0.001	0.0007	**0.0115**	0.0013	0.00019	**0.00035**	0.00023	**0.011**	0.0011	0.0005

Table 6: *Numerical results across disclosure regimes with duopoly sellers ($t = 1$, $\beta = 0.01$)*

regulation aimed at fostering seller surplus over platform profits: as noticed above, for some parameters constellation the platform selects more rigid regimes than the sellers themselves would. Leaving the control over consumer data sharing to the consumer would limit the platform's freedom to extract rent from sellers through higher fees and ultimately encourage more participation on both sides despite the regime selection looking disadvantageous for the sellers at first glance.

Overall, the misalignment between platform, consumers and sellers points at distortions introduced by the platform's regime selection. The platform has a tendency to restrict disclosure beyond the social optimum. The platform may do so at the expense of trade volume, since SE would lead to higher seller participation given the optimal fee selected by the platform. Comparison of the optimal regime selection rules indicates that CS-maximizing regulation would force the platform to enable information disclosure when consumer and platform interests are misaligned. Since in this case there is sometimes also a misalignment of the sellers' and platform interests this even holds for a PS-maximizing regulator. Then, mandating SE would also be the choice of a regulator who maximizes the sum of consumer surplus and producer surplus, and, in some parameter range, for a regulator maximizing total surplus (which includes platform profits).

Discussion of admissible parameter constellations. We must have that opportunity costs of joining the platform must be such that the model has an interior solution. The condition must reflect the fact that sellers benefit more from NE than from SE given consumer participation, and that sellers split the consumer base in half. Therefore, the candidate thresholds are

$$\overline{f}_s = (1 - \lambda^{NE})\frac{w}{2}, \quad \overline{f}_b = \lambda^{RE} w.$$

Furthermore, another condition must also hold: consumption benefit w must be high enough that the resulting Hotelling model features full coverage in equilibrium. For RE and SE this implies $w > \frac{1}{2}t$; for NE, it implies $w > \frac{5}{4}t$. Thus, we must have $w > \frac{5}{4}t$.

Unlike for the monopoly case, the restriction imposed by $\overline{f}_b$ and $\overline{f}_s$ and reflected in the sufficient condition

$$F = 2\overline{f}_s\overline{f}_b \geq \lambda^{RE}(1 - \lambda^{NE})w^2$$

turns out to be too restrictive. Under this condition, only parameter constellations such that both platform and sellers would optimally select NE can be considered, for all values of β. Such

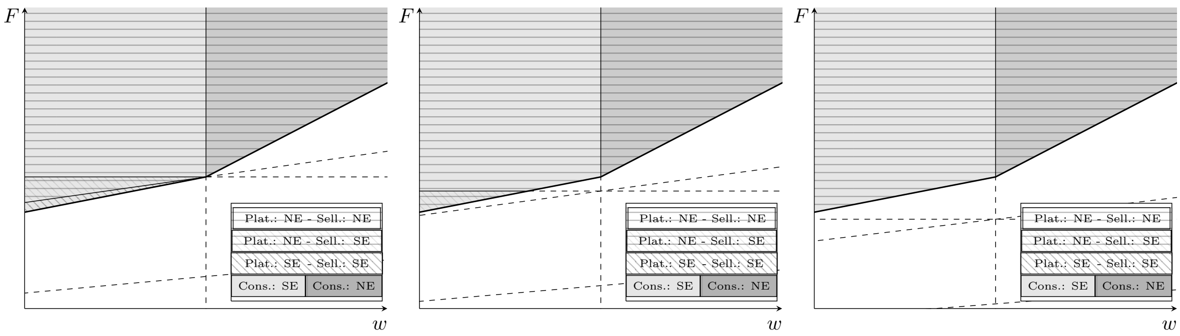

Figure 18: *Preferred disclosure regimes for β small (leftmost), intermediate (center), and large (rightmost) under duopoly sellers.*

a condition, however, is stricter than needed to guarantee interior solutions for equilibrium consumer and seller participation. When we consider the equilibrium values for n_s^* and n_b^*, instead, it can be shown that, for any disclosure regime $z \in \{NE, SE, RE\}$, different, more permissive constraints can be considered. In particular, for any regime z, $n_s, n_b \in (0,1)$ we require that

$$\overline{f}_s^z > (1-\lambda^z)\frac{w}{2}(2-\beta), \quad \overline{f}_b^z > \lambda^z w.$$

These restrictions represent three constraints that must be satisfied for the model to return admissible solutions. Taking the most constraining of the three for all values, $w > \frac{5}{4}t$, is sufficient.

Figure 18 illustrates the constraints and how they affect the possible equilibrium choices of disclosure regime for different values of β. First, we note that the restriction imposed by RE is never relevant, as the ones imposed by SE and RE are always tighter for admissible values of w. Second, for β small, there exist constellations of parameters such that the platform's interests are aligned with any regulator – the regulator may be CS- or PS-maximizing – and SE is selected. As β grows, the network effects become weaker and, for β large enough, the platform selects NE. For some constellations of parameters, this is in conflict with the optimal selection of both a CS- and a PS-maximizing regulator who would select SE instead. For even larger values of β, only values such that a PS-maximizing regulator would also select NE are possible. Still, this can be in conflict with the interests of a CS-maximizing regulator.

The form the constraints take reflect our modeling assumptions: competition taking place on the Hotelling line restricts the viable choices of w; moreover, if opportunity costs of joining for both consumer and seller sides were not to increase linearly, different constraints would apply. We note that the issue is computational and that a more general set-up would generate different statistics without affecting the misalignment highlighted in this paper.

5 Conclusion

In this paper we embed Ali et al. (2023)'s model of voluntary information disclosure by consumers in a platform model. Ali et al. (2023) have shown that consumers benefit from a richer set of messages when disclosing their preferences to sellers. In our setting, the platform enables transactions between consumers and sellers and, on top, is an information designer who decides on the extent to which consumers can voluntarily disclose information about their preferences to sellers. In return for its services, the platform takes a cut from sellers' profits. We consider two specifications. When the platform hosts at most one seller per product category, consumers always benefit from a disclosure regime that allows them to provide more information. In contrast, the sellers and the platform would generally opt for a more restrictive disclosure regime. Notably, for some parameter constellations, the platform restricts disclosure when sellers would prefer not to.

When there is a duopoly in each product category the same type of misalignment can arise. The platform tends to restrict the possibility of information disclosure more than what is in the interest of sellers and even more so than what is in the interest on consumers. Unlike the previous case, for some combinations of parameters all agents can align in their desire to restrict disclosure. While the platform never allows for richer disclosure than what is optimal for consumers and/or sellers, a regulator that maximizes consumer surplus, seller surplus, or any convex combination of consumer surplus, seller surplus, and platform profit, may want to intervene and force the platform to implement a more-flexible disclosure regime.

Our analysis could be generalized in several directions. First, we took a particular model of product market interaction (Hotelling duopoly and monopoly with linear demand). Second, we imposed uniform distributions of the consumers' and sellers' opportunity costs of participation. Third, following Ali et al. (2023), we allowed for three disclosure regimes. Generalizations in all these directions may lead to richer results.

We provided separate analyses as to whether sellers are monopolists or duopolists in their product category. To do so, we made assumptions guaranteeing that the platform will always host either one or two sellers, in any available product category. Under different assumptions, there would be a duopoly only for those product categories for which the opportunity cost to become active is low, and for an intermediate range of product categories sellers would be monopolists. It may be interesting to characterize the outcome under laissez-faire and regulation in this more complex environment.

Appendix

Chapter I Appendix

Simplified Framework: Monopoly Pricing

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

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

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

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

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

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

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

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

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

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 $\frac{3\alpha^2-2\alpha^3}{1+2\alpha-2\alpha^2} \le s < \alpha$, both p^D_L and p^E_L are feasible. Moreover, it holds $p_M = p^E_L < p^D_L$ for the whole segment. Therefore, it is sufficient to compare expected profits under p^E_L and p^D_L. Notice that p^D_L is such that searching again after a bad first realization is not possible. In this segment:

$$\alpha^2(2-p^D_L) + 2\alpha(1-\alpha)(1-p^D_L) - s < 0$$

Therefore, the seller compares:

$$\pi^E_L = (1-(1-\alpha)^4)p^E_L$$
$$\pi^D_I = \alpha^2(1+2(1-\alpha))p^D_L$$

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

$$\pi^D_I > \pi^E_L \iff \frac{4\alpha^2-2\alpha}{3\alpha-1} < s < \alpha$$

p_L^E is selected otherwise.

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

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

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

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

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

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

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

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

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

π_H^E is optimal for:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

which leads to expected profits:

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

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

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

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

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

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

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

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

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

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

which implies:

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

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

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

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

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

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

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

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

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

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

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

General Model: Search Dynamics

Proof of Proposition 3

Step 1: The first compound box Let $X_{1,1}$ be the compound box containing $(1,1)$ and infinitely many nested boxes containing $(1, j > 1)$, $(i > 1, 1)$. Suppose the consumer had already opened the box. Her current payoff is $k = \max\{u_0, A_1 + B_1\}$. To determine how she would act afterwards, consider the value function:

$$V(k) = \max\{k, -s + E[V(\max\{k, A_i + B_1\})], \\ -s + E[V(\max\{k, A_1 + B_j\})]\}.$$

Suppose $V(k) = k$. Then:

$$k > -s + E[V(\max\{k, A_i + B_1\})] = -s + E[\max\{V(A_i + B_1), k\}],$$
$$s > E[\max\{V(A_i + B_1) - k, 0\}] = \int_k^{\hat{y}} (V(A_i + B_1) - k)\, dF(y).$$

$$k > -s + E[V(\max\{k, A_1 + B_j\})] = -s + E[\max\{V(A_1 + B_j), k\}],$$
$$s > E[\max\{V(A_1 + B_j) - k, 0\}] = \int_k^{\hat{y}} (V(A_1 + B_j) - k)\, dF(y).$$

Therefore, there exist values r_A, r_B such that if $k > \max\{r_A, r_B\}$, $V(k) = k$. Suppose that $-s + E[V(\max\{k, A_i + B_1\})] > \max\{k, -s + E[V(\max\{k, A_1 + B_j\})]\}$. Then:

$$V(k) = -s + E[\max\{V(A_1 + B_j), V(k)\}],$$
$$s = E[\max V(A_1 + B_j, V(k))] \quad \rightarrow \quad V(k) = r_A$$

Suppose now that $-s + E[V(\max\{k, A_1 + B_j\})] > \max\{k, -s + E[V(\max\{k, A_i + B_1\})]\}$. Then:

$$V(k) = -s + E[\max\{V(A_i + B_1), V(k)\}],$$
$$s = E[\max V(A_i + B_1, V(k))] \quad \rightarrow \quad V(k) = r_B$$

To compute r_A and r_B, the optimal policy conditional on $V(k) = r_A$ and $V(k) = r_B$ respectively must be defined. Assuming that the consumer inspects products in increasing order of their indices when indifferent, I make the following:

Claim 1. *If $V(\max\{u_0, A_1 + B_1\}) = r_A$, $V(\max\{u_0, A_1 + B_1, A_1 + B_2\}) = \max\{A_1 + B_2, r_A\}$; if $V(\max\{u_0, A_1 + B_1\}) = r_B$, $V(\max\{u_0, A_1 + B_1, A_2 + B_1\}) = \max\{A_2 + B_1, r_B\}$.*

By contradiction, suppose that $V(\max\{u_0, A_1 + B_1\}) = r_A$. Then:

$$V(\max\{u_0, A_1 + B_1, A_1 + B_2\}) = \max\{u_0, A_1 + B_1, A_1 + B_2,$$
$$- s + E[V(\max\{k, A_i + B_1\})], -s + E[V(\max\{k, A_1 + B_j\})]\}.$$

This is immediate: because $V(\max\{u_0, A_1 + B_1\}) = r_A$, it must hold:

$$-s + E[V(\max\{k, A_1 + B_j\})] > \max\{u_0, A_1 + B_1, -s + E[V(\max\{k, A_i + B_1\})]\}.$$

To see why, suppose $\max\{u_0, A_1 + B_1\} = A_1 + B_1$. Then, by the same argument as above, for $V(\max\{u_0, A_1 + B_1\}) = r_A$ it must be that $k < r_A$. If $\max\{u_0, A_1 + B_1\} > A_1 + B_2$, the same condition applies. Otherwise, if $\max\{u_0, A_1 + B_1\} < A_2 + B_2$, then for $V(\max\{u_0, A_1 + B_1, A_2 + B, 2\}) = r_A$ to be true, $A_2 + B_2$ must also be below r_A. Because $V(\max\{u_0, A_1 + B_1\}) = r_A$, it must be that $A_1 + B_1 < r_A$. It follows that $V(\max\{u_0, A_1 + B_1, A_1 + B_2\}) = \max\{A_1 + B_2, r_A\}$.

In words: if it is optimal to inspect $A_1 + B_2$ after opening the compound box, it must also be optimal to inspect $A_1 + B_3$ if the consumer does not want to stop searching. Therefore, the optimal policy conditional on $V(\max\{u_0, A_1 + B_1\}) = r_A$ is a myopic policy in which the current highest realization is compared to the value of inspecting the next product. The consumer is indifferent between stopping at $(1, j)$ and inspecting $(1, j + 1)$, $j \geq 1$, if:

$$A_1 + B_j = - s + A_1 + B_j \int_0^{B_j} dF(y) + \int_{B_j}^{\hat{y}} B_{j+1}\, dF(y),$$
$$s = \int_{B_j}^{\hat{y}} (B_{j+1} - B_j)\, dF(y).$$

Let z be the value of B_j that satisfies the condition above. It follows that $r_A = A_1 + z$. In the same fashion, from $V(\max\{u_0, A_1 + B_1\}) = r_B$ one obtains that $r_B = z + B_1$. It follows that the value function representing the choice of opening the compound box $X_{1,1}$ and searching optimally in it is:

$$V(u_0) = \max\Big\{u_0, \max\left\{u_0, \max\{A_1, B_1\} + \max\{z, \min\{A_1, B_1\}\}\right\}\Big\},$$

because the consumer would always select $A_1 + z$ over $B_1 + z$ if and only if $A_1 > B_1$, and will stop at $(1, 1)$ if $\min\{A_1, B_1\} > z$ She would also take her outside option, u_0, if she opens the box and none of these options had value above it. The consumer is indifferent between opening the compound box and not opening if:

$$\begin{aligned} u_0 &= -s + u_0 \int_0^{u_0} dF(y) + \int_{u_0}^{2\hat{y}} w\, dH(w), \\ s &= \int_{u_0}^{2\hat{y}} (w - u_0)\, dH(w), \end{aligned}$$

where $w = \max\{A_1, B_1\} + \max\{z, \min\{A_1, B_1\}\} \in (z, 2\hat{y})$, and its CDF satisfies:

$$\begin{aligned} H(w) = &\int_0^z F_a \left(\int_0^B F_b\,(w-z)\ dF_a(A) + \int_B^{\hat{y}} F_b\,(w-z)\ dF_a(A)\right) dF_b(B) + \\ &+ \int_z^{\hat{y}} F_a \left(\int_0^z F_b\,(w-z)\ dF_a(A) + \int_z^{\hat{y}} F_b\,(w-A)\ dF_a(A)\right) dF_b(B). \end{aligned}$$

Keeping the standard nomenclature, I refer to the value u_0 that satisfies the above equation as the reservation value of the the compound box.

