Discussion Paper Series – CRC TR 224

Discussion Paper No. 421
Project B 03

Consumer Search, Steering and Choice Overload

Volker Nocke [1]
Patrick Rey [2]

April 2023

[1] Department of Economics and MaCCI, University of Mannheim, School of Economics, University of Surrey, Email: volker.nocke@gmail.com
[2] Toulouse School of Economics, University Toulouse Capitole, Email: patrick.rey@tse-fr.eu

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

Consumer Search, Steering and Choice Overload*

Volker Nocke[†] Patrick Rey[‡]

April 19, 2023

Abstract

We develop a model of within-firm sequential, directed search and study a firm's ability and incentive to steer consumers. We find that the firm often benefits from adopting a noisy positioning strategy, which limits the information available to consumers. This induces consumers to keep searching but discourages some of them from visiting the firm. This occurs even though the firm and the consumers have in common the interest of maximizing the probability of trade. Because of such noisy positioning, an increase in the size of the product line further discourages consumers from visiting the firm—consistent with *choice overload*.

JEL Classification: L12, L15, D42.

Keywords: Consumer search, sequential search, directed search, product variety, choice overload, multiproduct firm, platform, steering.

*Previous versions of this paper were circulated under the title "Consumer Search and Choice Overload." We wish to thank the Editor (Emir Kamenica), three anonymuous referees, Daniel Garcia, Paul Heidhues, Federico Innocenti, Bruno Jullien, Teddy Kim, Martin Peitz, Andrew Rhodes, Sandro Shelegia, John Thanassoulis, Nikhil Vellodi, Lucy White, Jidong Zhou, seminar participants at UT Austin, UC Berkeley, HU Berlin, Bocconi University, Cambridge University, CY Paris Cergy University, European University Institute, Goethe University Frankfurt, Hebrew University Jerusalem, ITAM, London School of Economics, LUISS Rome, University of Mannheim, University of Manchester, Paris School of Economics, University of Rochester, Santa Clara University, Tel Aviv University, Tilburg University, Toronto University, Toulouse School of Economics, University of Surrey, University of Washington, and the Consumer Search Digital Seminar Series as well as participants at the 2015 MaCCI Summer Institute in Competition Policy (Erftstadt), the 2015 MaCCI Workshop on Multiproduct Firms (Bad Homburg), the Spring 2021 CRC TR 224 conference, the 2021 Fall Meeting of the Standing Committee for Industrial Organization of the Verein für Socialpolitik (Düsseldorf), the 2021 Asia-Pacific Industrial Organization Conference, the 2021 University of Bergamo Winter IO Conference (Passo di Tonale), the 2022 Annual Meeting of the Standing Committee for Economic Theory of the Verein für Socialpolitik (Berlin), the 2022 MaCCI Summer Institute in Competition Policy (Puchberg), and the 2022 Workshop on Consumer Search and Switching Costs at NYU. Nocke gratefully acknowledges financial support from the German Research Foundation (DFG) through CRC TR 224 (Project B03). Rey gratefully acknowledges financial support from the European Research Council (ERC) under the European Community's Seventh Framework Programme (FP7/2007-2013) Grant Agreement N 340903 and from the ANR Investissements d'Avenir program under the grant CHESS ANR-17-EURE-0010.

†Department of Economics and MaCCI, University of Mannheim, Germany; School of Economics, University of Surrey, UK. Email: volker.nocke@gmail.com

‡Toulouse School of Economics, University Toulouse Capitole, France. Email: patrick.rey@tse-fr.eu

1 Introduction

In today's digital economy, consumers face a large and ever-growing number of products from which to choose. Finding out which product fits best (and what its final price is) has become a complicated and time-consuming task for which consumers increasingly rely on *steering* in the form of recommendations, rankings and the like. Indeed, experimental evidence suggests that, absent such help, consumers can suffer from *choice overload* and refrain from purchasing. For example, Iyengar and Lepper (2000) document that an increase in the number of flavors sold at a jam tasting booth in an upscale grocery store induced a significant reduction in the share of customers making a purchase.[1]

Firms' ability to steer consumer search—both online and offline—raises the concern that they may do so in a way that is self-serving rather than in consumers' interest. For instance, Petroski (2003) reports that supermarkets and other brick-and-mortar shops often place more popular products at the back of their stores. Likewise, McDevitt (2014) finds that plumbers appearing on Google's sponsored search links receive substantially worse ratings on Yelp than those that do not.[2] And in recent years, steering by large online platforms has been increasingly scrutinized by consumer protection and competition agencies,[3] and has led policy-makers to propose—and, in some instances, adopt—regulations aimed at curbing steering practices.[4] To investigate this issue, we develop a model of within-firm sequential, directed search and study a firm's ability and incentive to steer consumers.

The paper's main insight is that the firm often benefits from adopting a noisy positioning strategy, which limits the information available to consumers. This induces consumers to keep searching, but it also discourages some of them from visiting the firm. Such noisy positioning arises even though the firm and the consumers have in common the interest of maximizing the probability of trade—in particular, our setting abstracts from any self-preferencing[5] or bias in favor of particular products.[6] Because

[1]More recently, in a study of an online restaurant-to-consumer delivery platform, Natan (2022) finds that expanding the set of restaurants discourages consumer search.

[2]McDevitt (2014) also reports that plumbing services that are listed first in the Yellow pages receive more complaints, and charge higher prices, than their rivals. For more empirical evidence on steering see, e.g., Hannak et al. (2014).

[3]See, e.g., the June 2017 *Google Shopping* decision of the European Commission (Case AT.39740 *Google Search (Shopping)*, confirmed in November 2021 by the EU General Court), and the complaint filed by the US Department of Justice against Google in October 2020.

[4]For instance, the *Digital Markets Act* recently adopted by the European Union imposes a ban on self-preferencing on the largest online platforms, and the companion *Digital Services Act* requires them to clarify the role of their algorithms.

[5]See Lee and Musolff (2021) and Hunold et al. (2022) for empirical evidence on self-preferencing by online platforms.

[6]See, for instance, Shih, Kaufmann, and Spinola (2007) for empirical evidence on how online platforms steer users towards more profitable products; see also, e.g., Rayo and Segal (2010) for a

of noisy positioning, an increase in the *size* of the product line further discourages consumers from visiting the firm—a form of choice overload.

In the model, a number of available products differ in *popularity*, i.e., in the probability with which any consumer has a match with that product. Consumers differ in their match-conditional valuations and in their search costs. A monopolist decides which products to offer and at which prices. The firm also decides how to *position* the offered products—that is, which product to allocate to which *slot*. Slots may correspond to aisles or shelves in physical stores or to positions on recommendation lists or in rankings in e-commerce.

Consumers observe the size of the product line, but neither the identity of the products, nor their prices or positions. Upon inspecting a slot, a consumer learns whether or not he has a match with the product offered in that slot and, if so, his match-conditional valuation. He uses this information to update his beliefs about the not-yet-inspected slots. After each inspection, the consumer not only decides whether to continue searching but also which slot to inspect next, so search is not only sequential but also *directed*.

The equilibrium analysis in Section 3 focuses on two benchmarks. First, we show that there exists a unique pure-strategy equilibrium. In this *pure positioning* equilibrium, the firm offers all available products at the monopoly price. Conditional on monopoly pricing, this equilibrium maximizes consumer surplus: as consumers are perfectly informed, they inspect products in decreasing order of popularity and use an optimal stopping rule. Hence, if an inspection delivers no match, consumers with high search costs drop out whereas the others inspect the slot holding the next-most popular product.

Second, we show that, for any size of the product line, there exists a *random positioning* equilibrium in which the firm offers the most popular products at the monopoly price and uniformly randomizes over its positioning strategy—that is, it assigns each offered product to every slot with equal probability. As a result, consumers are indifferent as to the order in which to inspect slots. Those consumers who start searching continue to do so until finding a match. This is because consumers who do not find a match at a given inspection become increasingly optimistic about finding a match at the next one. Search therefore has the potential to be *addictive*: some of the consumers unwilling to start searching would, if coerced to make a first inspection, voluntarily keep searching afterwards—others (those with higher search costs) would need to experience more than one inspection before becoming addicted to searching.

Under random positioning, increasing the size of the product line has two opposing effects: on the one hand, it offers more opportunities for a match; on the other,

theoretical analysis in a Bayesian persuasion setting.

it reduces the average popularity among the products on offer and thus lowers the probability of finding a match on any given inspection. We show that the latter effect dominates for the marginal searcher, implying that the number of consumers who visit the firm is decreasing in the assortment size—consistent with choice overload.