Result 1. *Let $\underline{W}$ be the reservation value of $X_{1,1}$ and z the reservation value of any $y \in A \cup B$. The optimal policy with only one compound box $X_{1,1}$ is:*

- *Open $X_{1,1}$ if $u_0 < \underline{W}$, otherwise keep u_0,*
- *if $u_0 > \max\{A_1, B_1\} + \max\{z, \min\{A_1, B_1\}\}$, stop and keep the outside option, otherwise:*
  - *if $\max\{z, \min\{A_1, B_1\}\} = \min\{A_1, B_1\}$, stop and keep $A_1 + B_1$,*
  - *if $\max\{z, \min\{A_1, B_1\}\} = z$, inspect $(1, j)$ until $B_j \geq z$ is found if $A_1 > B_1$, and inspect $(i, 1)$ until $A_i \geq z$ is found if $A_1 < B_1$.*

Step 2: Uncorrelated compound boxes Let $\tilde{X}_{i,i}$, $i \geq 1$, be the compound box containing (i, i) and infinitely many compound boxes containing $(i, j > i)$, $(j > i, i)$. We want to show that, in this environment, the optimal search policy follows a myopic optimal policy such that if $u_0 < \underline{W}$, the consumer starts searching and stops after finding a product with *ex post* utility higher than the reservation value of all closed boxes.

Suppose the consumer opened $X_{1,1}$. Let $k = \max\{u_0, A_1 + B_1\}$; consider the value function:

$$\begin{aligned} V(k) = \max\{&k, -s + E[V(\max\{k, A_i + B_1\})], \\ &- s + E[V(\max\{k, A_1 + B_j\})], -s + E[V(\max\{k, A_i + B_j\})]\}. \end{aligned}$$

Compared the value function of the last paragraph, we must now also compare the first three options with the last one. Suppose once again that $V(k) = k$. Then, it holds:

$$\begin{aligned} &k > -s + E[V(\max\{k, A_i + B_j\})] = -s + E[\max\{V(A_i + B_1), k\}], \\ &s > E[\max\{V(A_i + B_1) - k, 0\}]. \end{aligned}$$

which once again implies that there exist a value $R_{2,2}$ such that if $k > \max\{r_A, r_B, R_{2,2}\}$, $V(k) = k$.

We must verify that r_A andr_B are still the value associated with a myopic policy. This is once again immediate: if $r_A > \max\{k, r_B, R_{2,2}\}$ (resp., $r_B > \max\{k, r_A, R_{2,2}\}$), then $V(\max\{u_0, A_1 + B + 1, A_1 + B_2\}) = \max\{A_1 + B_1, r_A\}$ (resp., $V(\max\{u_0, A_1 + B + 1, A_1 + B_2\}) = \max\{A_1 + B_1, r_B\}$). In words: if keeping A_1 or B_1 and inspecting B_2 or A_2 has value higher than searching $(2, 2)$, keeping the same attribute and inspecting A_3 or B_3 must also have a higher value. This implies that once an attribute is optimally kept, it is never abandoned.

To fully characterize the search process, we must compute the optimal policy after opening $X_{2,2}$. To do so, we prove the following:

Claim 2. *If* $V(\max\{u_0, A_1 + B_1\}) = R_{2,2}$*:*

$$\begin{aligned} V(\max\{u_0, A_1 + B_1, A_2 + B_2\}) = \max\{&A_2 + B_2, -s + E[V(\max\{A_2 + B_2, A_i + B_2\})], \\ &- s + E[V(\max\{A_2 + B_2, A_2 + B_j\})], -s + E[V(\max\{A_2 + B_2, A_i + B_j\})]\}, \end{aligned}$$

where once again A_i*,* B_j *are unsampled attributes.*

Since currently all compound boxes are uncorrelated, opening $X_{2,2}$ does not generate any new information about the content of $X_{1,1}$. This will not be the case when we prove the statement in the final step of the proof. For now: since we know that optimally keeping an attribute leads to an myopic optimal policy, the result follows from the same observation than before. In particular:

- If $V(\max\{u_0, A_1 + B_1, A_2 + B_2\}) = -s + E[V(\max\{A_2 + B_2, A_2 + B_j\})] = A_2 + z$, or $V(\max\{u_0, A_1 + B_1, A_2 + B_2\}) = -s + E[V(\max\{A_2 + B_2, A_i + B_2\})] = B_2 + z$ the consumer will open nested boxes myopically forever,
- If $V(\max\{u_0, A_1 + B_1, A_2 + B_2\}) = A_2 + B_2$, the consumer would stop,
- If $V(\max\{u_0, A_1 + B_1, A_2 + B_2\}) = -s + E[V(\max\{A_2 + B_2, A_i + B_j\})]$, the consumer would open the next compound box.

It must hold that $V(\max\{u_0, A_1 + B_1, A_2 + B_2\}) \neq \max\{u_0, A_1 + B_1, A_1 + z, B_1 + z\}$ because $V(\max\{u_0, A_1 + B_1\}) = R_{2,2}$. In particular, to see that $V(\max\{u_0, A_1 + B_1, A_2 + B_2\}) \neq \max\{A_1 + z, B_1 + z\}$ – that is, that it is suboptimal to go back to $X_{1,1}$ after opening $X_{2,2}$, suppose by contradiction that the optimal policy makes the consumer go back to $A_1 + z$ or $B_1 + z$ after opening $X_{2,2}$ with probability $q \in (0, 1]$ and proceed optimally afterwards. We already established that after inspecting $(1, 2)$ or $(2, 1)$ the consumer would optimally keep searching keeping either A_1 or B_1 fixed. If this is the case, since $R_{2,2} > r_A$ because the consumer opened $X_{2,2}$ instead of opening nested boxes, it must hold:

$$\begin{aligned} R_{2,2} =& q r_A + (1 - q)(R_{2,2} - r_A), \\ R_{2,2} q =& 2 q r_A - r_A, \\ R_{2,2} =& 2 r_A - \frac{r_A}{q} > r_A, \\ q >& 1. \end{aligned}$$

which is a contradiction.

It follows that the optimal policy after opening $X_{2,2}$ is to myopically select between $\max\{A_2, B_2\} + \max\{z, \min\{A_2, B_2\}\}$ and $R_{3,3}$. Since this was the same policy the consumer followed at $(1, 1)$, the consumer is once again following a myopic policy. Therefore, $R_{2,2} = \underline{W}$, which proves the claim.

Result 2. *Let* $\underline{W}_{i,i} = \underline{W}$ *be the reservation value of uncorrelated compound boxes* $\tilde{X}_{i,i}$ *and* z *the reservation value of any* $y \in A \cup B$*. The optimal policy with infinitely many* $\tilde{X}_{i,i}$ *is:*

- *Open* $\tilde{X}_{1,1}$ *if* $u_0 < \underline{W}$*, otherwise keep* u_0*,*
- *if* $\max\{A_1, B_1\} + \max\{z, \min\{A_1, B_1\}\} > \underline{W}$*:*
  - *if* $\max\{z, \min\{A_1, B_1\}\} = \min\{A_1, B_1\}$*, stop and keep* $A_1 + B_1$*,*
  - *if* $\max\{z, \min\{A_1, B_1\}\} = z$*, inspect* $(1, j)$ *until* $B_j > z$ *is found if* $A_1 > B_1$ *and inspect* $(i, 1)$ *until* $A_i > z$ *is found if* $A_1 < B_1$*,*
- *if* $\max\{A_1, B_1\} + \max\{z, \min\{A_1, B_1\}\} < \underline{W}$*, open* $\tilde{X}_{2,2}$ *and go back to the second point.*

Step 3: General model We now remove the assumption of compound boxes being uncorrelated. This implies that the consumer can move freely on the grid of products inspecting one or two new attributes as she sees fit. We want to show that the optimal search process still follows a process that can fully characterized with threshold rules.

Suppose for now that the consumer never optimally goes back to a previously discarded attribute. That is, suppose that $k = \max\{u_0, A_1 + B_1, A_1 + B_2\}$. Then:

$$\begin{aligned} V(k) = \max\{k, -s + E[V(\max\{k, A_1 + B_j\})], \\ -s + E[V(\max\{k, A_i + B_2\})], -s + E[V(\max\{k, A_i + B_j\})]\}. \end{aligned}$$

Notice that this value function does not allow the consumer to inspect combinations of discovered attributes. This will be addressed shortly. For now, we want to show that:

Claim 3. *If the consumer cannot backtrack to a combination of discovered attributes,* $V(\max\{u_0, A_1 + B_1\}) = r_A$ *implies:*

$$V(\max\{u_0, A_1 + B_1, A_1 + B_2\}) = \max\{A_1 + B_2, -s + E[V(\max\{A_1 + B_2, A_1 + B_j\})],$$

where B_j *are unsampled* B *attributes.*

In words: we want to show that the if the consumer cannot backtrack to a combination of undiscovered attributes, keeping an attribute and still searching follows an optimal policy that is still myopic.

From the last paragraph, we know that if $V(\max\{u_0, A_1 + B_1\}) = r_A$, it must hold:

$$\begin{aligned} V(\max\{u_0, A_1 + B_1, A_1 + B_2\}) \neq \max\{u_0, A_1 + B_1, \\ -s + E[V(\max\{k, A_i + B_1\})], -s + E[V(\max\{k, A_i + B_j\})]. \end{aligned}$$

We must now prove that $r_A > R_{2,2}$ implies that $V(\max\{u_0, A_1 + B_1, A_1 + B_2\}) \neq -s + E[V(\max\{k, A_i + B_2\})]$, that is, inspecting $(2, 2)$ after $(1, 2)$ cannot be optimal. Suppose by contradiction that the optimal policy was such that, after inspecting $(1, 2)$, the consumer would optimally inspect $(2, 2)$ with probability $q \in (0, 1)$ and then follow the optimal policy from there. Then, it must hold:

$$\begin{aligned} r_A =& qR_{2,2} + (1 - q)(r_A - R_{2,2}), \\ r_A q =& 2qR_{2,2} - R_{2,2}, \\ r_A =& 2R_{2,2} - \frac{R_{2,2}}{q} > R_{2,2}, \\ q >& 1. \end{aligned}$$

Which is clearly a contradiction: if $(1, 2)$ is optimally picked over $(2, 2)$, it can never be optimal to inspect $(2, 2)$ afterwards. This proves the claim.

When compound boxes share products, each combination of known attributes, inspected or not, becomes effectively an outside option. Suppose that the consumer opened $X_{1,1}$, $X_{2,2}$, and $X_{3,3}$. Then, the highest available payoff for the consumer is:

$$\begin{aligned} k = \max\{&u_0, A_1 + B_1, A_2 + B_2, A_3 + B_3, \\ &A_1 + B_2 - s, A_1 + B_3 - s, A_2 + B_3 - s, \\ &A_2 + B_1 - s, A_3 + B_1 - s, A_3 + B_2 - s\}. \end{aligned}$$

Only a subset of these can ever be relevant. Consider $A_2 + B_3 - s$ and $A_1 + B_3 - s$. If the consumer decides to backtrack to one of these two available payoffs after opening $X_{3,3}$ and observing the realization B_3, it is clear that she woulds select the former if $A_2 > A_1$ and the latter otherwise. Importantly, this information is known to the consumer before opening the compound box $X_{3,3}$.

We must now prove that if the consumer decides to open nested boxes, he never backtracks. Suppose the consumer optimally opened in sequence: $X_{1,1}$, $X_{2,2}$, $(2, 3)$. We want to show that the consumer does not backtrack to $(1, 2)$ or $(2, 1)$ nor to the newly discovered $(1, 3)$. For the first two, notice that the sequencing implies:

$$V(\max\{u_0, A_1 + B_1, A_2 + B_2\}) = r_{A_2} > \max\{R_{3,3} = R_{2,2}, A_1 + B_2 - s, A_2 + B_1 - s\},$$

where r_{A_2} is the continuation value of the optimal policy after keeping A_2.

For the last one, notice that $V(\max\{u_0, A_1 + B_1\}) = R_{2,2}$ implies $R_{2,2} > r_{A_1}$, where r_{A_1} is the continuation value of the optimal policy after keeping A_1. But then it holds:

$$r_{A_2} > R_{3,3} = R_{2,2} > r_{A_1} \iff A_2 > A_1 \iff A_2 + B_3 > A_1 + B_3 - s.$$

Since keeping $(2, 3)$ must better than backtracking to known $(1, 3)$, backtracking is never optimal. Therefore, once again keeping an attribute leads to a myopic optimal policy: once an attribute is kept, it is never dropped.

The choice of backtracking then must be pinned down by the highest past realization. Suppose the consumer opens $X_{3,3}$. If she chooses to keep A_3, she chooses between:

- $A_3 + B_3$, readily available,
- $A_3 + z$, opening nested boxes,
- $A_3 + B_1 - s$ or $A_3 + B_2 - s$, backtracking.

Without loss of generality, suppose $B_2 > B_1$. Suppose further that $B_3 < z$. Then, the consumer chooses $\max\{A_3 + z, A_3 + B_1 - s\}$. Notice that:

$$\max\{A_3 + z, A_3 + B_1 - s\} = A_3 + B_1 - s \quad \iff \quad B_1 \geq z + s.$$

Because the realization B_1 is known before opening the box, the consumer is already aware of whether she would backtrack or go forward. Moreover, if $V(k) = A_3 + z$, it is clear that the consumer would never choose to backtrack afterwards.

This confirms that the optimal policy conditional on inspecting a single attribute is myopic.

We can finally prove the main statement. In words, we now use the predictability of the search process when the highest past realization of A attributes, A^H, and B attributes, B^H, are above or below $z + s$ to define these thresholds.

Formally, suppose the consumer needs to decide whether to open $X_{i,i}$. let $k = \max\{u_0, \max_{j<i}\{u_{j,j}\}, A^H + B^H - s\}$ be the current highest sure payoff for the consumer. Define:

$$\begin{aligned} V(k) = \max\{&k, -s + E[V(\max\{k, A_i + B_k\})], -s + E[V(\max\{k, A_k + B_i\}), \\ &- s + E[V(\max\{k, A_k + B_k\})]], \quad \forall k > i. \end{aligned}$$

We already established that there exist a value $\hat{k}$ such that if $k > \hat{k}$, $V(k) = k$. Further, we proved above that if $V(k) = -s + E[V(\max\{k, A_k + B_i\})$ or $V(k) = -s + E[V(\max\{k, A_i + B_k\})$, the optimal policy has the consumer only opening nested boxes keeping the same attribute A_i or B_i fixed. Finally, we established that if A^H and/or B^H are above $z + s$, either one or both $-s + E[V(\max\{k, A_k + B_i\})$ and $-s + E[V(\max\{k, A_i + B_k\})$ will be dominated by backtracking. With these considerations we can define the value $\mathcal{W}$ such that the consumer opens compound box $X_{i,i}$ if and only if $\mathcal{W} > \max\{k, A_{i-1} + B_i, A_i + B_{i-1}\}$.

If $\max\{A^H, B^H\} < z + s$, the myopic (but incorrect) value of inspecting the next compound box is the same is in the last paragraph: $\underline{W}$. Suppose $X_{i,i}$ is opened next and $A_i > z + s$. The next box will have a different reservation value. The correct reservation value must account for the possibility of discovering attributes that change the search from that point onward. To do so, we must first compute the value of inspecting a compound box when $\max\{A^H, B^H\} > z + s$. We first do so assuming, as in the last paragraph, that boxes are uncorrelated but in different configurations depending on the number of attributes above $z + s$ that were found. Then, we combine them appropriately to produce the correct reservation values.

Suppose an attribute A^H was found above $z + s$. The consumer would go back to it rather than opening nested boxes unknown in their A component. Let $w_a(A^H)$ be the expected payoff of opening a compound box locked in this configuration. Let this box be $X_{i,i}$. Suppose $B_i < z$ is found. Then, the consumer must choose between opening nested boxes with score $r_{i,j>i} = A_i + z > A_i + B_i = u_{i,i}$ and backtracking to a box with score $r_{i,j<i} = A^H - s + B_i$. Therefore, the consumer inspects nested boxes if $A_i > A^H - s - (z - B_i)$, and backtrack otherwise. If $B_i > z$, the consumer stops at (i,i) if $A_i > A^H - s$ and backtrack otherwise. The CDF then can be obtained by fixing the value B_i and then integrating for it as it was done for configuration 1:

$$\begin{aligned}
H_a(w_a) = \int_0^z F_a \Bigg(& \int_0^{A^H - s - (z - B)} F_b \left(w_a - (A^H - s)\right) dF_a(A) + \\
& + \int_{A^H - s - (z - B)}^{\hat{y}} F_b \left(w_a - z\right) dF_a(A) \Bigg) dF_b(B) + \\
+ \int_z^{\hat{y}} F_a \Bigg(& \int_0^{A^H - s} F_b \left(w_a - (A^H - s)\right) dF_a(A) + \\
& + \int_{A^H - s}^{\hat{y}} F_b \left(w_a - A\right) dF_a(A) \Bigg) dF_b(B).
\end{aligned}$$

The equivalent formulation for $H_b(w_b)$, relevant if $A^H > z + s > B^H$ is:

$$\begin{aligned}
H_b(w_b) = \int_0^z F_b \Bigg(& \int_0^{B^H - s - (z - A)} F_a \left(w_b - (B^H - s)\right) dF_b(B) + \\
& + \int_{B^H - s - (z - A)}^{\hat{y}} F_a \left(w_b - z\right) dF_b(B) \Bigg) dF_a(A) + \\
+ \int_z^{\hat{y}} F_b \Bigg(& \int_0^{B^H - s} F_a \left(w_b - (B^H - s)\right) dF_b(B) + \\
& + \int_{B^H - s}^{\hat{y}} F_a \left(w_b - B\right) dF_b(B) \Bigg) dF_a(A).
\end{aligned}$$

Suppose now $\min\{A^H, B^H\} > z$: the consumer will not inspect any nested box of which he does not already know the value of. Now, the relevant thresholds determining whether something is kept or not are $A^H - s$ and $B_H - s$: (i,i) is only kept if both $A_i > A^H - s$ and

$B_i > B^H - s$. Otherwise, the highest between $A^H - s + B_i$ and $A_i + B^H - s$ is kept. Therefore:

$$\begin{aligned}
H_{a,b}(w_{a,b}) = \int_0^{B^H - s} F_a \Bigg(& \int_0^{A^H - B^H + B} F_b\left(w_{a,b} - (A^H - s)\right) dF_a(A) + \\
& + \int_{A^H - B^H + B}^{\hat{y}} F_b\left(w_{a,b} - (B^H - s)\right) dF_a(A) \Bigg) dF_b(B) + \\
& + \int_{B^H - s}^{\hat{y}} F_a \Bigg(\int_0^{A^H - s} F_b\left(w_{a,b} - (A^H - s)\right) dF_a(A) + \\
& + \int_{A^H - s}^{\hat{y}} F_b\left(w_{a,b} - A\right) dF_a(A) \Bigg) dF_b(B).
\end{aligned}$$

When boxes are assumed to be locked in any of these configurations, the value function governing search is exactly the same as the one in the last paragraph since all boxes are independent. Then, we construct myopic reservation values:

$$s = \int_{\underline{W}_\kappa}^{2\hat{y}} (w_\kappa - \underline{W}_\kappa) dH_\kappa(w_\kappa).,$$

with $\kappa \in \{\{a\}, \{b\}, \{a, b\}\}$.

Let $\underline{W}^*_{a,b}(A^H, B^H)$ be the expected equivalent of costly opening the next box on the search path:

$$\begin{aligned}
\underline{W}^*_{a,b}(A^H, B^H) = \underline{W}_{a,b}(A^H, B^H) & \int_0^{B^H} \int_0^{A^H} dF_a(A) dF_b(B) + \\
& + \int_0^{B^H} \int_{A^H}^{\hat{y}} \underline{W}_{a,b}(A, B^H)\, dF_a(A) dF_b(B) + \\
& + \int_{B^H}^{\hat{y}} \int_0^{A^H} \underline{W}_{a,b}(A^H, B)\, dF_a(A) dF_b(B) + \\
& + \int_{B^H}^{\hat{y}} \int_{A^H}^{\hat{y}} \underline{W}_{a,b}(A, B)\, dF_a(A) dF_b(B).
\end{aligned}$$

Consider the choice of the consumer. If she chooses to open the next compound box, $X_{i+1,i+1}$, she knows that she will stop only if $w_{i+1,i+1} > \underline{W}_{a,b}(\max\{A_{i+1}, A^H\}, \max\{B_{i+1}, B^H\})$. All future boxes will have this updated value. Suppose the consumer does open the box. The choice of opening the box after it will follow the same logic, with possibly new values of A^H, B^H. Notice that the value of this follow-up search is already incorporated in $\underline{W}^*_{a,b}(A^H, B^H)$, as it accounts for possible upward changes in the value of future boxes. Therefore, $\underline{W}^*_{a,b}(A^H, B^H)$ represents the value of inspecting the next box and following up optimally given the new information acquired with the new box, and fully capture the value of the search process from that point onward.

In the same way, for configuration 2 we write:

$$\begin{aligned}
\underline{W}^*_a(A^H) = & \underline{W}_a(A^H) \int_0^{z+s} \int_0^{A^H} dF_a(A) dF_b(B) + \\
& + \int_0^{z+s} \int_{A^H}^{\hat{y}} \underline{W}_a(A)\, dF_a(A) dF_b(B) + \\
& + \int_{z+s}^{\hat{y}} \int_0^{A^H} \underline{W}_{a,b}(A^H, B)\, dF_a(A) dF_b(B) + \\
& + \int_{z+s}^{\hat{y}} \int_{A^H}^{\hat{y}} \underline{W}_{a,b}(A, B)\, dF_a(A) dF_b(B).
\end{aligned}$$

An equivalent formulation can be found for configuration 3. Finally, for configuration 1, we write:

$$\begin{aligned}\underline{W}^* =&\underline{W}\int_0^{z+s}\int_0^{z+s} dF_a(A)dF_b(B)+\\&+\int_0^{z+s}\int_{z+s}^{\hat{y}} \underline{W}_a(A)\, dF_a(A)dF_b(B)+\\&+\int_{z+s}^{\hat{y}}\int_0^{z+s} \underline{W}_b(B)\, dF_a(A)dF_b(B)+\\&+\int_{z+s}^{\hat{y}}\int_{z+s}^{\hat{y}} \underline{W}_{a,b}(A,B)\, dF_a(A)dF_b(B).\end{aligned}$$

Therefore:

$$\mathcal{W}(A^H,B^H)=\begin{cases}\underline{W}^* & \text{if} \quad \max\{A^H,B^H\}<z+s,\\ \underline{W}^*_a(A^H) & \text{if} \quad A^H>z+s>B^H,\\ \underline{W}^*_B(B^H) & \text{if} \quad B^H>z+s>A^H,\\ \underline{W}^*_{a,b}(A^H,B^H) & \text{if} \quad \min\{A^H,B^H\}>z+s.\end{cases}$$

represents the history dependent value of the search process from that point onward accounting for all possible search paths conditional on them arising during the search process itself.

General Model: Monopoly Pricing

The proof of Proposition 4 comes in two 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 all such price vectors are dominated by the uniform price as per Proposition 4.

Optimal differentiated price vector Suppose the seller wanted to set differential prices for his infinite products. First, it is obvious that at least one product must be prices 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 $(1,2)$ because 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 seller 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 seller 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 seller 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 because the consumer has infinite better alternatives available.[46]

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:

$$\mathcal{W}(\mathbf{p^{unif}}) = \mathcal{W} - p^{unif} = 0, \quad \forall (i,j).$$

Where $\mathcal{W}$ is the initial value of inspecting a closed nested box net of prices. From the discussion above and from the proof of proposition 3, all updating to $\mathcal{W}$ is upward and, therefore, probability of trade in case of the highest uniform price is 1. Therefore, it must be shown that no pricing scheme with differential prices can generate a higher expected profit than $\mathcal{W}$.

To do so, it is sufficient to show that when differential prices are set, the value of inspecting any closed compound box is lower than with uniform prices. Suppose this is the case: the threshold above which a compound box is kept when locked in configuration 1 would then be lower. This cascades into a reduction of the overall value of searching and, therefore, reduces expected profits of the firm.

Notice that, as per the main text, without prices it holds:

$$w_{i,i} = \max\{A_i, B_i\} + \max\{z, \min\{A_i, B_i\}\},$$

when all products are priced uniformly, purchasing any of the products inside the compound box is equivalent. When they are not, instead, the price spread affects when consumers would keep searching instead of stopping at (i,i). Suppose $p_{i,i} = p$ and $p_{i,j\neq i} = p_{j\neq i,i} = p + \delta$ for some $\delta > 0$. Then, a consumer would elect to inspect nested boxes if $z - \delta > \min\{A_i, B_i\}$, because

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

now the prices associated with the nested boxes is higher. notice that this implies that $w_{i,i}$ when differential prices are set is equivalent to $w_{i,i}$ with uniform prices when search costs are higher, or:

$$\max\{A_i, B_i\} + \max\{z - \delta, \min\{A_i, B_i\}\} = \max\{A_i, B_i\} + \max\{z', \min\{A_i, B_i\}\},$$

where z' solves:

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

and because $z' = z - \delta$ and z is decreasing in s, the result follows.

Therefore, differential prices reduces the value of search when there are infinitely many attributes, which in turn means that any pricing vector with differential prices limits the value of search compared to an equivalent one with uniform prices. Therefore, any differential pricing vector that makes the consumer indifferent between searching or not (which makes it equivalent to $p^* = \mathcal{W}$) must generate lower expected profits than its equivalent counterpart.

Optimal uniform prices vector Finally, it must be shown that $p^* = \mathcal{W}$ is indeed optimal. To do so, it sufficient to show that $\mathcal{W}_{a,b}(z+s, z+s)$ is the lowest updated value a compound box can ever have and that $\mathcal{W}_{a,b}(z+s, z+s) \geq \mathcal{W}$.

For the former, recall that it holds:

$$\begin{aligned}
\mathcal{W}_{a,b}(A^H, B^H) = \underline{W}_{a,b}(A^H, B^H) \int_0^{B^H} \int_0^{A^H} dF_a(A) dF_b(B) + \\
+ \int_0^{B^H} \int_{A^H}^{\hat{y}} \underline{W}_{a,b}(A, B^H)\, dF_a(A) dF_b(B) + \int_{B^H}^{\hat{y}} \int_0^{A^H} \underline{W}_{a,b}(A^H, B)\, dF_a(A) dF_b(B) + \\
+ \int_{B^H}^{\hat{y}} \int_{A^H}^{\hat{y}} \underline{W}_{a,b}(A, B)\, dF_a(A) dF_b(B)
\end{aligned}$$

$$\begin{aligned}
\mathcal{W}_a(A^H) = \underline{W}_a(A^H) \int_0^{z+s} \int_0^{A^H} dF_a(A) dF_b(B) + \int_0^{z+s} \int_{A^H}^{\hat{y}} \underline{W}_a(A)\, dF_a(A) dF_b(B) + \\
+ \int_{z+s}^{\hat{y}} \int_0^{A^H} \mathcal{W}_{a,b}(A^H, B)\, dF_a(A) dF_b(B) + \int_{z+s}^{\hat{y}} \int_{A^H}^{\hat{y}} \mathcal{W}_{a,b}(A, B)\, dF_a(A) dF_b(B)
\end{aligned}$$

(which has an equivalent counterpart for $\mathcal{W}_b(B^H)$), and

$$\begin{aligned}
\mathcal{W} = \underline{W} \int_0^{z+s} \int_0^{z+s} dF_a(A) dF_b(B) + \int_0^{z+s} \int_{z+s}^{\hat{y}} \mathcal{W}_a(A)\, dF_a(A) dF_b(B) + \\
+ \int_{z+s}^{\hat{y}} \int_0^{z+s} \mathcal{W}_b(B)\, dF_a(A) dF_b(B) + \int_{z+s}^{\hat{y}} \int_{z+s}^{\hat{y}} \mathcal{W}_{a,b}(A, B)\, dF_a(A) dF_b(B),
\end{aligned}$$

and notice that if $A^H = z + s$ (equivalently, $B^H = z + s$), it holds:

$$H_{a,b} = H_a (= H_b) = H,$$

$$\underline{W}_{a,b}(z+s, z+s) = \underline{W}_a(z+s) (= \underline{W}_b(z+s)) = \underline{W}.$$

and that all are weakly increasing in A^H and/or B^H, with $\underline{W}$ being constant in both and the other being strictly increasing in either or both.

Therefore:

$$\mathcal{W}_{a,b}(z+s,z+s) = \underline{W}_{a,b}(z+s,z+s)\int_0^{z+s}\int_0^{z+s} dF_a(A)dF_b(B)+$$
$$+\int_0^{z+s}\int_{z+s}^{\hat{y}} \underline{W}_a(A)\, dF_a(A)dF_b(B)+$$
$$+\int_{z+s}^{\hat{y}}\int_0^{z+s} \underline{W}_b(B)\, dF_a(A)dF_b(B) + \int_{z+s}^{\hat{y}}\int_{z+s}^{\hat{y}} \underline{W}_{a,b}(A,B)\, dF_a(A)dF_b(B),$$

$$\mathcal{W} = \underline{W}\int_0^{z+s}\int_0^{z+s} dF_a(A)dF_b(B)+$$
$$+\int_0^{z+s}\int_{z+s}^{\hat{y}} \underline{W}_a(A)\, dF_a(A)dF_b(B)+$$
$$+\int_{z+s}^{\hat{y}}\int_0^{z+s} \underline{W}_b(B)\, dF_a(A)dF_b(B) + \int_{z+s}^{\hat{y}}\int_{z+s}^{\hat{y}} \underline{W}_{a,b}(A,B)\, dF_a(A)dF_b(B).$$

Therefore, $\underline{W}_{a,b}(z+s,z+s) = \underline{W}$ proves the result.