Fully randomizing over the positioning of the products makes consumers' search very inefficient. Indeed, the outcome is the same as if consumers, unable to distinguish between slots, engage in completely random search. Despite this, the firm may be better off under such random positioning than under pure positioning: while pure positioning maximizes the *extensive search margin*, random positioning maximizes instead the *intensive search margin* as all active consumers keep searching until finding a match.

In Section 4, we study the general class of *steering* equilibria (with monopoly pricing) in which the firm induces all consumers to inspect the slots in the same order—some may, however, stop earlier than others. We show that, in all such equilibria, the allocation of products to slots exhibits a *block structure*: the firm allocates the most popular products to the first block, the next-most popular products to the second one, and so on. Furthermore, active consumers inspect all slots in the first block until finding a match; if they do not find a match in the first block, consumers with high search costs drop out, whereas the others inspect the next block, and so on.

To steer consumers, the firm tends to allocate more popular products to the first slots. However, to induce *intensive search* (i.e., no drop-outs) within each block, it must introduce a sufficient amount of noise in its positioning. This is required to ensure that, in the absence of a match with the most promising slot, the marginal consumer (i.e., the consumer indifferent between inspecting or not that slot) becomes sufficiently optimistic about the next-most promising slot, and so on.

We show that, for any given block, the worst (for consumers and the firm) positioning strategy entails uniform randomization whereas the best entails just enough randomization to induce the marginal consumer to keep searching. In addition, expanding the size of a block discourages some consumers from visiting it.

Focusing on the most profitable block structure,[7] we first show that, holding fixed the distribution of search costs, expanding the blocks over which the firm randomizes its positioning can be profitable only if their products are neither too different not too similar. However, holding fixed the distribution of product popularities and assuming a constant hazard rate for the distribution of search costs, placing all slots in one single block is optimal for the firm if the hazard rate is large enough.

We consider various extensions and robustness checks in Section 5. We show that

[7] As mentioned above, pure positioning—where each block consists of a single slot—is best for consumers.

if the firm could perfectly condition on consumers' search costs, then the firm would always want to introduce enough noise to induce intensive (directed) search, and to offer the maximal number of products that consumers would be willing to inspect—a strategy that leaves very little surplus to consumers. We also discuss how the equilibrium analysis would be affected if the firm could disclose the identity of a product upon inspection, if there were aggregate uncertainty about the popularity of products, and if search costs were increasing or decreasing with the number of inspections. Finally, we consider the case where product selection and positioning is determined by a platform but product pricing is determined by individual sellers on the platforms. For that case, we also discuss what would happen if the platform were remunerated via "clicks" rather than sales.

Related Literature. Our paper is related to several strands of literature.[8] First, it builds on the consumer search literature. In the classic Wolinsky (1986) model of search among single-product firms and in the literature that builds upon it, products are symmetric in that consumers' valuations are drawn from the same distribution, and consumers sequentially search for the best match. We contribute to that literature in several ways. First, we introduce product heterogeneity in an analytically tractable way by assuming that products differ in *popularity* (measured by the probability with which a consumers likes a product), as in Kamenica (2008) and Chen and He (2011); we also allow consumers to differ not only in their match values but also in their search costs.[9] Second, most of this literature studies non-directed search, in which consumers randomly pick the next firm to visit. Our focus on directed search is instead related to the notion of *firm prominence*.[10] For example, Armstrong, Vickers and Zhou (2009) show that, in addition to a symmetric equilibrium featuring random search, there also exists an equilibrium with price dispersion in which all consumers correctly expect some firm to be cheaper than others, inducing them to inspect it first. More recently, Anderson and Renault (2021) consider a setting with asymmetric firms and provide conditions under which any search order can be sustained in equilibrium. However, this literature focuses on the role of prices as a driver of search between firms; we study instead the role of positioning as steering search within a firm. Third, the existing literature mostly focuses on (competing) single-product firms, and the few papers that do consider multiproduct firms (e.g., Zhou, 2014; Rhodes, 2015; Rhodes et al., 2021) typically assume that, upon visiting a multiproduct firm, a consumer automatically

[8]We focus here on steering and consumer search behavior by fully rational consumers. There is also a large literature studying consumers' behavioral biases that firms can exploit.

[9]Moraga-Gonzalez et al. (2017) introduce heterogeneous search costs in a version of the Wolinsky model with an infinite number of firms.

[10]See Armstrong (2017) for a survey.

learns his match values (and prices) for all of the products offered by that firm (i.e., there is no within-firm search).[11]

Our analysis of the firm's positioning strategies is related to recent models on consumer obfuscation and search diversion. Petrikaitė (2018) studies a Wolinsky-type model of sequential search where a monopolist sells two (symmetric) products and chooses not only prices but also product-specific search costs.[12] Consumers observe these search costs and then decide which product to inspect first; that is, search is directed. In equilibrium, the firm engages in price discrimination by introducing a search cost for one of the two products and offering it at a lower price. By contrast, we consider the case of heterogeneous products; furthermore, as search costs are fixed (and identical across products but heterogeneous across consumers), the firm has no incentive to engage in price discrimination. Hagiu and Jullien (2011) analyze search diversion by a monopolistic intermediary. On the platform, there are two independent sellers and the intermediary receives an exogenous seller-specific fee for making a consumer visit a seller. The intermediary observes consumers' types and ex ante commits to seller-specific diversion probabilities (i.e., probabilities with which a consumer has to visit first the less preferred seller). Hagiu and Jullien show that, if prices are exogenous, the intermediary diverts consumers with positive probability toward the firm from which it receives the higher fee.[13] By contrast, in our model, the firm receives profits from sales rather than from "clicks" and earns the same profit on all sales; furthermore, it does not observe consumers' types and cannot force consumers to inspect products in any particular order.

Our paper is also related to the small literature on sponsored-link advertising (or position) auctions and consumer search, most notably Chen and He (2011).[14] In their model, consumers are willing to inspect a given number of distinguishable slots on a platform and single-product sellers, with products that differ in popularity, participate in auctions for these slots. There exists a separating equilibrium (akin to our pure positioning equilibrium), in which the seller offering the most popular product wins the auction for the advertising slot that consumers visit first, as well as a pooling equilibrium (akin to our random positioning equilibrium), in which sellers bid zero

[11] One exception is Rhodes (2015) who studies the pricing and advertising decisions of a multi-product retailer offering n symmetric products. Contributing to the literature on price (rather than match) search with homogeneous goods, pioneered by Stahl (1989), Hämäläinen (2022) allows not only for inter-firm search but also for intra-firm search.

[12] Gamp (2019) analyzes a setup that is, in some ways, more general than Petrikaitė (2018)'s but assumes that prices are observable.

[13] If prices are endogenous, then some consumer types are diverted even if the intermediary receives the same fee from both sellers. More recently, Teh and Wright (2022) study steering of consumers by a monopoly intermediary on a platform where firms pay commission fees to influence the intermediary's ranking of products.

[14] For a related paper on auction design, see Athey and Ellison (2011).

and search is random. In our model, search cost heterogeneity generates endogenous intensive and extensive search margins. Moreover, the platform is vertically integrated and does not commit to any (given) mechanism when allocating products to slots.

Choice overload—that consumers may sometimes be less likely to choose any product if they face more options[15]—has been documented in several field experiments. For instance, besides the above-mentioned study by Iyengar and Leppar (2000), Boatwright and Nunes (2001) report that, in an experiment run by an online grocery, halving product choice resulted in 11 percent higher sales.[16]

Kamenica (2008) shows how choice overload can arise when consumers make inferences from the size of the product line. Products differ in how likely they appeal to consumers, as in our model; however some consumers know which product they like, whereas others do not, and must therefore choose randomly. As the firm always selects the most popular products, expanding the product line gives more choice to informed consumers, but lowers the *average popularity* of the selected products, which reduces the demand from uninformed consumers.

Kuksov and Villas-Boas (2010) and Villas-Boas (2009) develop alternative models of contextual inference, in which products are instead horizontally differentiated. Kuksov and Villas-Boas (2010) allow consumers to choose between picking a product at random or engaging in costly sequential search. Expanding the product line leads the firm to offer more *niche* products (i.e., products that are particularly attractive for specific consumers, but not so attractive for most of them), which reduces again the average popularity but also affects the *variance* of the utility offered by the selected products, thereby encouraging consumers to search rather than to pick at random. Villas-Boas (2009) endogenizes prices in a monopoly setting in which consumers can, by incurring a one-time search cost, learn the prices and locations of all of the products on offer, as well as their own locations. Expanding the product line helps consumers find a better fit, but also allows the firm to appropriate the associated benefits by increasing its prices. Anticipating this, consumers are less inclined to search.