First, we must compute CDFs for all four configurations in isolation. For configuration 1, both are provided in the main text:

$$w_{i,i} = \max\{A_i, B_i\} + \max\{z, \min\{A_i, B_i\}\}.$$

$$H(w) = \int_0^z F_a\left(\int_0^B F_b\left(w-z\right)\, dF_a(A) + \int_B^{\hat{y}} F_b\left(w-z\right)\, dF_a(A)\right) dF_b(B)+$$
$$+\int_z^{\hat{y}} F_a\left(\int_0^z F_b\left(w-z\right)\, dF_a(A) + \int_z^{\hat{y}} F_b\left(w-A\right)\, dF_a(A)\right) dF_b(B).$$

Consider now configuration 2: in this configuration, an attribute A^H was found above $z+s$. Therefore, the consumer would go back to it rather than opening nested boxes unknown in their A component. Let $w_a(A^H)$ be the expected payoff of opening a compound box locked in this configuration. W.L.O.G., let this box be $X_{i,i}$. Suppose $B_i < z$ is found. Then, the consumer must choose between opening nested boxes with score $r_{i,j>i} = A_i + z > A_i + B_i = u_{i,i}$ and backtracking to a box with score $r_{i,j<i} = A^H - s + B_i$. Therefore, the consumer inspects nested boxes if $A_i > A^H - s - (z - B_i)$, and backtrack otherwise. If $B_i > z$, the consumer stops at (i,i) if $A_i > A^H - s$ and backtrack otherwise. The CDF then can be obtained by fixing the value B_i and then integrating for it as it was done for configuration 1:

$$H_a(w_a) = \int_0^z F_a\left(\int_0^{A^H-s-(z-B)} F_b\left(w_a - (A^H - s)\right)\, dF_a(A)+\right.$$
$$\left.+\int_{A^H-s-(z-B)}^{\hat{y}} F_b\left(w_a - z\right)\, dF_a(A)\right) dF_b(B)+$$
$$+\int_z^{\hat{y}} F_a\left(\int_0^{A^H-s} F_b\left(w_a - (A^H - s)\right)\, dF_a(A) + \int_{A^H-s}^{\hat{y}} F_b\left(w_a - A\right)\, dF_a(A)\right) dF_b(B).$$

The equivalent formulation for $H_b(w_b)$ is:

$$H_b(w_b) = \int_0^z F_b\left(\int_0^{B^H-s-(z-A)} F_a\left(w_b-(B^H-s)\right)\, dF_b(B)+ \right.$$
$$\left. + \int_{B^H-s-(z-A)}^{\hat{y}} F_a\left(w_b-z\right)\, dF_b(B)\right)\, dF_a(A)+$$
$$+\int_z^{\hat{y}} F_b\left(\int_0^{B^H-s} F_a\left(w_b-(B^H-s)\right)\, dF_b(B) + \int_{B^H-s}^{\hat{y}} F_a\left(w_b-B\right)\, dF_b(B)\right)\, dF_a(A).$$

Consider now configuration 4: in this configuration, both an attribute A^H and an attribute B^H were found above $z+s$. The consumer will not inspect any nested box of which he does not already know the value of. Now, the relevant thresholds determining whether something is kept or not are A^H-s and B_H-s: (i,i) is only kept if both $A_i > A^H-s$ and $B_i > B^H-s$. Otherwise, the highest between A^H-s+B_i and A_i+B^H-s is kept. Therefore:

$$H_{a,b}(w_{a,b}) = \int_0^{B^H-s} F_a\left(\int_0^{A^H-B^H+B} F_b\left(w_{a,b}-(A^H-s)\right)\, dF_a(A)+ \right.$$
$$\left. + \int_{A^H-B^H+B}^{\hat{y}} F_b\left(w_{a,b}-(B^H-s)\right)\, dF_a(A)\right)\, dF_b(B)+$$
$$+\int_{B^H-s}^{\hat{y}} F_a\left(\int_0^{A^H-s} F_b\left(w_{a,b}-(A^H-s)\right)\, dF_a(A) + \int_{A^H-s}^{\hat{y}} F_b\left(w_{a,b}-A\right)\, dF_a(A)\right)\, dF_b(B).$$

Given any of these CDFs, we can define and solve the value function governing search following McCall (1970) and Kohn and Shavell (1974). In particular, we want to find $\underline{W}_.$ that solves:

$$\underline{W}_. = -s + \max\{w_., E[\underline{W}_.]\}.$$

It is immediate to see that locking boxes in one of the above configuration, the optimal strategy was the consumer stop after finding $w_. > E[\underline{W}_.]$ and keep searching otherwise. Therefore, we just need to compute $E[\underline{W}_.]$. To do so, we proceed as follows. First, because locking the boxes makes them independent, the probability of each compound box opened being the one at which the consumer stops is the same. Formally, let $E[\underline{W}_.|N]$ be the expected value of the payoff of the consumer after he accepts the N^th offer he receives, it holds:

$$\begin{aligned} E[\underline{W}_.|N] &= E[w_.|N] - sN; \\ E[\underline{W}_.] &= E[E[w_.|N]] - sE[N] \\ &= E[w_.|w_. > \underline{W}_.] - \frac{s}{Pr[w_. > \underline{W}_.]}. \end{aligned}$$

Therefore, $\underline{W}_.$ satisfies:

$$s = \int_{\underline{W}_.}^{2\hat{y}} (w_. - \underline{W}_.)dH_.(w_.).$$

For all four configurations. These static reservation values can then be combines as per the main text to obtain the actual reservation values $\mathcal{W}_.$

Proof of Proposition 3 We are now in a position to prove Proposition 3. To do so, it must be shown that for all configurations it is optimal to keep searching after finding $w_. < \mathcal{W}_.$ and stop otherwise. I consider configuration 4 first because it is the last possible configuration of all unopened compound boxes, then argue that the result implies the same outcome for the configurations that can turn into it. To ease notation, I drop the suffix.

Notice first that, for any A^H, B^H, it holds $\mathcal{W}(A^H, B^H) > \underline{W}(A^H, B^H)$: indeed, $\mathcal{W}(A^H, B^H)$ includes any update generated by realizations higher than A^H or B^H. Then, opening a compound box leads to one of three outcomes. First, if $w < \underline{W}(A^H, B^H)$, the consumer would keep searching. This follows from the definition of $\underline{W}(A^H, B^H)$ in the locked configuration setting. If $w > \mathcal{W}(A^H, B^H)$, the consumer would clearly stop at that compound box. It must be shown that the consumer should keep searching if $\underline{W}(A^H, B^H) < w < \mathcal{W}(A^H, B^H)$ is found.

Suppose the consumer is about to open compound box $X_{i,i}$ and found $\underline{W}(A^H, B^H) < w < \mathcal{W}(A^H, B^H)$. If this is true, it must be that either $A_i > A^H$, $B_i > B^H$, or both. Therefore, the next box will be defined by "locked" reservation value:

$$\underline{W}(\max\{A_i, A^H\}, \max\{B_i, B^H\}) > \underline{W}(A^H, B^H),$$

and dynamic reservation value

$$\mathcal{W}(\max\{A_i, A^H\}, \max\{B_i, B^H\}) > \mathcal{W}(A^H, B^H).$$

If the consumer were to stop searching, he would do so solving a dynamic programming problem with an unchanging value function, namely one that assumes all future boxes to have value $\underline{W}(\max\{A_i, A^H\}, \max\{B_i, B^H\})$. Instead, the continuation value of the search process should account for the possibility of opening the next box to update the value of all unopened boxes to $\underline{W}(\max\{A_i, A_{i+1}, A^H\}, \max\{B_i, B_{i+1}, B^H\})$. But this is exactly the definition of $\mathcal{W}(\max\{A_i, A^H\}, \max\{B_i, B^H\})$. Therefore, the consumer should only stop after finding $w_{i,i} > \mathcal{W}(\max\{A_i, A^H\}, \max\{B_i, B^H\})$.

The same logic can be applied backwards: configurations 2 and 3 both have a dynamic reservation value higher than their "locked" reservation value. Moreover, the dynamic reservation value includes the probability that opening $X_{i,i}$ changes the configuration of unopened boxes from 2 or 3 to 4 and the value of searching on boxes with that configuration from that point onward. Therefore, the dynamic reservation value correctly captures the continuation value of the search process: $\mathcal{W}_a$ and $\mathcal{W}_b$ are the appropriate reservation values governing search. The same applies to configuration 1 as well.

Once a compound box is selected, optimal search inside of it (no matter the configuration therein) follows standard myopic search logic: a nested box is inspected if its reservation value $r_{i,j}$ or $r_{j,i}$ is higher than the highest realized payoff $u_{i,i}$, $u_{i,j}$, or $u_{j,i}$. Because it holds:

$$r_{i,j} = \begin{cases} A_i + B_j - s & \text{if} \quad j < i \\ A_i + z & \text{if} \quad j > i \end{cases}$$

$$r_{j,i} = \begin{cases} A_j + B_i - s & \text{if} \quad j < i \\ B_i + z & \text{if} \quad j > i \end{cases}$$

if the highest $A_{j<i} > z + s$ (resp. $B_{j<i} > z + s$), the consumer would backtrack to a product known to beat all closed nested boxes both known and unknown in their realization. Otherwise, unknown nested boxes would be opened until a second attribute is found to beat the reservation value of all subsequent ones.

To finally prove Proposition 3, it must be shown that when a compound box is optimally kept, it is never abandoned. Suppose $X_{i,i}$ is kept and, in particular, A_i is kept. If B_i is optimally kept as well, the search process stops. Suppose instead $B_i < z$ so that the consumer inspects nested boxes $(i, j > i)$. Suppose $B_{i+1} > z$ is found. It must be shown that the consumer does not want to open any other box. If some $\hat{A}_{j<i} > z + s$ was found, this follows from the fact that if $\hat{A}_{j<i} + z$ was better than the next closed compound box, it would have been kept. The

consumer optimally keeping A_i implies that $A_i + z > \underline{W}_a > \hat{A}j < z + z$. Therefore, it must only be shown that $A_i + B_{i+1} > z + B_{i+1}$, or that $A_i > z$ is necessary for A_i to be kept.

Suppose by contradiction that was not the case. Then, since $A_i \leq z$ was kept but $B_i < z$ was not, it must hold that $2z \geq \mathcal{W}$. Notice however that it holds:

$$2s = \int_{\tilde{z}}^{2\hat{y}} (2y - \tilde{z}) dF(y) \qquad \Longleftrightarrow \qquad \tilde{z} = 2z.$$

Since $2z$ is the certain equivalent of spending two times the search cost to obtain two times the reward associated with one attribute, $2z < \underline{W} < \mathcal{W}$ because $\underline{W}$ is the certain equivalent of discovering the realization of two i.i.d. attributes for one search cost. This contradicts $2z \geq \mathcal{W}$, which proves the result.

Chapter II Appendix

Proof of Lemma 1 and Proposition 1.

Proof. From the indifferent consumers' locations as derived in expressions (8), we build the demands of the three firms. First, the demand platform is simply composed by the consumers who join the platform, i.e., $D_p^{bln} = x_{pb}^{bln} - x_{ap}^{bln}$. Instead, the expressions of the two music labels' demands are more complex, as they include not only the direct consumption by users who buy directly from them but also the share of their content streamed to the platform's users. Formally:

$$D_a^{bln} = x_{ap}^{bln} + \int_{x_{ap}^{bln}}^{x_{pb}^{bln}} (\lambda^*(x_i) + \varepsilon(x_i))\, dx_i$$

$$D_b^{bln} = 1 - x_{pb}^{bln} + \int_{x_{ap}^{bln}}^{x_{pb}^{bln}} (1 - \lambda^*(x) - \varepsilon(x_i))\, dx_i$$

What matters from the platform perspective is not the individual level of bias that each consumer support, but the total mass of demand that, via the biased recommendation, it is able to shift from one seller to the other. In other words, provided that the total mass of demand does not exceed the aggregate participation constraint of the consumers, we can treat it as a uniform value ε.[47] Accordingly, the demand functions of the two music labels change as follows:

$$D_a^{bln} = x_{ap}^{bln} + \int_{x_{ap}^{bln}}^{x_{pb}^{bln}} \lambda^*(x_i)\, dx_i + \varepsilon$$

$$D_b^{bln} = 1 - x_{pb}^{bln} + \int_{x_{ap}^{bln}}^{x_{pb}^{bln}} (1 - \lambda^*(x))\, dx_i + \varepsilon$$

We use these demands and the consumers' locations in expression (8) to obtain the profit functions of the three firms:

$$\pi_a^{bln} = \frac{pb - pa + t^2}{2t^2} p_a, \qquad \pi_b^{bln} = \frac{pa - pb + t^2}{2t^2} p_b,$$

$$\pi_p^{bln} = \frac{t - \sqrt{p_p - p_a} - \sqrt{p_p - p_b}}{t} p_p - \frac{p_b - p_a - t(2\sqrt{p_p - p_a} - t)}{2t^2} p_a - \frac{p_a - p_b - t(2\sqrt{p_p - p_b} - t)}{2t^2} p_b$$

Simple maximization with respect to the prices yields the following:

$$p_a^{bln} = t + \frac{2\varepsilon t}{3}, \quad p_b^{bln} = t - \frac{2\varepsilon t}{3}, \quad p_p^{bln} = \frac{10t}{9} + \varepsilon^2 t$$

[47] For a thorough discussion of the conditions on the bias, refer to Section 4.

.

Using these prices in the functions of firms' profits and consumers' locations we obtain Lemma 1.

We can now go backward to the first stage of the game, where the platform sets and commits to a specific intensity of bias. From Lemma 1, the problem of the platform is the following:

$$\max_{\varepsilon} \quad \pi_p^{bln} = \frac{t(1-39\varepsilon^2)}{27}$$

It is straightforward to observe that the function has a unique maximum in $\varepsilon = 0$. This proves Proposition 1. ■

Proof of Lemma 2 and Proposition 2

Proof. Proof of Lemma 2 follows a similar logic applied to prove Corollary 1. Assume the platform biases her recommendation system by transferring additional demand ε to the firm with low-quality content (firm a). The new demand functions are

$$D_a^{\varepsilon,hq} = x_{ap}^{\varepsilon,hq} + \int_{x_{ap}^{hq}}^{x_{pb}^{\varepsilon,hq}} \lambda^*(x)\, dx + \varepsilon$$

$$D_b^{\varepsilon,hq} = 1 - x_{pb}^{\varepsilon,hq} + \int_{x_{ap}^{\varepsilon,hq}}^{x_{pb}^{\varepsilon,hq}} (1-\lambda^*(x))\, dx - \varepsilon$$

$$D_p^{\varepsilon,hq} = x_{pb}^{\varepsilon,hq} - x_{ap}^{\varepsilon,hq}$$

We use the location of indifferent consumers derived in (9) and apply the same backward induction logic used in the case with no bias. From the system of the first-order condition, the profit-maximizing prices are as derived in equation (15). We use those prices in the payoff functions in equations (12)-(13)-(14) and in (9) to derive Lemma 2.