Our paper builds on these earlier contributions. As in Kamenica (2008), choice overload arises because expanding the product line reduces the average popularity of products. Furthermore, consumers can engage in costly sequential search, as in Kuksov and Villas-Boas (2010), and prices are endogenous, as in Villas-Boas (2009). The novelty resides in the focus on directed search and steering incentives. Our analysis reveals that the attractiveness of a product line to consumers depends not only on its

[15] A related notion of choice overload focuses instead on (self-reported) satisfaction with the chosen option; see, for example, Scheibehenne et al. (2010).

[16] While similar findings have been reported in other contexts, the meta-analytic review by Scheibehenne et al. (2010) paints a more nuanced and mixed picture. See also the discussion in Section 3 below.

size but also on the amount of noise in its positioning.

2 Setting

We consider an industry with N available products, indexed by $i \in \mathcal{I}_N \equiv \{1, \ldots, N\}$.

Supply. A monopolist chooses which of these products to offer, and at what prices.[17] There is no fixed cost associated with offering a product, and all products involve the same constant unit cost of production, which is normalized to zero.

Demand. A unit mass of consumers, each with a unit demand, differ in search costs and valuations. Each consumer's search cost, c, is independently drawn over $\mathbb{R}_+$ from a distribution with c.d.f. $G(\cdot)$.[18] As in Chen and He (2011), each product i is characterized by a probability $\mu_i \in (0,1)$ with which any consumer has a match with that product, where $1 > \mu_i > \cdots > \mu_N > 0$; that is, products are labeled in descending order of popularity and all match probabilities lie strictly between 0 and 1. A consumer values a product only if he has a match with that product, in which case the consumer's valuation, v, is drawn over $\mathbb{R}_+$ from a distribution with c.d.f. $F(\cdot)$. Matches and match-conditional valuations are i.i.d. across consumers but, conditional on having a match, a consumer's valuation is the same for all products.[19]

Sequential search. Consumers first learn their search costs and observe the number of products offered by the firm (but not their identity). They then decide whether or not to start a search. If a consumer chooses not to search, his payoff is zero; otherwise, he pays his search cost c and inspects one of the products. Upon inspection, he observes the price of the product and learns whether or not he has a match with that product. If he does, he learns his match-conditional valuation v and then decides whether or not to purchase the product, in which case he stops searching. In the absence of a purchase, the consumer decides whether to inspect another product, thereby incurring again the search cost c, and so on. That is, we assume that there is a constant (per-product) search cost.[20] There is perfect recall: the consumer has the option of purchasing any previously inspected product. Importantly, consumers never learn the identity of inspected products; instead, they update their beliefs on the basis of products' prices, whether or not they had a match and, if so, their match-conditional

[17] Throughout the paper, we assume that the price of a product is the same for all consumers. In particular, the monopolist cannot condition prices on consumers' search histories or other characteristics.

[18] See Section 5.1 for a discussion of the case of homogeneous search costs.

[19] The assumption that match-conditional valuations are the same across products follows Chen and He (2011) and is made for tractability. In particular, as we will see it ensures that the firm charges a uniform price—see Lemma 1—and that consumers stop searching upon finding a match. It also ensures that the monopoly price is the same for all products, which avoids signalling issues.

[20] See Section 5.3 for a discussion of the case of non-constant search costs.

valuations.

Positioning. Finally, the firm can influence consumers' search by placing its products in specific spots. For example, supermarkets place their products in given aisles and shelves, and online platforms can make recommendations. Formally, we assume that there are as many distinguishable *slots* as products, and that the firm assigns (possibly randomly) products to slots. Active consumers must in turn choose not only how many slots to inspect, but also which ones and in which order.[21]

Formally, for any product portfolio $\mathcal{I} \subseteq \mathcal{I}_N$, let $\mathcal{S}(\mathcal{I})$ denote the set of sequences in $\mathcal{I}$; a slot assignment (or product placement) is of the form $\boldsymbol{\sigma} \equiv (\sigma_1, \ldots, \sigma_{|\mathcal{I}|}) \in \mathcal{S}(\mathcal{I})$, meaning that slot k holds product $\sigma_k \in \mathcal{I}$.

Timing. The timing is as follows:

Stage 1 The firm chooses the size $n \in \mathcal{I}_N$ of its product line, which is publicly observed.

Stage 2 The firm privately observes the identity of the products, chooses its product portfolio (i.e., it chooses $\mathcal{I}$ subject to $|\mathcal{I}| = n$), positions and prices the selected products; consumers then observe their search costs and sequentially decide which slots to inspect (if any).

Specifically, in stage 2 the firm can choose deterministic or random positioning and pricing strategies. As for consumers' search strategies, they consist in deciding—for any history of inspected products, realized matches, observed prices and observed valuation (in case of a previous match)—whether to inspect an additional slot and, if so, which one.

Equilibrium. We focus on Perfect Bayesian Equilibria (PBE). For any size of the product line chosen in stage 1, stage 2 forms a proper subgame of incomplete information. Hence, the continuation equilibrium strategies of the firm (product selection, positioning and pricing decisions) and of the consumers (search and purchasing decisions) must constitute a PBE of the subgame. In this subgame, consumers never observe the composition of the product portfolio nor the positions of the selected products but, upon inspection, observe the prices charged by the firm, whether there is a match, and moreover learn their valuation upon the first match.

As long as they observe prices that are consistent with the firm's equilibrium strategy,[22] consumers update their beliefs according to Bayes' rule, using all relevant information. In particular, whether a match occurs may convey information about the

[21] For the most part, we assume that the firm cannot condition its positioning strategy on consumer characteristics. However, we study in Section 5.1 the case where positioning can be tailored to consumers' search costs.

[22] As $0 < \mu_i < 1$ for every product i, any observed match sequence is consistent with any strategy of the firm.

popularity of the inspected product. The observed prices could also be informative if they were to differ across products. By contrast, as a consumer's match valuation is the same for all products, the realization of this valuation does not add any information to that conveyed by the sequence of matches.

When instead consumers observe unexpected prices, Bayes' rule has no bite but consumers need to form beliefs about the implications of these prices for the selection, positioning and pricing of the products in yet-uninspected slots. This issue does not arise in traditional between-firm search models, where the price of one firm does not convey any information about other firms' prices, or in the few papers on within-firm search, which focus on non-sequential search (that is, incurring the search cost provides information about all of the firm's offerings).[23] A similar issue arises however in the context of vertical relations, where upon receiving an unexpected offer from a supplier, a downstream firm must form beliefs about the implications for the offers made to its rivals. We will adopt the assumption of *passive beliefs* commonly made in such settings.[24] This means here that, when encountering one or more out-of-equilibrium prices during their search process, consumers stick to the belief that the firm has selected the subgame equilibrium product portfolio and set the subgame equilibrium prices for all the not-yet-inspected products.[25]

From now on, "equilibrium" thus means PBE with passive beliefs. To streamline the exposition, we adopt the tie-breaking convention that, whenever consumers are indifferent between alternative decisions, they select the option most favorable to the firm. Also, whenever consumers inspect slots in a specific order (as is the case in all the equilibria considered below, except the fully random positioning ones studied in Section 3.2), without loss of generality we label the slots according to the search sequence—i.e., "slot k" refers to the k^{th} inspected slot.

Finally, let

$$\pi(p) \equiv p[1 - F(p)] \quad \text{and} \quad s(p) \equiv \int_p^\infty (v - p) dF(v)$$

denote the expected profit and consumer surplus, respectively, generated by a match with a product priced at p. We assume that $\pi(p)$ has a unique maximizer, denoted

[23] See, e.g., Zhou (2014), Rhodes (2015), or Rhodes, Watanabe, and Zhou (2021).

[24] Introduced by Hart and Tirole (1990), this assumption is also consistent with the *contract equilibrium* or "Nash-in-Nash" approach pioneered by Crémer and Riordan (1987) and Horn and Wolinsky (1988), and adopted by most theoretical and empirical papers on vertical relations.

[25] McAfee and Schwartz (1994) proposed the notion of wary beliefs, according to which consumers expect the firm to have chosen the optimal product portfolio and prices, given the observed out-of-equilibrium price(s). Unfortunately, this avenue is hardly tractable, even in simple sender-receiver settings—in particular, the firm's optimal strategy depends on consumers' interpretations of unexpected prices, which in turn requires a fixed-point approach.

p^m, and let $\pi^m \equiv \pi(p^m)$ and $s^m \equiv s(p^m)$.

3 Benchmarks

We first characterize two polar types of equilibria, in which the firm positions its products either deterministically (pure positioning) or fully randomly (random positioning hereafter).

3.1 Pure positioning

We start with deterministic (i.e., pure-strategy) equilibria. We first show that the firm necessarily charges monopoly prices:

Lemma 1 (monopoly pricing) *For any given product portfolio size $n \in \mathcal{I}_N$ adopted in the first stage and any pure-strategy equilibrium of the continuation subgame, the firm offers all selected products at the monopoly price p^m.*

Proof. See Appendix A. ■

By construction, if the firm offers all selected products at the same price, it cannot do better than charging the monopoly price. Lemma 1 shows further that the firm cannot gain from charging different prices across products. This is because consumers' response works against the interest of the firm and in particular prevents any profitable discrimination, as consumers with greater overall values (i.e., lower costs or higher valuations) are more prone to keep searching, and thus likely to find lower prices. To see this, consider the subgame for $n = 2$, a candidate continuation equilibrium in which the firm offers products i and j at prices p_i and $p_j < p_i$, and consider those consumers who have a match (and thus learn their valuations) on the first inspection. Obviously, consumers who first inspect the lower-priced product, j, stop searching, regardless of their cost or valuations. For those who inspect product i first, the expected benefit of inspecting the other product is equal to $-c + \mu_j [\max\{v - p_j, 0\} - \max\{v - p_i, 0\}]$, which is decreasing in c and weakly increasing in v, and strictly so if $v \in (p_i, p_j)$. As a result consumers with lower costs or higher valuations are more likely to find the lower price p_j.[26]

[26]Price discrimination could, however, be profitable if valuations were positively correlated with search costs, as consumers with higher valuations would then be less likely to search intensively. Furthermore, if consumers' match-conditional valuations were independent across products, then the "monopoly price" would differ across consumers, as those with lower search costs are then more likely to search more intensively and draw higher valuations.

Building on this first insight, we now show that, for any size of the product line, there exists a unique pure positioning equilibrium in the continuation subgame. For any $i \in \mathcal{I}_N$, let (with P for *Pure positioning*)

$$c_i^P \equiv \mu_i s^m \text{ and } \lambda_i \equiv 1 - \mu_i$$

respectively denote the search cost threshold below which consumers would be willing to inspect product i and the probability that this product does not deliver a match; also, for any $\mathcal{I} \subseteq \mathcal{I}_N$, let

$$\Lambda\left(\mathcal{I}\right) \equiv \prod_{i \in \mathcal{I}} \lambda_i$$

denote the probability that all products in $\mathcal{I}$ fail to deliver a match (with the convention $\Lambda\left(\varnothing\right) = 1$). We have:

Proposition 1 (pure positioning) *For any given product portfolio size $n \in \mathcal{I}_N$ adopted in the first stage, there exists a unique pure-strategy equilibrium in the continuation subgame. In this equilibrium, the firm selects the n most popular products (i.e., $\mathcal{I} = \mathcal{I}_n$) and offers them at the monopoly price (i.e., $p_i = p^m$ for every $i \in \mathcal{I}_n$), consumers with a search cost $c \leq c_k^P$ inspect the k most popular products by decreasing order of popularity until finding a match, and the firm's profit is (with the convention $\mathcal{I}_0 = \varnothing$):*

$$\Pi^P\left(n\right) \equiv \sum_{k=1}^{n} \Lambda\left(\mathcal{I}_{k-1}\right) G\left(c_k^P\right) \mu_k \pi^m.$$

Proof. See Appendix B. ■

From Lemma 1, in any pure positioning equilibrium the firm offers the selected products at the same price p^m. It follows that consumers' interests are somewhat aligned with that of the firm, in that they all seek to maximize the probability of a match. In particular, consumers inspect first the slots that are expected to hold the most popular of the selected products, and some consumers who start searching stop doing so even in the absence of a match; specifically, consumers with a search cost $c \in \left(c_k^P, c_{k+1}^P\right)$ inspect the k most popular products until finding a match, but stop searching after that. This, in turn, induces the firm not only to offer the most popular products (i.e., $\mathcal{I} = \mathcal{I}_n$), but also to position them by popularity.

Remark 1 (observable product portfolio) *If consumers observed the composition of the product line, then offering the most popular products would maximize not only the probability of a match but also the number of searchers. Hence, Proposition 1 would a fortiori hold; a similar comment applies to the analysis below.*

It follows from Proposition 1 that the overall game also has a unique pure-strategy equilibrium:

Corollary 1 (efficiency) *There exists a unique pure-strategy equilibrium, in which the firm offers all available products and active consumers inspect them by decreasing order of popularity. This equilibrium therefore offers maximal choice to consumers and enables them to make fully informed decisions; it thus induces efficient search patterns and, conditional on monopoly pricing, maximizes consumers' expected surplus.*

Proof. See Appendix C. ■

Hence, there always exists an equilibrium in which the firm adopts a pure strategy and offers all products at the monopoly price. Furthermore, as this equilibrium is perfectly informative, it enables consumers to inspect a given slot when and only when it is efficient for them to do so—hence, conditional on monopoly pricing, this equilibrium maximizes their expected surplus.

Despite these efficiency benefits, the firm may find it profitable to introduce noise in the positioning of its products, which limits the information available to consumers and encourages them to search more intensively.

3.2 Random positioning

Obviously, if in stage 1 the firm chooses to offer a single product (i.e., $n = 1$), the equilibrium described by Proposition 1, in which the firm offers the most popular product at the monopoly price, constitutes the unique continuation equilibrium. We now show that, for any given $n \in \mathcal{I}_N \diagdown \{1\}$, there also exist *random positioning* equilibria in which the firm sticks to monopoly pricing,[27] but uniformly randomizes over its positioning strategy—that is, it assigns each selected product to every slot with equal probability.[28]

We first note that the absence of positioning induces an *intensive search* pattern:

Lemma 2 (intensive search) *For any given product portfolio size $n \in \mathcal{I}_N \diagdown \{1\}$ adopted in the first stage, and any continuation equilibrium featuring monopoly pricing and random positioning, any consumer who starts searching keeps searching until finding a match.*

[27]In the absence of positioning, Nocke and Rey (2021) show that, as in the case of pure positioning, consumers with higher valuations tend to inspect more products and are thus more likely to encounter lower prices, which again rules out any profitable price discrimination; charging a uniform price equal to p^m thus maximizes the profit generated by active consumers. Product-specific prices can however allow consumers to identify inspected products, which in turn may encourage more consumers to become active. See also Remark 4 and the discussion of disclosure policies in Section 5.4.

[28]Specifically, each of the $n!$ possible slot assignments $\boldsymbol{\sigma} \in \mathcal{S}(\mathcal{I})$ is selected with probability $1/(n!)$.

Proof. See Appendix D. ■

When an inspection does not produce a match, consumers assign a higher probability to the inspected product being one of the less popular ones; they thus become more *optimistic* about the remaining products, which in turn encourages them to keep searching until finding a match.