As mentioned in the main text, we treat ε as a general mass of demand that the platform can shift from one seller to the other. One should notice that, individually, each consumer is not willing to accept a mix that contains too much content a — i.e., $\varepsilon(x_i) \leq \bar{\varepsilon}(x_i)$. We account for that assuming that the platform can re-distribute the bias towards consumers according to their participation constraints. In other words, the condition we impose is that the total demand shifted by the platform cannot exceed the aggregate participation constraint of the consumers, as stated in Condition 1. Moreover, we know the bias must also satisfied the participation of the consumers as derived in expression (16).

Taking all these conditions into consideration, we can now proceed backward to the first stage of the game. As before, from Lemma 2, the problem of the platform is the following:

$$\max_{\varepsilon}^{p} \quad \pi_p^{hq} = \frac{t + 3\varepsilon^p(7v_b - 13\varepsilon^p)}{27} - \frac{v_b^2}{36t}$$
$$s.t. \qquad \varepsilon^p < \min\{\varepsilon^c, \varepsilon^s\}$$

Standard maximization yields the unconstrained profit-maximizing level of bias: $\varepsilon^p = \frac{7v_b}{26t}$. We use this value and prove Proposition 2. ■

Proof of Lemma 3 and Proposition 3

Proof. The proof of Lemma 3 is identical to the ones provided for Lemma 2 up until stage 0, with two adjustments: quality differential, $v_b - v_a = v_b$, is now equal to unconstrained $v_b - v_a$, and sellers' profit is reduced by $I(v_j) = \phi_j\, v_j^2$.

From the proof of Lemma 2, we know that adjusted unconstrained bias can be written as:

$$\varepsilon^p = \frac{7(v_b - v_a)}{26t},$$

which is positive (resp. negative) if $v_b > v_a$ (resp. if $v_b < v_a$): the bias favors the lower quality product. Plugging in this expression in equations 17, one obtains:

$$\pi_a^{eq} = \frac{2(v_a - v_b)^2}{169t} + \frac{t}{2} + \frac{1}{13}((2 - 13v_a)v_a - 2v_b)$$

$$\pi_b^{eq} = \frac{(13t + 2(v_b - v_a))^2}{338t} - \phi_b v_b^2$$

The system of equations that solves these two expressions together leads to the first set of equations in Lemma 3. The second set of equations comes from F.O.C. of the system of the following equations:

$$\pi_a^{eq,out} = \alpha\frac{(3t - (v_b - v_a))^2}{18t} - v_a^2$$

$$\pi_b^{eq,out} = \alpha\frac{(3t + (v_b - v_a))^2}{18t} - (\phi_b v_b)^2$$

Proposition 3 presents two results: aggregate investment in quality is higher off the platform than on the platform for α high enough, and for all values of ϕ_b there exist values of α such that b invests more off the platform than on it, and a invests less off the platform than on it. We prove the two results one at a time.

For the former: aggregate investments on and off the platform are simply:

$$v_{aggr}^{eq} = v_a^{eq} + v_b^{eq} = \frac{169t\,(1 + \phi_b) - 8}{13((169t - 2)\phi_b - 2)}$$

$$v_{aggr}^{eq,out} = v_a^{eq,out} + v_b^{eq,out} = \frac{\alpha\,(9t\,(1 + \phi_b) - 2\alpha)}{3\,((18t - \alpha)\phi_b - \alpha)}$$

The latter equation is increasing in α: the numerator is increasing by virtue of the fact that $\alpha \in (0, 1)$; the denominator shrinks in α. We can then find the value $\underline{\alpha}$ that equates v_{aggr}^{eq} and $v_{aggr}^{eq,out}$:

$$\underline{\alpha} = -\frac{(1 + \phi_b)\left(-19773t^2\phi_b - 273t(1 + \phi_b) + 24\right)}{52((169t - 2)\phi_b - 2)} + \\ + \frac{\sqrt{9(1 + \phi_b)^2(8 - 13t((507t + 7)\phi_b + 7))^2 - 5616t\phi_b((169t - 2)\phi_b - 2)(169t(1 + \phi_b) - 8)}}{52((169t - 2)\phi_b - 2)}$$

Finally, it can be shown that $\underline{\alpha} \in (0, 1)$ for ϕ_b is high enough, and strictly increasing in ϕ_b, which proves the first result.

To prove the second statement we follow the same intuition: first, we find $\underline{\alpha}_a$ and $\underline{\alpha}_b$ at which seller a and b respectively would have invested the same on and off the platform:

$$\underline{\alpha}_a = \frac{39t(507t + 7)\phi_b^2 + 3(91t - 4)\phi_b - 3}{26((169t - 2)\phi_b - 2)} + \\ + \frac{\sqrt{(\phi_b(13t((507t + 7)\phi_b + 7) - 4) - 4)^2 - 312t\phi_b(169t\phi_b - 4)((169t - 2)\phi_b - 2)} - 12}{26((169t - 2)\phi_b - 2)}$$

$$\underline{\alpha}_b = \frac{3\left(6591t^2\phi_b + 91t(1 + \phi_b)\right)}{26((169t - 2)\phi_b - 2)} + \\ - \frac{3\left(\sqrt{(4(1 + \phi_b) - 13t((507t + 7)\phi_b + 7))^2 - 312t(169t - 4)\phi_b((169t - 2)\phi_b - 2)} + 4(1 + \phi_b)\right)}{26((169t - 2)\phi_b - 2)}$$

It can be showed that $\lim_{\phi_b \to 1} \underline{\alpha}_a = \lim_{\phi_b \to 1} \underline{\alpha}_b$, and in particular:

$$\lim_{\phi_b \to 1} \underline{\alpha}_a = \lim_{\phi_b \to 1} \underline{\alpha}_b = \frac{3}{26}\left(\frac{6591t^2 - 494t + 8}{4 - 169t} + 39t + 2\right) = \frac{6}{13}$$

Finally, we show that $\underline{\alpha}_a$ (resp. $\underline{\alpha}_b$) is decreasing (resp. increasing) in ϕ_b, which proves the result. To simplify the proof, assume $t = 1$ without loss of generality. One can see that in this case, $\underline{\alpha}_a$ exists provided that $\phi_b \geq \tilde{\tilde{\phi}} \approx 0.1287$. Moreover, the two thresholds become:

$$\underline{\alpha}_a\Big|_{t=1} = \frac{3\left(\phi_b(6682\phi_b + 87) - 4 - \sqrt{\phi_b(\phi_b(52\phi_b(858637\phi_b - 146979) + 267985) - 3192) + 16}\right)}{26(2 - 167\phi_b)}$$

$$\underline{\alpha}_b\Big|_{t=1} = \frac{3\left(\sqrt{(4(1+\phi_b) - 13(514\phi_b + 7))^2 - 51480\phi_b(167\phi_b - 2)} - 6591\phi_b - 87(1+\phi_b)\right)}{26(2 - 167\phi_b)}$$

The expressions of the first derivatives w.r.t. ϕ are cumbersome, but with some algebra, it is possible to show that they write:

$$\frac{\partial \underline{\alpha}_a\big|_{t=1}}{\partial \phi_b} \equiv \frac{495(27602\phi_b + 393)}{2(2 - 167\phi_b)^2\sqrt{3999836\phi_b^2 + 140548\phi_b + 841}} - \frac{6435}{2(2 - 167\phi_b)^2} > 0$$

$$\frac{\partial \underline{\alpha}_b\big|_{t=1}}{\partial \phi_b} \equiv \frac{771\phi_b(167\phi_b - 4) + 57}{(2 - 167\phi_b)^2} + \frac{3\phi_b(\phi_b((31414589 - 286784758\phi_b)\phi_b - 881874) + 10363) - 60}{(2 - 167\phi_b)^2\sqrt{\phi_b(\phi_b(52\phi_b(858637\phi_b - 146979) + 267985) - 3192) + 16}} < 0$$

the inequalities hold $\forall\, \phi_b \geq \tilde{\tilde{\phi}}$. ■

Extension: On-platform search

Proof. The level of search costs that makes the platform's users indifferent between accepting the bias or paying in order to get the efficient mix is obtained by solving the following:

$$v + (1 - \lambda(x_i) - \varepsilon(x_i))v_b - p_p - t(x_i - (1 - \lambda(x_i) - \varepsilon(x_i)))^2 = v + (1 - \lambda(x_i))v_b - p_p - k - t(x_i - (1 - \lambda(x_i))^2$$

$$k < \varepsilon(x_i)(v_b - t(2(1 - \lambda(x_i) - x) - \varepsilon(x_i))$$

Using the efficient mix bundle $\lambda^{hq}(x_i) = 1 - x_i - \frac{v_b}{2t}$ into the above threshold, we obtain:

$$k < \bar{k} \equiv t(\varepsilon^2(x_i))$$

■

Extension: Different timing - Lemma 4

Proof. Define the demand of the sellers and the platform as D_a, D_b, and D_p, respectively. Consider a situation in which $v_b > v_a = 0$, so that, in equilibrium, absent any bias, $D_a < D_b$ and $p_a < p_b$. Because the platform pays p_a and p_b to the sellers in royalties, it has an incentive to increase the share of content a (the cheapest) in the mix offered to consumers. Also, define ε^c as the total demand on the platform that can be steered toward the cheapest inferior good a without altering D_p.

The bias enters the profit functions of the sellers by altering their demand function. In fact, $D'_{a,\varepsilon} > 0$ and $D'_{b,\varepsilon} < 0$. The two sellers anticipate the bias and modify their prices accordingly. Seller a, who is benefiting from the demand shock, increases the price to $p_a(\varepsilon^c) > p_a$, whereas seller b lower her price to $p_b(\varepsilon^c) < p_b$. This is so because the bias enters the demand function inelastically - i.e., as long as $p_a(\varepsilon^c) < p_b(\varepsilon^c)$, the entire mass ε^c shifts toward good a.

Intuitively, the two scenarios described in the Lemma emerge. First, the demand shift is not enough to change the ranking of the prices. In this case, the resulting equilibrium is such that

$$p_a^* \equiv p_a(\varepsilon^c) < p_b(\varepsilon^c) \equiv p_b^*$$

such that seller a and the platform are better off. In contrast, seller b and consumers are worse off (increasing the recommendation bias lowers consumer surplus).

Second, the demand shift is so significant that the ranking of prices changes. In such a case, an equilibrium in pure strategy no longer exists. In fact, in anticipation of ε^c, sellers change their prices but to such an extent that

$$p_a^* \equiv p_a(\varepsilon^c) \geq p_b(\varepsilon^c) \equiv p_b^*$$

Observing these prices, the platform implements a recommendation bias that goes in the opposite direction of the one anticipated by the seller, promoting content b - which is now the cheapest. Clearly, this cannot be an equilibrium. ■

Extension: Asymmetric costs - Proposition 4

Proof. We derive the relevant functions as we did for the main model. Since by assumption it holds $V_a = V_b = v$, we use the indifferent consumer as derived in expression (8). The demands of sellers are likewise the ones derived in the baseline with homogeneous goods. Demand of the platform is given by $D_p^{dc} = x_{pb}^{dc} - x_{ap}^{dc}$, where the apex ^{dc} stands for "different costs".

Since by assumption it holds $c_a > c_b = 0$, profit function of a must account for the marginal cost and in particular:

$$\pi_a^{dc} = (p_a - c_a)\left(x_{ap}^{dc} + \int_{x_{ap}^{dc}}^{x_{pb}^{dc}} \lambda^*(x)\, dx - \varepsilon\right)$$

Profit functions of b and p are unchanged:

$$\pi_b^{dc} = p_b\left(1 - x_{pb}^{dc} + \int_{x_{ap}^{dc}}^{x_{pb}^{dc}} (1 - \lambda^*(x))\, dx + \varepsilon\right)$$

$$\pi_p^{dc} = p_p\left(x_{pb}^{dc} - x_{ap}^{dc}\right) - p_a\left(\int_{x_{ap}^{dc}}^{x_{pb}^{dc}} \lambda^*(x)\, dx - \varepsilon\right) - p_b\left(\int_{x_{ap}^{dc}}^{x_{pb}^{dc}} (1 - \lambda^*(x))\, dx + \varepsilon\right)$$

We notice that a will see the platform bias consumption away from him since, by standard Hotelling logic, $p_a > p_b$ whenever $c_a > c_b$.

After substituting $\lambda^*(x) = (1 - x)$, standard F.O.C. arguments lead to equilibrium prices:

$$p_a^{dc} = \frac{2c_a}{3} - \frac{2t\varepsilon}{3} + t, \qquad p_b^{dc} = \frac{1}{3}(c_a + t(2\varepsilon + 3)),$$

$$p_p^{dc} = \frac{c_a^2}{16t} - \frac{1}{2}c_a(1 - \varepsilon) + t\left(\varepsilon^2 + \frac{10}{9}\right),$$

and profits:

$$\pi_a = \frac{(c_a + t(2\varepsilon - 3))^2}{18t}, \qquad \pi_b = \frac{(c_a + t(2\varepsilon + 3))^2}{18t}, \qquad \pi_p = \frac{t}{27} - \frac{(c_a - 52t\varepsilon)(c_a - 4t\varepsilon)}{144t};$$

The latter equation immediately leads to the platform profit maximizing bias by standard F.O.C. argument:

$$\varepsilon^p = \frac{7c_a}{52t}$$

To make consumers join, it must hold the following:

$$\varepsilon^p < \varepsilon^c = \int_{x_{ap}^{dc}}^{x_{pb}^{dc}} \bar{\varepsilon}(x)dx$$

where $\bar{\varepsilon}(x)$ is defined as the larger (absolute) bias consumer x is willing to accept before choosing to leave the platform.

Finally, since equilibrium profits in the subgame in which sellers choose not to join the platform are:

$$\pi_a^{dc,out} = \frac{\alpha(c_a - 3t)^2}{18t}, \qquad \pi_b^{dc,out} = \frac{\alpha(c_a + 3t)^2}{18t},$$

ε^p is constrained by ε^s satisfying:

$$\varepsilon^s = \frac{3t - c_a - \sqrt{\alpha(c_a - 3t)^2}}{2t}$$

The result as stated in Proposition 4 follows immediately.

■

Chapter III Appendix

Monopoly sellers

This appendix provides the proof for Proposition 5.