Remark 2 (search addiction) *That consumers become more optimistic in the absence of a match gives rise to a search addiction pattern: consumers who do not want to start searching may choose to keep searching if coerced to do a first search, the result of which turns out to be unsuccessful.*

Building on Lemma 2, we now show that there exists an essentially unique equilibrium featuring fully random positioning and monopoly pricing, in which the firm offers the n most popular products. For any product portfolio $\mathcal{I}$, let

$$\mathrm{M}(\mathcal{I}) \equiv 1 - \Lambda(\mathcal{I}) \tag{1}$$

denote the probability of having at least one match in that portfolio, and

$$\Gamma(\mathcal{I}) \equiv \frac{1}{|\mathcal{S}(\mathcal{I})|} \sum_{\sigma \in \mathcal{S}(\mathcal{I})} \left(1 + \sum_{k=1}^{n-1} \prod_{i=1}^{k} \lambda_{\sigma_i}\right) \tag{2}$$

denote the expected number of inspections. Finally, let (with R for *Random positioning*)

$$\mathrm{M}_n^R \equiv \mathrm{M}(\mathcal{I}_n) \text{ and } c_n^R \equiv \frac{\mathrm{M}(\mathcal{I}_n)}{\Gamma(\mathcal{I}_n)} s^m \tag{3}$$

respectively denote the probability of a match with the n most popular products, and the search cost threshold below which consumers would be willing to inspect these products. We have:[29]

Proposition 2 (random positioning) *For any given product portfolio size $n \in \mathcal{I}_N \setminus \{1\}$ adopted in the first stage, there exist continuation equilibria featuring monopoly pricing and random positioning. These equilibria only differ in the sequence with which consumers inspect slots (any random or deterministic sequence being admissible): in all of them, the firm offers the n most popular products, consumers with search cost below c_n^R inspect slots until finding a match, and the firm's profit is*

$$\Pi^R(n) \equiv G\left(c_n^R\right) \mathrm{M}_n^R \pi^m.$$

[29]The equilibria identified by Proposition 2 are analogous to the babbling equilibria of cheap talk games, in which the sender's "messages" (here, for example, the firm's recommendations) provide no information and is thus discarded by the receiver (here, the consumers).

Proof. See Appendix E. ■

The comparison between the equilibria characterized by Propositions 1 and 2 highlights a trade-off between *extensive* and *intensive* search margins: pure positioning encourages more consumers to start searching (extensive margin), by providing them with better information, whereas random positioning encourages active consumers to keep searching until finding a match (intensive margin), by making them more optimistic in the absence of a match.[30]

It follows that, depending on the match probabilities and the distribution of search costs, either type of equilibrium can be more profitable.[31]

Remark 3 (indistinguishable slots) *The random positioning equilibrium in which consumers also uniformly randomize over their search sequences replicates the unique "belief-proof" equilibrium that would arise if consumers could not distinguish between slots.*[32]

Remark 4 (uniform pricing) *The same equilibrium is also the only one featuring uniform pricing (i.e., all selected products are offered at the same price) as well as uniformly random positioning and search. However, there may exist equilibria with non-uniform pricing. In the Online Appendix, we provide an example in which (i) the firm offers products* 1 *and* 2 *at prices $p_1 = p^m$ and $p_2 < p^m$, uniformly randomizing over their positions, and (ii) active consumers uniformly randomize over their search sequence but, being able to so identify inspected products through their prices, stop searching when encountering the more popular product on the first inspection. As this feature encourages more consumers to start searching, this equilibrium can be more profitable than the equilibrium with random positioning and uniform monopoly pricing.*

It can be checked that the threshold c_n^R identified by Proposition 2 is decreasing in n, implying that the absence of positioning generates a form of *choice overload*:

---

[30]Compared with $\Pi^R(n) = G\left(c_n^R\right) \mathrm{M}_n^R \pi^m$, the profit generated by pure positioning can be expressed as $\Pi^P(n) = G\left(c_1^P\right) \mathrm{M}_n^P \pi^m$, where $G\left(c_1^P\right)$ is the number of active searchers and $\mathrm{M}_n^P \equiv \sum_{k=1}^{n} \Lambda\left(\mathcal{I}_{k-1}\right) G\left(c_k^P\right) \mu_k / G\left(c_1^P\right)$ is the average probability of a match among active searchers; for any $n \in \mathcal{I}_N$, $c_n^R < c_1^P$ but $\mathrm{M}_n^R > \mathrm{M}_n^P$.

[31]For instance, holding fixed the popularity of the various products, pure positioning is more profitable whenever $G\left(c_n^R\right)$ is small enough, whereas random positioning is instead more profitable when $G\left(c_n^R\right)$ is sufficiently close to $G\left(c_1^P\right)$.

[32]"Belief-proof equilibria" refers here to equilibria that survive replacing consumers off-equilibrium beliefs and behavior with any consistent beliefs and any behavior induced by these beliefs. See Nocke and Rey (2021) for a formal treatment.

Corollary 2 (choice overload) *If a random positioning equilibrium is played in all continuation subgames, then expanding the product line reduces the number of active consumers.*

Proof. See Appendix F. ■

Expanding the product line gives rise to two opposing effects on consumers' incentives to search: on the one hand, it increases the probability of finding a match ($M(\mathcal{I}_n)$); on the other, by decreasing the average popularity of products offered, it increases the expected number of inspections ($\Gamma(\mathcal{I}_n)$). While the first effect outweighs the second for consumers with sufficiently low search costs, Corollary 2 reveals that the opposite is true for the marginal searcher c_n^R.

To see why, note first that the cost of the marginal consumer (weakly) *exceeds* the average popularity of the selected products.[33] It follows that the consumer who is barely willing to inspect an assortment of the n most popular products would not inspect product $n+1$ if he knew the position of that product. As random positioning further reduces consumers' expected value from search, it follows that expanding the product line discourages consumer participation.

Under pure positioning, expanding the product line is always profitable. By contrast, in the absence of positioning, such expansion triggers another trade-off between the extensive and intensive search margins: it discourages some consumers from searching ($G(c_n^R)$ is decreasing in n) but induces active consumers to inspect more products and thus increases their probability of finding a match (M_n^R increases with n). Interestingly, in her recent study of an online delivery platform, Natan (2022) finds that expanding the assortment size indeed reduces the number of active consumers (extensive margin) but tends to increase the duration of their searches (intensive margin).[34]

Because of this trade-off, an increase in the size of the product line may thus decrease or increase the firm's sales. This is consistent with the experimental evidence on choice overload: while some studies (e.g., Iyengar and Leppar, 2000; Boatwright and Nunes, 2001) find a negative effect of assortment size on sales, others (e.g., Kahn and Wansink, 2004) find the opposite. In their meta-study, Scheibehenne et al. (2010) conclude: "The overall mean effect size [...] in our meta-analysis was virtually zero." Our analysis can help predicting when the effect of assortment size on sales is likely to be negative—namely, when products are randomly placed, and the distribution of

[33] Indeed, the overall (net) value from search can be expressed as the sum of the value of the first inspection and the continuation value from subsequent ones; as the latter cannot be negative, it follows that, for the marginal consumer (whose overall value is zero), the value of the first inspection cannot be positive.

[34] See columns (1) and (2) in Tables 5 and 12 for consumer participation, and columns (3) and (4) in Table OA11 for search intensity.

cognitive costs (c) exhibits a large hazard rate or the assortment size is already large.[35]

4 Steering

The equilibria identified so far resolve the trade-off between extensive and intensive search margins in arguably extreme ways: pure positioning maximizes consumer participation (as no consumer with a cost $c > c_1^P$ would ever inspect any product, even the most popular one), whereas random positioning maximizes participating consumers' probability of a match (as they keep searching until finding a match). We now explore alternative resolutions of this trade-off.

The above equilibria exhibit *uniform pricing*; prices therefore do not affect consumers' search patterns. Furthermore, with passive beliefs, a deviation from a uniform equilibrium price $p \neq p^m$ to p^m would also have no impact on consumers' search behavior, and increase the expected profit generated by a match; hence, the firm optimally charges the monopoly price. Yet, the firm can steer consumers' search through its positioning strategy. Specifically, pure positioning induces consumers to start with the most popular products. Under random positioning, they are indifferent among all search sequences, and thus willing to inspect slots in a specific order; this indifference moreover relies critically on positioning being uniformly random: any other (e.g., partly deterministic and/or non-uniformly random) positioning strategy is likely to induce consumers to start again with the most promising slots. In what follows, we therefore focus on *steering equilibria* where prices are uniform (and thus equal to the monopoly price) but the positioning strategy induces consumers to inspect slots in a specific order.

4.1 Preliminaries

We first show that any such equilibrium has a *block* structure of the form

$$\boldsymbol{\mathcal{B}} = (\mathcal{B}_1, \dots, \mathcal{B}_{|\boldsymbol{\mathcal{B}}|}),$$

where $\bar{\mathcal{B}}_k \equiv \cup_{h=1}^k \mathcal{B}_h$ is the set of the $\sum_{h=1}^k |\mathcal{B}_h|$ most popular products. We have:

Lemma 3 (steering) *For any given size $n \in \mathcal{I}_N$ adopted in the first stage and any steering continuation equilibrium, there exists $\boldsymbol{\mathcal{B}} = (\mathcal{B}_1, \dots, \mathcal{B}_{|\boldsymbol{\mathcal{B}}|})$ and $\hat{\mathbf{c}} = (\hat{c}_1, \dots, \hat{c}_{|\boldsymbol{\mathcal{B}}|}) \in \mathbb{R}_+^{|\boldsymbol{\mathcal{B}}|}$ satisfying $\hat{c}_1 > \dots > \hat{c}_{|\boldsymbol{\mathcal{B}}|}$ such that, for any $k \in \mathcal{I}_{|\boldsymbol{\mathcal{B}}|}$:*

[35] Abstracting from integer constraints, the impact of an increase in the size of the product line on sales is given by $d\Pi^R \simeq \left[d\mathrm{M}(\mathcal{I}_n)/dn - \gamma\left(c_n^R\right)\mathrm{M}(\mathcal{I}_n)\right] G\left(c_n^R\right) dn$, where $\gamma(c) \equiv g(c)/G(c)$ denotes the reverse hazard rate of the search cost distribution and $d\mathrm{M}(\mathcal{I}_n)/dn$ is positive but tends to 0 (as $\mathrm{M}(\cdot)$ is bounded by 1) as n tends to ∞.

- *the firm assigns the products from $\mathcal{B}_k$ to the slots $\mathcal{B}_k$, and does so randomly whenever $|\mathcal{B}_k| > 1$;*

- *active consumers inspect slots in increasing order and, in the absence of a match on previous inspections, those with a search cost $c \leq \hat{c}_k$ inspect all the slots $\mathcal{B}_k$ until finding a match.*

Proof. See Appendix G. ■

The firm thus assigns blocks of similarly popular products to designated blocks of slots—pure positioning corresponds to $\boldsymbol{\mathcal{B}} = (\{i\})_{i \in \mathcal{I}_n}$, and random positioning to $\boldsymbol{\mathcal{B}} = (\mathcal{I}_n)$. In what follows, $\mathcal{B}_k$ will refer interchangeably to the corresponding slots or products.

The block boundaries are driven by consumers' search patterns. Specifically, active consumers all inspect the slots in the same order and stop searching as soon as they obtain a match, as all products offer the same match-value and are offered at the same price. In the absence of a match, however, consumers with higher search costs may stop searching before those with lower costs; the blocks are then defined so that no consumer stops searching within a block, whereas some stop searching between blocks. That is, all consumers inspecting the first slot in $\mathcal{B}_k$ keep searching within $\mathcal{B}_k$, but some of them stop afterwards and never visit the next blocks. This, in turn, incentivizes the firm to assign the $|\mathcal{B}_1|$ most popular products to the first block, $\mathcal{B}_1$, the next $|\mathcal{B}_2|$ most popular products to the second block, $\mathcal{B}_2$, and so forth. For instance, if some consumers only inspect slot 1, then $\mathcal{B}_1 = \{1\}$. If instead all active consumers are willing to inspect slot 2 as well, and some stop after that, then $\mathcal{B}_1 = \{1, 2\}$; and so forth.

Finally, whenever a block holds more than one product, ensuring that consumers who start inspecting it keep doing so until finding a match requires the firm to randomize over its positioning strategy within that block; otherwise some consumers, having inspected the most popular products of the block, would then stop searching even in the absence of a match. That is, the firm must introduce noise in its positioning within each such block.

4.2 Intensive search

From Lemma 3, any steering equilibrium is characterized by a block structure $\boldsymbol{\mathcal{B}}$ and search cost thresholds $\hat{\mathbf{c}} = \left(\hat{c}_1, \ldots, \hat{c}_{|\boldsymbol{\mathcal{B}}|}\right)$. The resulting profit can be expressed as

$$\Pi\left(\boldsymbol{\mathcal{B}}, \hat{\mathbf{c}}\right) \equiv \sum_{k=1}^{|\boldsymbol{\mathcal{B}}|} \Lambda\left(\bar{\mathcal{B}}_{k-1}\right) G\left(\hat{c}_k\right) \mathrm{M}\left(\mathcal{B}_k\right) \pi^m, \tag{4}$$

By revealed preferences, an increase in $\hat{c}_k$ means that more consumers are interested in block $\mathcal{B}_k$; furthermore, as $\Pi(\boldsymbol{\mathcal{B}}, \hat{\mathbf{c}})$ increases with every $\hat{c}_k$, fostering participation benefits the firm as well. It follows that, keeping the block structure constant, increasing participation enhances the firm's profit as well as consumer surplus: for any two equilibria with participation thresholds $\hat{\mathbf{c}}$ and $\hat{\mathbf{c}}'$, the latter Pareto-dominates the former whenever $\hat{\mathbf{c}}' > \hat{\mathbf{c}}$.

To gain further insights, we consider a given block $\mathcal{B} = \left\{i_1, \ldots, i_{|\mathcal{B}|}\right\}$, where $\mu_{i_1} > \cdots > \mu_{i_{|\mathcal{B}|}}$, and study a variant of our model in which these products are the only available ones, and the firm offers all of them (i.e., it chose $n = |\mathcal{B}|$ in stage 1). We will say that a steering equilibrium features *intensive search* over $\mathcal{B}$ if active consumers inspect *all* slots in $\mathcal{B}$ until finding a match.

We have seen that *random positioning* sustains such an equilibrium. As it provides no information to consumers, it actually constitutes the *worst* one—for the firm as well as for consumers:

Lemma 4 (worst intensive search) *Fix $\mathcal{B} \subseteq \mathcal{I}_N$ such that $|\mathcal{B}| > 1$. Among the steering equilibria featuring intensive search over $\mathcal{B}$, the random positioning equilibrium minimizes consumer participation.*

Proof. See Appendix H. ■

The proof is straightforward: it suffices to note that consumers can always secure the payoff achievable under random positioning by simply adopting a uniformly random search strategy. It follows that, if consumers prefer to direct their search, this must increase their expected surplus, which boosts participation and the firm's profit.

Characterizing the best equilibrium featuring intensive search is substantially more tricky. As the number of slot sequences increases exponentially with the size of $\mathcal{B}$, the firm must choose a large number of positioning probabilities and faces an even larger number of incentive constraints.[36] To make progress, it is useful to start with a simpler problem, in which the firm could *control* the order in which slots can be inspected, and must only induce them to keep searching until finding a match. Let

$$\bar{\lambda}(\mathcal{B}) \equiv \sqrt[|\mathcal{B}|]{\Lambda(\mathcal{B})}$$

denote the geometric mean of the no-match probabilities of the products in $\mathcal{B}$ and define

$$\hat{c}(\mathcal{B}) \equiv \left[1 - \bar{\lambda}(\mathcal{B})\right] s^m \text{ and } \hat{\Pi}(\mathcal{B}) \equiv G\left(\left[1 - \bar{\lambda}(\mathcal{B})\right] s^m\right) \mathrm{M}(\mathcal{B}) \pi^m.$$

[36]Specifically, if $\mathcal{B}$ is of size n, there are $n! - 1$ independent slot assignment probabilities (as each product can be assigned to any slot) and, accounting for the possibility that consumers may not inspect all slots, there are $\sum_{k=1}^{n} \frac{n!}{(n-k)!} - 1 = n! - 1 + \sum_{k=1}^{n-1} \frac{n!}{(n-k)!}$ incentive constraints.