Profit-maximizing platform First, by plugging in the profit-maximizing fee in the expressions for n_s:

$$\begin{aligned} n_s =& \frac{(1-\alpha_s)\beta(1-\lambda^z)w^z\overline{f}_b}{\overline{f}_s\overline{f}_b - (1-\alpha_s)(1-\beta)(1-\lambda^z)\lambda^z(w^z)^2} \\ =& \frac{\frac{\overline{f}_s\overline{f}_b}{2\overline{f}_s\overline{f}_b-(1-\beta)(1-\lambda^z)\lambda^z(w^z)^2}\beta(1-\lambda^z)w^z\overline{f}_b}{\overline{f}_s\overline{f}_b - \frac{\overline{f}_s\overline{f}_b}{2\overline{f}_s\overline{f}_b-(1-\beta)(1-\lambda^z)\lambda^z(w^z)^2}(1-\beta)(1-\lambda^z)\lambda^z(w^z)^2} \\ =& \frac{\beta(1-\lambda^z)w^z\overline{f}_b}{2[\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^z)\lambda^z(w^z)^2]} \end{aligned}$$

and, since it holds:

$$n_b = \frac{n_s\lambda^z w^z}{\overline{f}_b}.$$

we obtain platform profits as

$$\Pi^z = \overline{f}_s \frac{\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^z)\lambda^z(w^z)^2}{2\overline{f}_s\overline{f}_b} \left(\frac{\beta(1-\lambda^z)w^z\overline{f}_b}{2[\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^z)\lambda^z(w^z)^2]} \right)^2,$$

which simplifies to:

$$\Pi^z = \frac{\beta^2(1-\lambda^z)^2(w^z)^2\overline{f}_b}{4[\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^z)\lambda^z(w^z)^2]}.$$

We must have that $\pi^{z1} > \pi^{z1}$ if and only if

$$\begin{aligned} &(1-\lambda^{z_1})^2(w^{z_1})^2(\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_2})\lambda^{z_2}(w^{z2})^2) > \\ &\qquad (1-\lambda^{z_2})^2(w^{z_2})^2(\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_1})\lambda^{z_1}(w^{z_1})^2), \end{aligned}$$

which, after rearranging, is equivalent to

$$\overline{f}_s\overline{f}_b[[(1-\lambda^{z_1})w^{z_1}]^2 - [(1-\lambda^{z_2})w^{z_2}]^2] >$$
$$(1-\beta)(w^{z_1})^2(w^{z_2})^2(1-\lambda^{z_1})(1-\lambda^{z_2})[(1-\lambda^{z_1})-(1-\lambda^{z_2})].$$

Consider λ^z, w^z for all regimes (see Lemma 5). Since $\lambda^{RE} = \lambda^{NE}$ and $w^{RE} > w^{NE}$, RE is preferred to NE as long as:

$$\frac{4}{9}\overline{f}_s\overline{f}_b[(w^{RE})^2 - (w^{NE})^2] > 0,$$

which is always satisfied. Since $\lambda^{RE} > \lambda^{SE}$ and $w^{RE} = w^{SE} = w = \frac{1}{2}\overline{v}$, instead, RE is preferred to SE as long as

$$\overline{f}_s\overline{f}_b[[(1-\lambda^{RE})]^2 - [(1-\lambda^{SE})]^2] >$$
$$(1-\beta)w^2(1-\lambda^{RE})(1-\lambda^{SE})[(1-\lambda^{RE})-(1-\lambda^{SE})]$$

Plugging in the values for λ^{RE} and λ^{SE}, this is equivalent to

$$\frac{17}{6}\overline{f}_s\overline{f}_b < w^2(1-\beta),$$

which is the expression reported in Proposition 5.

Consumer surplus Consumer surplus can be obtained by combining expected utility of joining the platform for both the group of β consumers and the group of $1-\beta$ consumers. For the latter, this is straightforward: since the β consumers only join the platform if the product category they are interested in is present on the platform and have zero opportunity cost of joining by assumption, their expected utility under disclosure regime z is

$$CS^z_\beta = \beta(\lambda^z w^z n^z_s)$$

For the former, we must consider that all $1-\beta$ might not find their product category of interest, and that their opportunity cost of joining is sunk even if they do not find anything to purchase. Given our distributional assumptions, consumers who join have average opportunity cost equal to $\frac{\lambda^z w^z n^z_s}{2}$. Therefore:

$$CS^z_{1-\beta} = (1-\beta)n^z_b\left[n^z_s(\lambda^z w^z - \frac{\lambda^z w^z n^z_s}{2}) - (1-n^z_s)\frac{\lambda^z w^z n^z_s}{2}\right]$$
$$(1-\beta)n^z_b\left[\lambda^z w^z n^z_s - \frac{\lambda^z w^z (n^z_s)^2}{2}) - \frac{\lambda^z w^z n^z_s}{2} + \frac{\lambda^z w^z (n^z_s)^2}{2}\right]$$

Since it holds $n^z_b = \frac{\lambda^z w^z n^z_s}{f_b}$, combining the two equations above, we obtain

$$CS^z = (1-\beta)\frac{(\lambda^z w^z n^z_s)^2}{2f_b} + \beta(\lambda^z w^z n^z_s).$$

Therefore, the consumer-preferred disclosure regime can be found by direct comparison of $\lambda^z w^z n^z_s$ across regimes. In particular, it holds that $CS^{z1} > CS^{z1}$ if and only if

$$\lambda^{z_1}(1-\lambda^{z_1})(w^{z_1})^2(\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_2})\lambda^{z_2}(w^{z2})^2)$$
$$> \lambda^{z_2}(1-\lambda^{z_2})(w^{z_2})^2(\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_1})\lambda^{z_1}(w^{z_1})^2),$$

which can be rewritten as

$$\overline{f}_s\overline{f}_b\left[\lambda^{z_1}(1-\lambda^{z_1})(w^{z_1})^2 - \lambda^{z_2}(1-\lambda^{z_2})(w^{z_2})^2\right] > 0.$$

Consider again the equilibrium values reported in Lemma 5. It is immediate to see that RE is always preferred to NE, as $\lambda^{RE} = \lambda^{NE}$ and $w^{RE} = w^{NE}$. RE is also preferred to SE if and only if

$$\lambda^{RE}(1-\lambda^{RE}) - \lambda^{SE}(1-\lambda^{SE}) > 0,$$

and since $\lambda^{RE} = \frac{1}{3}$ and $\lambda^{SE} = \frac{1}{4}$, it follows that the above is always satisfied. Hence, CS is always maximized by $z = RE$.

Producer surplus Producer surplus is straightforward to obtain. Given our distributional assumption on the opportunity cost of joining the platform, a seller in a product category gets on average half of the gains from trade retained after accounting for the platform fee. Therefore:

$$PS^z = n_s^z\left(\frac{1}{2}(1-\alpha_s^*)(1-\lambda^z)w^z(\beta+(1-\beta)n_b)\right)$$

and since $n_s = \frac{(1-\alpha_s^*)(1-\lambda^z)w^z(\beta+(1-\beta)n_b)}{\overline{f}_s}$:

$$PS^z = (n_s^z)^2\frac{\overline{f}_s}{2}$$

Since $\overline{f}_s$ is the same in all disclosure regimes, the one preferred by the sellers can be identified by direct comparison of n_s^z across the three regimes. In particular, regime z_1 is preferred to regime z_2 if and only if

$$\frac{\beta(1-\lambda^{z_1})w^{z_1}\overline{f}_b}{2[\overline{f}_s\overline{f}_b-(1-\beta)(1-\lambda^{z_1})\lambda^{z_1}(w^{z_1})^2]} > \frac{\beta(1-\lambda^{z_2})w^{z_2}\overline{f}_b}{2[\overline{f}_s\overline{f}_b-(1-\beta)(1-\lambda^{z_2})\lambda^{z_1}(w^{z_2})^2]}.$$

Since β and $\overline{f}_b$ are the same in all regimes, this condition is equivalent to

$$\begin{aligned}&\overline{f}_s\overline{f}_b[(1-\lambda^{z_1})w^{z_1}-(1-\lambda^{z_2})w^{z_2}]\\&>(1-\beta)w^{z_1}w^{z_2}(1-\lambda^{z_1})(1-\lambda^{z_2})[\lambda^{z_2}w^{z_2}-\lambda^{z_1}w^{z_1}].\end{aligned}$$

Taking the values λ^z and w^z from Lemma 5, it is immediate to see that RE is always preferred to NE since $\lambda^{RE} = \lambda^{NE}$ and $w^{RE} > w^{NE}$; RE is also preferred to SE if and only if

$$\overline{f}_s\overline{f}_b[(1-\lambda^{RE})-(1-\lambda^{SE})] > (1-\beta)w^2(1-\lambda^{RE})(1-\lambda^{SE})[\lambda^{SE}-\lambda^{RE}]$$

or, equivalently,

$$2\overline{f}_s\overline{f}_b > (1-\beta)w^2,$$

which is the expression given in Proposition 5.

Details on the consumer-seller interaction with duopoly sellers

This appendix reproduces the findings of Ali et al. (2023) with duopoly sellers which, combined, provide a formal proof of Lemma 6. Values for π and u as reported in the main text are derived for each of the three disclosure regimes. Values for $w = \pi + u$ and shares retained by consumers (λ) and sellers ($1-\lambda$) follow immediately.

Under the duopoly specification, consumers are uniformly distributed on $[0,1]$ and a consumer with characteristic x obtains utility $V - tx - p_1$ and $V - t(1-x) - p_2$ buying from sellers 1 and 2, respectively. t is measures the degree of product differentiation and V is the stand-alone utility assumed to be high enough to cover the market. Consumers make disclosure decisions and then sellers set prices simultaneously, where prices can be conditioned on the information received from consumers. Active disclosure leads to different equilibrium values under our assumptions on the equilibrium selection. Under each disclosure regime, trade is always efficient and, thus, total gains from trade can be shown to be $w = V - \frac{t}{4}$.

The "no disclosure" case is the standard Hotelling setting: when consumers cannot disclose their location, sellers set the same equilibrium price to all consumers and split demand in the middle, every consumer purchases from the closer seller at uniform equilibrium prices $p_1^* = p_2^* = t$. A seller's equilibrium profit is $\pi^{NE} = \frac{1}{2}t$ and consumers can be shown to obtain on average $u = V - (5/4)t$.

Formally, the consumers' participation condition is:

$$u^{NE} = \int_0^{\hat{x}} (V - tx - p_1)dx + \int_{\hat{x}}^1 (V - t(x-1) - p_2)dx \geq c$$

where $\hat{x}$ is the consumer indifferent between purchasing from either seller given prices, that is:

$$\hat{x} : V - t\hat{x} - p_1 = V - t(1-\hat{x}) - p_2 \quad \rightarrow \quad \hat{x} = \frac{1}{2} + \frac{p_2 - p_1}{2t}.$$

Expected utility is then

$$u^{NE} = V - \hat{x}p_1 - (1-\hat{x})p_2 - \frac{t\hat{x}^2}{2} - \frac{t(1-\hat{x})^2}{2}.$$

The maximization problem of seller $i = 1, 2$ is straightforward; given a fixed demand (which then acts as a scalar and can be omitted):

$$\max_{p_i} \pi_i = p_i \left[\frac{1}{2} + \frac{p_j - p_i}{2t}\right]$$

where the second component is the standard Hotelling demand that depends on $\hat{x}$.

From the system of first-order conditions of profit maximization, we obtain equilibrium prices $p_1^* = p_2^* = t$, which, after plugging them into the expression for u^{NE} and π^{NE}, gives values:

$$u^{NE} = V - \frac{5}{4}t \qquad \pi^{NE} = \frac{1}{2}t$$

Under the simple evidence disclosure regime, consumers can choose to disclose only their exact location to either, neither or both sellers by sending private messages $M_1(x), M_2(x)$ to sellers 1 and 2, respectively. Optimal disclosure by consumers in this setting is given in the following result:

Lemma 7. *(Propositions 6 and 7 in Ali et al., 2023) With simple evidence, the consumers' preferred equilibrium strategy is partial pooling and contains the following messages:*

$$(M_1^*(x), M_2^*(x)) = \begin{cases} ([0,1], \{x\}) & \text{if } x \in [0, \underline{x}], \\ (\{x\}, \{x\}) & \text{if } x \in (\underline{x}, \overline{x}), \\ (\{x\}, [0,1]) & \text{if } x \in [\overline{x}, 1]. \end{cases}$$

All consumers purchase from the closest seller.

In the above, $\underline{x}$ and $\overline{x}$ represent the consumer type who is indifferent between the general offer by seller 1 and a zero price offer by seller 2 and viceversa.

The intuition is the following: consumers have an incentive to disclose their location if and only if by doing so they are offered the product at a lower prices. Consumers located close to either of the sellers cannot realistically threaten to purchase from the distant seller. They choose not to disclose their location to the close seller and purchase at the general price offered by the closer seller. They do, however, have an incentive to disclose their location to the distant seller, since this will trigger a personalized zero price offer, which makes the constraint more binding. The general price offered by the seller is a constrained monopoly price directed at the segment of consumers closest to them, with the constraint stemming from the zero price offer made by the competitor.

Consumers located close to the middle of the distribution have an incentive to share their location to both sellers: by doing so, they communicate to their preferred seller that they are close enough to the competition to purchase from them if their price is low enough. This implies that sellers will compete in an asymmetric Bertrand model for each of these consumer locations: the consumer located at the center will receive a zero price offer from both sellers and those relatively closer to the center will face a price not far above zero (as in Thisse and Vives, 1988).

Compared to the no evidence case, the overall result of this interaction is a drop in prices for all consumers. The resulting profit and expected utility are $\pi^{SE} = \frac{3}{16}t$ and $u = V - (5/8)t$. Thus, consumers have a higher expected utility from trade than under no evidence, while simple evidence has a negative effect on sellers' profits because of more intense competition.

The expected utility of consumers depends on three threshold values for x: the consumer indifferent between seller 1's general offer and a zero price offer by seller 2 (*i.e.* $\underline{x}$), the consumer indifferent between seller 2's general offer and a zero price offer by seller 1 (i.e. $\overline{x}$) and the consumer indifferent between a zero price offer by both sellers (i.e. $\hat{x} = 1/2$). The expected utility under simple evidence is:

$$u^{SE} = \int_0^{\underline{x}} (V - tx - p_1)dx + \int_{\underline{x}}^{\hat{x}} (V - tx - p_1(x))dx + \int_{\hat{x}}^{\overline{x}} (V - t(x-1) - p_2(x))dx + \int_{\overline{x}}^{1} (V - t(x-1) - p_2)dx,$$

where $p_1(x)$ and $p_2(x)$ are the personalized prices consumer will get if they disclose to both sellers. These personalized prices are determined by:

$$V - tx - p_1(x) = V - t(1-x) \quad \leftrightarrow \quad p_1(x) = t(1-2x)$$
$$V - tx = V - t(1-x) - p_2(x) \quad \leftrightarrow \quad p_2(x) = t(2x-1)$$

The relevant thresholds can be obtained by setting equal the general offer of the closer seller and a zero price offer from the other (or, in the case of $\hat{x}$, confronting a zero price offer from both sellers), leading to:

$$\underline{x} = \frac{1}{2} - \frac{p_1}{2t} \qquad \hat{x} = \frac{1}{2} \qquad \overline{x} = \frac{1}{2} + \frac{p_2}{2t}$$

Plugging in all of the above and solving, we obtain that

$$u^{SE} = V - \frac{t}{4} - \underline{x}p_1 - (1-\overline{x})p_2 - t([x - x^2]_{\underline{x}}^{\hat{x}} + [x^2 - x]_{\hat{x}}^{\overline{x}})$$
$$= V - \frac{t}{4} - \left(\frac{1}{2} - \frac{p_1}{2t}\right)p_1 - \left(\frac{1}{2} - \frac{p_2}{2t}\right)p_2 - t\left(\frac{p_1^2 + p_2^2}{4t^2}\right),$$

which leads to

$$u^{SE} = V - \frac{2(p_1 + p_2) + t}{4} + \frac{p_1^2 + p_2^2}{4t}.$$

The sellers' maximization problem with given demand under "simple evidence" is

$$\max_{p_i} \pi_i = \left[p_i\left(\frac{1}{2} - \frac{p_i}{2t}\right) + \left(\frac{p_i^2}{4t}\right)\right],$$

where $(\frac{p_i^2}{4t})$ are the profits sellers make by selling to consumers at personalized prices. The optimal base price ($p_1 = p_2 = \frac{t}{2}$) can be found from the first-order conditions; it then follows:

$$f_b^{SE} = V - \frac{5}{8}t \qquad \pi^{SE} = \frac{3}{16}t$$

Consider now the consumers' "rich evidence" disclosure strategy. When consumers have more control over their information, they can achieve the same equilibrium found in the "simple evidence" case, but there exists an equilibrium in which their expected utility is even higher. As shown in Ali et al. (2023), the optimal disclosure strategy for consumers in this case is to partially pool messages to the closer seller in a way that mirrors the optimal constrained price on the remaining demand given other consumers' messages, a procedure that, without additional constraints to the messages, can be iterated infinitely. This generates an infinite set of prices offered by both sellers targeting different consumer segments. Prices decrease as the targeted segment is further away from the seller's location.

Consumers' expected utility reflects the associated partition. We define prices $p_{i,k}$ and thresholds $\underline{x}_k$ and $\overline{x}_k$, $k = 0, 1, ...$ as follows:

- $\{p_{i,k}\}_{k=0,1,...}$ offered by seller $i = 1, 2$ is such that $p_{i,k} \geq p_{i,k+1} \geq 0\ \ \forall k$,
- $\{\underline{x}_k\}_{k=0,1,...}$ are such that $\underline{x}_0 = 0$ and $\underline{x}_k : \frac{1}{2} - \frac{p_{1,k-1}}{2t}\ \ \forall k > 0$,
- $\{\overline{x}_k\}_{k=0,1,...}$ are such that $\overline{x}_0 = 1$ and $\overline{x}_k : \frac{1}{2} + \frac{p_{2,k-1}}{2t}\ \ \forall k > 0$.

Each price $p_{i,k}$ determines a segment that ends at threshold location x_{k+1} and starts at location x_k, which is either the same location at which the segment determined by the previous, weakly higher price $p_{i,k-1}$ ends or the relevant border location ($\underline{x}_0 = 0$ or $\overline{x}_0 = 1$). The threshold themselves are such that consumers at the end of a segment are indifferent between the relevant "group" price offered by the closer seller and a zero price offered by the distant one. As k grows, prices go down and segments shrink in size and get closer and closer to the center of the distribution; at the limit, $p_{i,\infty} = 0, i = 1, 2$ and $\underline{x}_\infty = \overline{x}_\infty = \frac{1}{2}$.

The following lemma summarizes Ali et al. (2023)'s findings:

Lemma 8. *(Propositions 6 and 8 in Ali et al., 2023) With rich evidence, there exists an equilibrium in which a consumer's reporting strategy is to send the following message to both sellers:*

$$M^*(x) = \begin{cases} (\underline{x}_k, \underline{x}_{k+1}] & \text{if } \underline{x}_k < x \leq \underline{x}_{k+1}, \\ (\overline{x}_{k+1}, \overline{x}_k] & \text{if } \overline{x}_{k+1} > x \geq \overline{x}_k, \\ \frac{1}{2} & \text{if } x = 1/2. \end{cases}$$

After receiving such messages, the closer seller i charges p_i^k and distant seller j charges 0. All consumers purchase from the closest seller. This is the consumer-preferred equilibrium.

Since consumers are not limited to fully reveal their location anymore, they can make better use of the information asymmetry to lower prices even further compared to the "simple evidence" case. To understand why this works, consider again the equilibrium disclosure in the simple evidence case. Consumers located at the extremes purchase at the constrained monopoly price as mentioned above. Consumers located in the center could only share their exact location and, by doing so, trigger low personalized prices by making said constraint more binding. Suppose however that sellers could not exactly identify consumers beyond them being "central", "very close", or "very far". When setting the price for the central group of consumers, they would set another constrained monopoly price to maximize profit in that section. The consumers' optimal strategy is then to distinguish themselves: those that would happily purchase at such a price and those who would do better by disclosing their "even more central" location. This procedure can be iterated infinitely on the "leftover" segment of consumers: in equilibrium, then, the only consumers fully revealing their location to both sellers are the central ones ($x = \frac{1}{2}$), who then receive a zero price offer from both sellers. By mirroring the constrained monopoly prices given the truncated distribution of other consumers, then, consumers can achieve a higher expected utility, again at the expense of sellers. The resulting profit and expected utility can be shown to be $\pi^{RE} = \frac{1}{6}t$ and $u = V - (7/12)t$ respectively.

Formally, expected utility given the buyers' preferred disclosure strategy is

$$u^{RE} = \sum_{k=0}^{\infty} \left\{ \int_{\underline{x}_k}^{\underline{x}_{k+1}} (V - tx - p_{1,k})dx \right\} + \\ + \sum_{k=0}^{\infty} \left\{ \int_{\overline{x}_{k+1}}^{\overline{x}_k} (V - t(1-x) - p_{1,k})dx \right\}$$

which can be rewritten as follows:

$$u^{RE} = V - \frac{t}{4} \\ - \left\{ p_{1,0} \left[\frac{1}{2} - \frac{p_{1,0}}{2t} \right] + p_{1,1} \left[\frac{1}{2} - \frac{p_{1,1}}{2t} - \frac{1}{2} + \frac{p_{1,0}}{2t} \right] + \dots \right\} \\ - \left\{ p_{2,0} \left[\frac{1}{2} - \frac{p_{2,0}}{2t} \right] + p_{2,1} \left[\frac{1}{2} - \frac{p_{2,0}}{2t} - \frac{1}{2} + \frac{p_{2,1}}{2t} \right] + \dots \right\}$$

Thus, we have

$$u^{RE} = V - \frac{2(p_1 + p_2) - t}{4} + \frac{\sum_{k=0}^{\infty} [p_{1,k}(p_{1,k} - p_{1,k+1}) + p_{2,k}(p_{2,k} - p_{2,k+1})]}{2t}.$$

The sellers' maximization problems given fixed demand are:

$$\max_{\{p_{1,k}\}_{k=0,1,\dots}} \pi_1 = \left[\sum_{k=0}^{\infty} p_{1,k} \left(\underline{x}_{k+1} - \underline{x}_k \right) \right]$$

$$\max_{\{p_{2,k}\}_{k=0,1,\dots}} \pi_2 = \left[\sum_{k=0}^{\infty} p_{2,k} \left(\overline{x}_k - \overline{x}_{k+1} \right) \right]$$

The optimal disclosure inducing optimal constrained monopoly pricing on the truncated distribution of consumers makes solving the maximization problem straightforward. Rearranging the system of first-order conditions of profit maximization, gives equilibrium prices:

$$p^*_{i,k} = \frac{t}{2^{k+1}} \qquad k = 0, 1, 2, ...$$

and the respective segments of consumers are immediately identifiable. For buyers closer to seller 1 (seller 2), the optimal message defines the relevant segment and can be expressed as:

$$M(x) = \{[\underline{x}_k, \underline{x}_{k+1}) : \underline{x}_k \leq x \leq \underline{x}_{k+1}\}$$

$$(M(x) = \{(\overline{x}_{k+1}, \overline{x}_k] : \overline{x}_{k+1} \leq x \leq \overline{x}_k\})$$

for $k = 0, 1, 2, ...$

The per-consumer profit can be then expressed as:

$$\pi_i = \frac{t}{2} \cdot \frac{1}{4} + \frac{t}{4} \cdot \frac{1}{8} + \frac{t}{8} \cdot \frac{1}{16} + \dots \\ = \frac{t}{2} \sum_{k=1}^{\infty} \left(\frac{1}{4} \right)^k = \frac{t}{6}$$

Finally, equilibrium prices lead to $u^{RE} = V - \frac{7}{12}t$ and $\pi^{RE} = \frac{1}{6}t$.

Duopoly sellers

This appendix provides the proof for Proposition 6.

Profit-maximizing platform Platform profits are given by

$$\Pi^z = \frac{\beta^2(1-\lambda)^2 w^2 \overline{f}_b}{4[2\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^z)\lambda^z(w^z)^2]}.$$

Since β, $\overline{f}_b$, and w are common across regimes, it holds that $\pi^{z_1} > \pi^{z_2}$ if and only if

$$(1-\lambda^{z_1})^2(2\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_2})\lambda^{z_2}w^2) > \\ (1-\lambda^{z_2})^2(2\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_1})\lambda^{z_1}w^2),$$

which, after rearranging, is equivalent to

$$2\overline{f}_s\overline{f}_b[(1-\lambda^{z_1})^2 - (1-\lambda^{z_2})^2] > \\ (1-\beta)w^2(1-\lambda^{z_1})(1-\lambda^{z_2})[(1-\lambda^{z_1}) - (1-\lambda^{z_2})].$$

Consider λ^z, w^z for all the regimes as in Lemma 6. The expression above implies that RE is preferred to SE as long as

$$128\overline{f}_s\overline{f}_b - 3(1-\beta)t(8w-3t) > 9(18\overline{f}_s\overline{f}_b - (1-\beta)t(3w-t))$$

or, equivalently,

$$3t(1-\beta)w > 34\overline{f}_s\overline{f}_b.$$

SE is preferred to NE as long as

$$9(2\overline{f}_s\overline{f}_b - (1-\beta)t(w-t)) > 128\overline{f}_s\overline{f}_b - (1-\beta)3t(3w-t))$$

or, equivalently,

$$3t(1-\beta)w > 22\overline{f}_s\overline{f}_b.$$

Finally, RE is preferred to NE as long as:

$$2\overline{f}_s\overline{f}_b - (1-\beta)t(w-t) > 18\overline{f}_s\overline{f}_b - (1-\beta)t(3w-t))$$

or, equivalently,

$$t(1-\beta)w > 8\overline{f}_s\overline{f}_b.$$

The former two conditions are equivalent to the threshold conditions reported in Proposition 6. The last condition can be shown to be never relevant: since it holds that RE is preferred to NE if and only if:

$$w > \frac{8\overline{f}_s\overline{f}_b}{(1-\beta)t} = \tilde{w}$$

and since $w_1^P < \tilde{w} < w_2^P$, the result in Proposition 6 follows immediately.

Consumer surplus All observations made for the monopoly case still apply, and the expression for consumer surplus is unchanged:

$$CS^z = (1-\beta)\frac{(\lambda^z w^z n_s^z)^2}{2\overline{f}_b} + \beta(\lambda^z w^z n_s^z)$$

therefore, it holds that $CS^{z1} > CS^{z1}$ if and only if

$$\overline{f}_s \overline{f}_b (w)^2 \left[\lambda^{z_1}(1-\lambda^{z_1}) - \lambda^{z_2}(1-\lambda^{z_2})\right] > 0.$$

Since it holds that

$$\lambda^{RE}(1-\lambda^{RE}) = \frac{t}{3w}\frac{3w-t}{3w}, \qquad \lambda^{SE}(1-\lambda^{SE}) = \frac{3t}{8w}\frac{8w-3t}{8w}, \qquad \lambda^{NE}(1-\lambda^{NE}) = \frac{t}{w}\frac{w-t}{w},$$

and since t and w are the same in all regimes, it follows that RE is preferred to SE if and only if

$$\frac{3w-t}{9} > \frac{3(8w-3t)}{64} \quad \Longleftrightarrow \quad w < \frac{17}{24}t;$$

SE is preferred to NE if and only if

$$\frac{3(8w-3t)}{64} > w-t \quad \Longleftrightarrow \quad w < \frac{11}{8}t;$$

and, finally, RE is preferred to NE if and only if

$$\frac{3w-t}{9} > w-t \quad \Longleftrightarrow \quad w < \frac{4}{3}t.$$

Since $\frac{17}{24} < \frac{4}{3} < \frac{11}{8}$, the last condition is never relevant. The result given in Proposition 6 follows immediately.

Producer surplus The observations made in the monopoly seller case still apply, and the expression for producer surplus is unchanged since surplus of both sellers per product category are included:

$$PS^z = (n_s^z)^2 \frac{\overline{f}_s}{2}$$

Since $\overline{f}_s$ is the same in all disclosure regimes, the one preferred by the sellers can be identified by direct comparison of n_s^z across the three regimes. Regime z_1 is preferred to regime z_2 if and only if

$$\frac{\beta(1-\lambda^{z_1})w^{z_1}\overline{f}_b}{2[2\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_1})\lambda^{z_1}(w^{z_1})^2]} > \frac{\beta(1-\lambda^{z_2})w^{z_2}\overline{f}_b}{2[2\overline{f}_s\overline{f}_b - (1-\beta)(1-\lambda^{z_2})\lambda^{z_1}(w^{z_2})^2]}.$$

Since β, $\overline{f}_b$, and w are the same in all regimes, this condition is equivalent to

$$2\overline{f}_s\overline{f}_b > (1-\beta)w^2(1-\lambda^{z_1})(1-\lambda^{z_2}).$$

We have that $(1-\lambda^{RE}) = \frac{t}{3w}$, $(1-\lambda^{SE}) = \frac{3t}{8w}$, and $(1-\lambda^{NE}) = \frac{t}{w}$. Hence, RE is preferred to SE if and only if $2\overline{f}_s\overline{f}_b > (1-\beta)\frac{1}{8}t^2$; SE is preferred to NE if and only if $2\overline{f}_s\overline{f}_b > (1-\beta)\frac{3}{8}t^2$; SE is preferred to NE if and only if $2\overline{f}_s\overline{f}_b > (1-\beta)\frac{1}{3}t^2$. The last condition is never relevant. Therefore, the sellers' preferred disclosure regime is the one stated in Proposition 6.

References

Acquisti, A., Taylor, C., and Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54:442–492.

Adams, W. J. and Yellen, J. L. (1976). Commodity bundling and the burden of monopoly. *Quarterly Journal of Economics*, 90:475–498.

Aguiar, L. and Waldfogel, J. (2021). Platforms, power, and promotion: Evidence from Spotify playlists. *Journal of Industrial Economics*, 69:653–691.

Aguiar, L., Waldfogel, J., and Waldfogel, S. (2021). Playlisting favorites: Measuring platform bias in the music industry. *International Journal of Industrial Organization*, 78:1027–1065.

Aguiar, L., Waldfogel, J., and Zeijen, A. (2023). Platform power struggle: Spotify and the major record labels, 2017-2020. *Unpublished manuscript.*

Ali, S. N., Lewis, G., and Vasserman, S. (2023). Voluntary disclosure and personalized pricing. *Review of Economic Studies*, 90:538–571.

Anderson, S. P. and Bedre-Defolie, Ö. (forthcoming). Hybrid platform model: Monopolistic competition and a dominant firm. *RAND Journal of Economics.*

Anderson, S. P. and Coate, S. (2005). Market provision of broadcasting: A welfare analysis. *Review of Economic Studies*, 72:947–972.

Anderson, S. P., Engers, M., and Savelle, D. (2020). Ordered search: Equilibrium and optimum. *Unpublished manuscript.*

Anderson, S. P., Engers, M., and Savelle, D. (2021). Consumer journeys with nested Weitzman scores and informational complementarities. *Unpublished manuscript.*

Anderson, S. P. and Gabszewicz, J. J. (2006). The media and advertising: A tale of two-sided markets. *Handbook of the Economics of Art and Culture*, 1:567–614.