We have:

Proposition 3 (controlled search) *Fix $\mathcal{B} \subseteq \mathcal{I}_N$ such that $|\mathcal{B}| > 1$, and suppose that slots can only be inspected in a given order (consumers remaining free to start or stop searching). Among the equilibria featuring intensive search over $\mathcal{B}$, the Pareto-efficient ones yield a participation threshold and profit of $\hat{c}(\mathcal{B})$ and $\hat{\Pi}(\mathcal{B})$, respectively.*

Proof. See Appendix I. ■

Let β_k denote the probability that the first k slots deliver no match, for $k \in \mathcal{I}_{|\mathcal{B}|}$. Inspecting up to k slots thus delivers a match with probability $1 - \beta_k$ and generates an expected number of inspections equal to $1 + \beta_1 + \cdots + \beta_{k-1}$.[37] The value of doing so can therefore be expressed as (with the convention $\beta_0 = 1$):

$$V_k(c) = (1 - \beta_k) s^m - \left(\sum_{i=1}^{k} \beta_{i-1} \right) c, \tag{5}$$

which is decreasing in c. Consumer participation is thus characterized by a threshold $\hat{c}$, which is the largest cost satisfying the participation constraint $V_{|\mathcal{B}|}(c) \geq 0$, or, noting that, by construction, $\beta_{|\mathcal{B}|} = \Lambda(\mathcal{B})$:

$$[1 - \Lambda(\mathcal{B})] s^m - \left(\sum_{i=0}^{|\mathcal{B}|-1} \beta_i \right) \hat{c} \geq 0. \tag{IR}$$

Ensuring that the marginal consumer favors inspecting all slots rather than only $k \in \mathcal{I}_{|\mathcal{B}|-1}$ of them yields the incentive constraints $V_{|\mathcal{B}|}(\hat{c}) \geq V_k(\hat{c})$, or:[38]

$$[\beta_k - \Lambda(\mathcal{B})] s^m \geq \left(\sum_{i=k}^{|\mathcal{B}|-1} \beta_i \right) \hat{c}. \tag{IC_k}$$

Suppose now that the firm can choose at will $(\beta_k)_{k \in \mathcal{I}_{|\mathcal{B}-1|}}$. Maximizing participation calls for minimizing these no-match probabilities, and so the $|\mathcal{B}| - 1$ incentive constraints are all binding: that is, the marginal consumer is indifferent not only between starting or not a search, but also about how many slots to inspect. This pins down a unique solution, $\beta_k = \hat{\beta}_k \equiv \left[\bar{\lambda}(\mathcal{B})\right]^k$ for $k \in \mathcal{I}_{|\mathcal{B}-1|}$, and provides an upper bound on participation, $\hat{c} = \hat{c}(\mathcal{B})$.

[37]The consumer inspects the first slot with probability 1, an additional one with probability β_1, and so on, and the last one with probability β_{k-1}.

[38]This condition, in turn, ensures that infra-marginal consumers are also willing to inspect all slots.

The last part of the proof is constructive, and consists in exhibiting a positioning strategy $\hat{\boldsymbol{\rho}}(\mathcal{B})$ that implements the desired no-match probabilities $(\hat{\beta}_k)_{k\in\mathcal{I}_{|\mathcal{B}-1|}}$. Interestingly, it is actually possible to find positioning strategies that not only achieve that, but moreover induce consumers to inspect slots in the desired order, even if they do not have to. Specifically, consider the positioning strategy $\hat{\boldsymbol{\rho}}(\mathcal{B})$ that selects each slot assignment $\boldsymbol{\sigma} \in S(\mathcal{B})$ with probability $\hat{\rho}_{\boldsymbol{\sigma}}(\mathcal{B}) \equiv \omega_{\boldsymbol{\sigma}}(\mathcal{B}) / \sum_{\tilde{\boldsymbol{\sigma}}\in S(\mathcal{B})} \omega_{\tilde{\boldsymbol{\sigma}}}(\mathcal{B})$, where:

$$\omega_{\boldsymbol{\sigma}}(\mathcal{B}) \equiv \prod_{i\in I_{|\mathcal{B}|}} \lambda_{\sigma_i}^{\frac{i-1}{|\mathcal{B}|}}. \tag{6}$$

We have:

Proposition 4 (best intensive search) *Fix $\mathcal{B} \subseteq \mathcal{I}_N$ such that $|\mathcal{B}| > 1$. There exists a steering equilibrium featuring intensive search over $\mathcal{B}$ in which:*

(i) the firm adopts the positioning strategy $\hat{\boldsymbol{\rho}}(\mathcal{B})$,

(ii) the participation threshold and profit are $\hat{c}(\mathcal{B})$ and $\hat{\Pi}(\mathcal{B})$, respectively.

Proof. See Appendix J. ∎

The strategy $\hat{\boldsymbol{\rho}}(\mathcal{B})$ assigns each product in $\mathcal{B}$ to every designated slot with positive probability. It however puts more weight on slot assignments that position more popular products in earlier slots. For instance, if $\mathcal{B} = \{1,2\}$, then product 1 is more likely to be assigned to the first slot: $\hat{\boldsymbol{\rho}}_{(1,2)}(\mathcal{B}) = \sqrt{\lambda_2}/\left(\sqrt{\lambda_1}+\sqrt{\lambda_2}\right) > 1/2$.

The precise weights strike a balance between two conflicting forces. On the one hand, maximizing participation (extensive search margin) calls for minimizing the expected number of inspections and thus for positioning the popular products in the first slots. On the other hand, inducing active consumers to keep searching (intensive search margin) calls for positioning these products with some probability in the remaining slots. From Lemma 2, opting for random positioning satisfies the latter requirement, but does so with slack: as consumers become increasingly optimistic in the absence of a match, they actually become more prone to keep searching. It is therefore possible to assign the popular products to the later slots with slightly lower probability, and still induce consumers to keep searching until finding a match.

The positioning strategy $\hat{\boldsymbol{\rho}}(\mathcal{B})$ pushes this logic to the limit and "barely" preserves the intensive search margin: the more popular products are assigned to the later slots with just enough probability that the marginal consumer (i.e., the consumer initially indifferent between starting a search or not) remains indifferent, after unsuccessful inspections, between inspecting or not the remaining slots.[39] Specifically, the marginal consumer initially expects the first slot to deliver a match with greater probability, equal to $1-\bar{\lambda}(\mathcal{B})$. However, if the first slot does not produce a match, the logic of Lemma 2 applies: the consumer becomes more optimistic about the remaining slots, and now expects the second slot to deliver a match with the same probability $1-\bar{\lambda}(\mathcal{B})$; and so on.[40] Formally, letting $m_k \in \{0,1\}$ denote whether inspecting the k^{th} slot yields a match ($m_k = 1$) or not ($m_k = 0$), we have, for $k = 1, \ldots, |\mathcal{B}| - 1$:[41]

$$\Pr\left[m_k = 1 \mid m_1 = \cdots m_{k-1} = 0\right] = 1 - \bar{\lambda}(\mathcal{B}) > \max_{h>k} \Pr\left[m_h = 1 \mid m_1 = \cdots m_{k-1} = 0\right].$$

Remark 5 (*myopic search*) *Consumers' best response to the positioning strategy $\hat{\boldsymbol{\rho}}(\mathcal{B})$ coincides with the optimal myopic strategy: consumers with $c \leq \hat{c}(\mathcal{B}) = \left[1-\bar{\lambda}(\mathcal{B})\right] s^m$ inspect the first slot, which offers the highest probability of a match, equal to $1-\bar{\lambda}(\mathcal{B})$, as they would if they were to inspect a single slot. In the absence of a match, these consumers then inspect the second slot, which offers the highest revised probability of a match, equal again to $1-\bar{\lambda}(\mathcal{B})$, as they would if they were to inspect a single additional slot. And so forth.*

Remark 6 (*controlled search*) *It follows from Propositions 3 and 4 that dictating the order in which consumers can inspect slots (e.g., by arranging aisles in a specific manner, or by conditioning access to a given slot on having visited the previous one) does not benefit the firm, in the sense that the most profitable steering equilibrium generates the same profit, $\hat{\Pi}(\mathcal{B})$, even if the firm must incentivize consumers to inspect slots in the desired order.[42] In other words, the key incentive problem is to induce consumers to keep searching, rather than to induce them to adopt a particular order.*

[39] For $n = 2$, there is a single decision variable (say, the probability of positioning the more popular product in the first slot), which is uniquely pinned down by the need to induce active consumers to inspect both slots in the absence of a match. For $n > 2$, multiple positioning strategies support maximal participation.

[40] Consumers thus become more optimistic about the remaining slots in the absence of a match, as in the case of random positioning. However, contrary to that case, there is no "search addiction": any consumer with a cost exceeding $\hat{c}(\mathcal{B})$ would be unwilling to keep searching, even if forced to inspect the first slot.

[41] See Remark 9 in Appendix J.

[42] The firm could however do better if it were able to force consumers to inspect slots in a given order *and* to do so until a first match. In that case, pure positioning would be optimal, as it minimizes the expected number of inspections needed to obtain a match. By contrast, forcing consumers to inspect all slots, regardless of whether they already obtained a match, would be counterproductive and reduce participation.