Anderson, S. P. and Neven, D. J. (1989). Market efficiency with combinable products. *European Economic Review*, 33:707–719.

Anderson, S. P. and Renault, R. (2006). Advertising content. *American Economic Review*, 96:93–113.

Armstrong, M. (2006a). Competition in two-sided markets. *RAND Journal of Economics*, 37:668–691.

Armstrong, M. (2006b). *Recent Developments in the Economics of Price Discrimination*, volume 2 of *Econometric Society Monographs*, pages 97–141. Cambridge University Press.

Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association*, 15:989–1024.

Armstrong, M. and Zhou, J. (2011). Paying for prominence. *Economic Journal*, 121:F368–F395.

Armstrong, M. and Zhou, J. (2022). Consumer information and the limits to competition. *American Economic Review*, 112:534–577.

Arya, A., Mittendorf, B., and Sappington, D. E. (2007). The bright side of supplier encroachment. *Marketing Science*, 26:651–659.

Bao, R., Li, S., and Yu, J. (2022). Search with interdependent values. *Unpublished manuscript.*

Belleflamme, P. and Peitz, M. (2019). Managing competition on a two-sided platform. *Journal of Economics & Management Strategy*, 28:5–22.

Belleflamme, P. and Peitz, M. (2021). *The Economics of Platforms: Concepts and Strategy.* Cambridge University Press.

Bergemann, D. and Bonatti, A. (2019). Markets for information: An introduction. *Annual Review of Economics*, 11:85–107.

Bergemann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. *American Economic Review*, 105:921–957.

Berton, F., Dughera, S., and Ricci, A. (2021). Are unions detrimental to innovation: Theory and evidence. *IZA Discussion Papers Series - 14102*.

Bounie, D., Dubus, A., and Waelbroeck, P. (2021). Selling strategic information in digital competitive markets. *RAND Journal of Economics*, 52:283–313.

Bourreau, M. and Gaudin, G. (2022). Streaming platform and strategic recommendation bias. *Journal of Economics & Management Strategy*, 31:25–47.

Bourreau, M., Kraemer, J., and Hofmann, J. (2021). Prominence-for-data schemes in digital platform ecosystems: Implications for platform bias and consumer data collection. *Unpublished manuscript.* Available at SSRN 3867580.

Brander, J. A. and Eaton, J. (1984). Product line rivalry. *American Economic Review*, 74:323–334.

Brynjolfsson, E., Hu, Y., and Simester, D. (2011). Goodbye Pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Science*, 57:1373–1386.

Casner, B. (2020). Seller curation in platforms. *International Journal of Industrial Organization*, 72:1026–1059.

Chen, Y. and Schwartz, M. (2013). Product innovation incentives: Monopoly vs. competition. *Journal of Economics & Management Strategy*, 22:513–528.

Chiang, W. K., Chhajed, D., and Hess, J. D. (2003). Direct marketing, indirect profits: A strategic analysis of dual-channel supply-chain design. *Management Science*, 49:1–20.

Choi, J. P. and Jeon, D.-S. (2023). Platform design biases in ad-funded two-sided markets. *RAND Journal of Economics*, 54:240–267.

Choi, M., Dai, A. Y., and Kim, K. (2018). Consumer search and price competition. *Econometrica*, 86:1257–1281.

Cremer, H. and Thisse, J.-F. (1991). Location models of horizontal differentiation: A special case of vertical differentiation models. *Journal of Industrial Economics*, 39:383–390.

Crémer, J., de Montjoye, Y.-A., and Schweitzer, H. (2019). Competition policy for the digital era. *Technical report.* Available at https://op.europa.eu/en/publication-detail/-/publication/21dc175c-7b76-11e9-9f05-01aa75ed71a1.

Datta, H., Knox, G., and Bronnenberg, B. J. (2018). Changing their tune: How consumers' adoption of online streaming affects music consumption and discovery. *Marketing Science*, 37:5–21.

De Cornière, A. and Taylor, G. (2019). A model of biased intermediation. *RAND Journal of Economics*, 50:854–882.

Döpper, H. and Rasch, A. (2024). Combinable products, price discrimination, and collusion. *International Journal of Industrial Organization*, 94:1030–1071.

Doval, L. (2018). Whether or not to open Pandora's box. *Journal of Economic Theory*, 175:127–158.

Eaton, B. C. and Lipsey, R. G. (1979). The theory of market pre-emption: The persistence of excess capacity and monopoly in growing spatial markets. *Economica*, 46:149–158.

Elliott, M. and Galeotti, A. (2019). Market segmentation through information. *Unpublished manuscript.* Available at SSRN 3432315.

Fleder, D. and Hosanagar, K. (2009). Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. *Management Science*, 55:697–712.

Fletcher, A., Ormosi, P. L., Savani, R., and Castellini, J. (2023). Biased recommender systems and supplier competition. *Unpublished manuscript.* Available at SSRN 4319311.

Freimane, M. (2022). Substituting away? The effect of platform bargaining regulation on content display. *Unpublished manuscript.*

Fudenberg, D. and Tirole, J. (2000). Customer poaching and brand switching. *RAND Journal of the Economics*, 31:634–657.

Gal-Or, E. and Dukes, A. (2003). Minimum differentiation in commercial media markets. *Journal of Economics & Management Strategy*, 12:291–325.

Garcia, D. and Shelegia, S. (2018). Consumer search with observational learning. *RAND Journal of Economics*, 49:224–253.

Greminger, R. P. (2022). Optimal search and discovery. *Management Science*, 68:3904–3924.

Grout, P. A. (1984). Investment and wages in the absence of binding contracts: A Nash bargining approach. *Econometrica*, 52:449–460.

Haan, M. A., Moraga-González, J. L., and Petrikaitė, V. (2018). A model of directed consumer search. *International Journal of Industrial Organization*, 61:223–255.

Hagenbach, J. and Koessler, F. (2017). Simple versus rich language in disclosure games. *Review of Economic Design*, 21:163–175.

Hagiu, A. (2009). Two-sided platforms: Product variety and pricing structures. *Journal of Economics & Management Strategy*, 18:1011–1043.

Hagiu, A. and Wright, J. (2015). Multi-sided platforms. *International Journal of Industrial Organization*, 43:162–174.

Halmenschlager, C. and Mantovani, A. (2017). On the private and social desirability of mixed bundling in complementary markets with cost savings. *Information Economics and Policy*, 39:45–59.

Haucap, J. and Wey, C. (2004). Unionisation structures and innovation incentives. *Economic Journal*, 114:C149–C165.

Heidhues, P., Köster, M., and Kőszegi, B. (2023). Steering fallible consumers. *Economic Journal*, 133:1430–1465.

Hidir, S. and Vellodi, N. (2021). Privacy, personalization, and price discrimination. *Journal of the European Economic Association*, 19:1342–1363.

Hodgson, C. and Lewis, G. (2020). You can lead a horse to water: Spatial learning and path dependence in consumer search. *Unpublished manuscript.*

Hoernig, S. H. and Valletti, T. M. (2007). Mixing goods with two-part tariffs. *European Economic Review*, 51:1733–1750.

Hoernig, S. H. and Valletti, T. M. (2011). When two-part tariffs are not enough: Mixing with nonlinear pricing. *B.E. Journal of Theoretical Economics*, 11. Article 21.

Ichihashi, S. (2020). Online privacy and information disclosure by consumers. *American Economic Review*, 110:569–595.

Jeon, D.-S., Lefouili, Y., and Madio, L. (2022). Platform liability and innovation. *CESifo Working Paper Series - 9984.*

Johnson, J. P. and Myatt, D. P. (2002). Multiproduct quality competition: Fighting brands and product line pruning. *Unpublished manuscript.*

Johnson, J. P., Rhodes, A., and Wildenbeest, M. R. (2020). Platform design when sellers use pricing algorithms. *CEPR Discussion Paper Series - DP15504.*

Karle, H., Peitz, M., and Reisinger, M. (2020). Segmentation versus agglomeration: Competition between platforms with competitive sellers. *Journal of Political Economy*, 128:2329–2374.

Ke, T. T. and Lin, S. (2022). Informational complementarity. *Management Science*, 66:3699–3716.

Kline, P., Petkova, N., Williams, H., and Zidar, O. (2019). Who profits from patents? Rent-sharing at innovative firms. *Quarterly Journal of Economics*, 134:1343–1404.

Kohn, M. G. and Shavell, S. (1974). The theory of search. *Journal of Economic Theory*, 9:93–123.

Konishi, H. and Sandfort, M. T. (2002). Expanding demand through price advertisement. *International Journal of Industrial Organization*, 20:965–994.

Kumar, N. and Ruan, R. (2006). On manufacturers complementing the traditional retail channel with a direct online channel. *Quantitative Marketing and Economics*, 4:289–323.

Lambertini, L. and Mantovani, A. (2009). Process and product innovation by a multiproduct monopolist: A dynamic approach. *International Journal of Industrial Organization*, 27:508–518.

Lancaster, K. J. (1966). A new approach to consumer theory. *Journal of Political Economy*, 74:132–157.

Lefez, W. (2022). Price recommendations and the value of data: A mechanism design approach. *Unpublished manuscript.*

Lin, P. (2004). Process and product R&D by a multiproduct monopolist. *Oxford Economic Papers*, 56:735–743.

Liu, Q. and Serfes, K. (2004). Quality of information and oligopolistic price discrimination. *Journal of Economics & Management Strategy*, 13:671–702.

Mayzlin, D. and Shin, J. (2011). Uninformative advertising as an invitation to search. *Marketing science*, 30:666–685.

McAfee, R. P., McMillan, J., and Whinston, M. D. (1989). Multiproduct monopoly, commodity bundling, and correlation of values. *Quarterly Journal of Economics*, 104:371–383.

McCall, J. J. (1970). Economics of information and job search. *Quarterly Journal of Economics*, 84:113–126.

Montes, R., Sand-Zantman, W., and Valletti, T. (2019). The value of personal information in online markets with endogenous privacy. *Management Science*, 65:1342–1362.

Mukherjee, A. and Pennings, E. (2011). Unionization structure, licensing and innovation. *International Journal of Industrial Organization*, 29:232–241.

Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 18:301–317.

Nocke, V., Peitz, M., and Stahl, K. (2007). Platform ownership. *Journal of the European Economic Association*, 5:1130–1160.

Nocke, V. and Rey, P. (forthcoming). Consumer search and choice overload. *Journal of Political Economy.*

Nocke, V. and Schutz, N. (2018). Multiproduct-firm oligopoly: An aggregative games approach. *Econometrica*, 86:523–557.

Oestreicher-Singer, G. and Sundararajan, A. (2012). Recommendation networks and the long tail of electronic commerce. *MIS Quarterly*, 36:65–83.

Peitz, M. and Sobolev, A. (2022). Inflated recommendations. *CEPR Discussion Paper Series - DP17260.*

Peitz, M. and Valletti, T. M. (2008). Content and advertising in the media: Pay-tv versus free-to-air. *International Journal of Industrial Organization*, 26:949–965.

Pepall, L., Richards, D., and Norman, G. (2014). *Industrial Organization: Contemporary Theory and Empirical Applications.* John Wiley & Sons.

Petrikaitė, V. (2018). Consumer obfuscation by a multiproduct firm. *RAND Journal of Economics*, 49:206–223.

Popper, B. (2016). Spotify‘s discover weekly reaches 40 million users and 5 billion tracks streamed. *The Verge, May*, 25.

Preuss, M. (2023). Search, learning and tracking. *RAND Journal of Economics*, 54:54–82.

Rhee, B.-D. and Park, S.-Y. (2000). Online store as a new direct channel and emerging hybrid channel system. *Unpublished manuscript.*

Rhodes, A. (2015). Multiproduct retailing. *Review of Economic Studies*, 82:360–390.

Rhodes, A., Watanabe, M., and Zhou, J. (2021). Multiproduct intermediaries. *Journal of Political Economy*, 129:421–464.

Robert, J. and Stahl, D. O. (1993). Informative price advertising in a sequential search model. *Econometrica*, 61:657–686.

Ronayne, D. and Taylor, G. (2022). Competing sales channels with captive consumers. *Economic Journal*, 132:741–766.

Shen, J. (2015). Ex-ante preference in a consumer search market. *Unpublished manuscript.*

Smolin, A. (2020). Disclosure and pricing of attributes. *Unpublished manuscript.*

Sovinsky Goeree, M. (2008). Limited information and advertising in the US personal computer industry. *Econometrica*, 76:1017–1074.

Sutton, J. (1986). Vertical product differentiation: Some basic themes. *American Economic Review*, 76:393–398.

Teh, T.-H. (2022). Platform governance. *American Economic Journal: Microeconomics*, 14:213–254.

Teh, T.-H. and Wright, J. (2022). Intermediation and steering: Competition in prices and commissions. *American Economic Journal: Microeconomics*, 14:281–321.

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

Thomes, T. P. (2013). An economic analysis of online streaming music services. *Information Economics and Policy*, 25:81–91.

Tsay, A. A. and Agrawal, N. (2004). Channel conflict and coordination in the e-commerce age. *Production and Operations Management*, 13:93–110.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica*, 47:641–654.

Wolinsky, A. (1986). True monopolistic competition as a result of imperfect information. *Quarterly Journal of Economics*, 101:493–511.

Zhou, J. (2014). Multiproduct search and the joint search effect. *American Economic Review*, 104:2918–2939.

Eidesstattliche Versicherung gemäß §8 Absatz 2 Buchstabe a) der Promotionsordnung der Universität Mannheim zur Erlangung des Doktorgrades der Volkswirtschaftslehre (Dr. rer. pol.)

1. Bei der eingereichten Dissertation mit dem Titel "Three Essays on Consumer Search and Platform Design" handelt es sich um mein eigenständig erstelltes Werk, das den Regeln guter wissenschaftlicher Praxis entspricht.

2. Ich habe nur die angegebenen Quellen und Hilfsmittel benutzt und mich keiner unzulässigen Hilfe Dritter bedient. Insbesondere habe ich wörtliche und nicht wörtliche Zitate aus anderen Werken als solche kenntlich gemacht.

3. Die Arbeit oder Teile davon habe ich bislang nicht an einer Hochschule des In- oder Auslands als Bestandteil einer Prüfungs- oder Qualifikationsleistung vorgelegt.

4. Die Richtigkeit der vorstehenden Erklärung bestätige ich.

5. Die Bedeutung der eidesstattlichen Versicherung und die strafrechtlichen Folgen einer unrichtigen oder unvollständigen eidesstattlichen Versicherung sind mir bekannt. Ich versichere an Eides statt, dass ich nach bestem Wissen die reine Wahrheit erklärt und nichts verschwiegen habe.

   Die eingereichten Dissertationsexemplare sowie der Datenträger gehen in das Eigentum der Universität über.

Mannheim

Jacopo Gambato

Curriculum Vitae

2018 - 2024	University of Mannheim Center for Doctoral Studies in Economics Ph.D. in Economics (Dr. rer. pol.)
2020 - 2024	ZEW - Leibniz Centre for European Economic Research Junior Research Group Competition and Innovation Researcher
2015 - 2017	Bocconi University Master of Science in Economics and Social Sciences
2012 - 2015	Bocconi University Bachelor of Science in Economics and Social Sciences
2015	University of British Columbia Exchange Bachelor Student