4.3 Optimal steering

Building on the previous insights (and, in particular, on the additive separability of the profit function $\Pi(\boldsymbol{\mathcal{B}}, \hat{\mathbf{c}})$), the following proposition highlights the key features of the best steering equilibria from the firm's standpoint (with S for *Steering*):

Proposition 5 (most profitable steering) *For any given size $n \in \mathcal{I}_N$ adopted in the first stage, the most profitable steering continuation equilibria are such that, for some $\boldsymbol{\mathcal{B}}^S = (\mathcal{B}_1^S, \ldots, \mathcal{B}_{|\boldsymbol{\mathcal{B}}^S|}^S)$ and for every $k \in \mathcal{I}_{|\boldsymbol{\mathcal{B}}^S|}$:*

- *the firm assigns products $\mathcal{B}_k^S$ to slots $\mathcal{B}_k^S$ and obtains a profit equal to*

$$\sum_{k=1}^{|\mathcal{B}^S|} \Lambda\left(\bar{\mathcal{B}}_{k-1}^S\right) \hat{\Pi}(\mathcal{B}_k^S);$$

- *consumers with a search cost $c \leq \hat{c}(\mathcal{B}_k^S)$ inspect the slots of the first k blocks in increasing order until finding a match.*

The most profitable steering equilibria of the overall game entails the firm offering all products: $n = N$.

Proof. See Appendix J. ■

From Lemma 3, in the most profitable steering continuation equilibrium the products are structured in blocks $(\mathcal{B}_1^S, \ldots, \mathcal{B}_k^S)$, within which the firm positions the products so as to induce active consumers to search intensively until finding a match. Given the additive separability of the profit expression (4), profit is then maximal when the firm adopts a positioning strategy that maximizes participation—such as $\hat{\boldsymbol{\rho}}(\cdot)$—so that any consumer with a search cost below $\hat{c}(\mathcal{B}_k^S)$ visits block $\mathcal{B}_k^S$.

It also follows from Lemma 3 that, for any portfolio size n adopted in stage 1, the firm selects the n most popular products. To see why profit is maximal when the firm offers all products, it therefore suffices to note that, starting from an incomplete portfolio where the firm only offers $n < N$ products, adding the next-most popular product and assigning it to a specific slot, to be inspected last, does not affect consumers' search behavior among the first n products, and yet encourages some of those consumers who did not obtain a match with these products to inspect the additional one.

In the remainder of this section, we provide additional insights regarding the best block structure for the firm, $\boldsymbol{\mathcal{B}}^S$ (i.e., the block structure sustaining the most profitable steering equilibria). As we have seen, the comparison between the pure and random positioning equilibria highlights a trade-off between extensive and intensive search margins. Likewise, expanding the size of a block requires introducing noise in the positioning of a larger set of products and gives rise to a similar tradeoff: it induces intensive search over a larger product portfolio but discourages consumer participation; in particular, merging two successive blocks reduces the number of consumers willing to inspect them:

Corollary 3 (noisy positioning discourages participation) *For any block structure $\boldsymbol{\mathcal{B}} = (\mathcal{B}_1, \ldots, \mathcal{B}_{|\boldsymbol{\mathcal{B}}|})$ and any $k \in \mathcal{I}_{|\boldsymbol{\mathcal{B}}|} \setminus \{|\boldsymbol{\mathcal{B}}|\}$:*

$$\hat{c}(\mathcal{B}_k) > \hat{c}(\mathcal{B}_k \cup \mathcal{B}_{k+1}).$$

Proof. We have:

$$\hat{c}(\mathcal{B}_k) = \left[1 - \bar{\lambda}(\mathcal{B}_k)\right] s^m > \left[1 - \bar{\lambda}(\mathcal{B}_k \cup \mathcal{B}_{k+1})\right] s^m = \hat{c}(\mathcal{B}_k \cup \mathcal{B}_{k+1}). \quad \blacksquare$$

The design of the optimal block structure is driven both by the attractiveness of the products and the distribution of consumers' search costs. To gain further insights, we first hold the search cost distribution fixed and study the impact of product popularity on the scope for noisy positioning. Specifically, consider a steering equilibrium characterized by a block structure that includes two consecutive blocks $\mathcal{B}_1 = \{h, \ldots, k\}$ and $\mathcal{B}_2 = \{k+1, \ldots, \ell\}$. The contribution of these two blocks to the most profitable equilibrium based on this block structure is

$$\Lambda(\mathcal{I}_{h-1}) \left[\hat{\Pi}(\mathcal{B}_1) + \Lambda(\mathcal{B}_1)\hat{\Pi}(\mathcal{B}_2)\right].$$

If instead the two blocks are combined, so as to generate intensive search throughout $\mathcal{B} = \mathcal{B}_1 \cup \mathcal{B}_2$, the contribution of the enlarged block $\mathcal{B}$ is $\Lambda(\mathcal{I}_{h-2})\hat{\Pi}(\mathcal{B})$. It follows that doing so is profitable if $\Delta(\mathcal{B}_1, \mathcal{B}_2) \geq 0$, where

$$\Delta(\mathcal{B}_1, \mathcal{B}_2) \equiv \hat{\Pi}(\mathcal{B}_1 \cup \mathcal{B}_2) - \left[\hat{\Pi}(\mathcal{B}_1) + \Lambda(\mathcal{B}_1)\hat{\Pi}(\mathcal{B}_2)\right].$$

The following proposition shows that this can only be the case when the products of the two blocks are neither too different nor too similar:[43]

Proposition 6 (product popularities) *For any successive blocks $\mathcal{B}_1$ and $\mathcal{B}_2$ of given sizes n_1 and n_2:*

(i) $\Delta(\mathcal{B}_1, \mathcal{B}_2) < 0$ if $\bar{\lambda}(\mathcal{B}_1)$ is close enough to 0 or $\bar{\lambda}(\mathcal{B}_2)$ is close enough to 1;

[43] The proposition is loosely stated for ease of exposition. Formally, for any $\mathcal{B}_1$ and $\mathcal{B}_2$ of given sizes n_1 and n_2, the following holds for $\lambda_i = \bar{\lambda}(\mathcal{B}_i)$: $\forall \lambda_2 \in (0,1)$, $\exists \hat{\lambda}_1, \tilde{\lambda}_1 \in (0, \lambda_2)$ such that $\Delta(\mathcal{B}_1, \mathcal{B}_2) < 0$ for $\lambda_1 \in (0, \hat{\lambda}_1) \cup (\tilde{\lambda}_1, \lambda_2)$; and $\forall \lambda_1 \in (0,1)$, $\exists \tilde{\lambda}_2, \hat{\lambda}_2 \in (\lambda_1, 1)$ such that $\Delta(\mathcal{B}_1, \mathcal{B}_2) < 0$ for $\lambda_2 \in (\lambda_1, \tilde{\lambda}_2) \cup (\hat{\lambda}_2, 1)$.

(ii) $\Delta(\mathcal{B}_1, \mathcal{B}_2) < 0$ *if* $\bar{\lambda}(\mathcal{B}_2)$ *and* $\bar{\lambda}(\mathcal{B}_1)$ *are close enough to each other.*[44]

Proof. See Appendix L. ■

Combining the two blocks involves a trade-off between participation and search intensity. Indeed, the resulting change in profit is:

$$\Delta(\mathcal{B}_1, \mathcal{B}_2) = G(\hat{c}(\mathcal{B}))\,\mathrm{M}(\mathcal{B}) - G(\hat{c}(\mathcal{B}_1))\,\mathrm{M}(\mathcal{B}_1) - \Lambda(\mathcal{B}_1)\,G(\hat{c}(\mathcal{B}_2))\,\mathrm{M}(\mathcal{B}_2),$$

which, using $\mathrm{M}(\mathcal{B}) = \mathrm{M}(\mathcal{B}_1) + \Lambda(\mathcal{B}_1)\,\mathrm{M}(\mathcal{B}_2)$, can be expressed as $B(\mathcal{B}_1, \mathcal{B}_2) - C(\mathcal{B}_1, \mathcal{B}_2)$, where

$$C(\mathcal{B}_1, \mathcal{B}_2) \equiv [G(\hat{c}(\mathcal{B}_1)) - G(\hat{c}(\mathcal{B}))]\,\mathrm{M}(\mathcal{B}_1)\;(> 0) \tag{7}$$

denotes the cost of reducing the participation threshold (from $\hat{c}(\mathcal{B}_1)$ to $\hat{c}(\mathcal{B})$), and

$$B(\mathcal{B}_1, \mathcal{B}_2) \equiv [G(\hat{c}(\mathcal{B})) - G(\hat{c}(\mathcal{B}_2))]\,\Lambda(\mathcal{B}_1)\,\mathrm{M}(\mathcal{B}_2) \tag{8}$$

denotes instead the benefit from more intensive search: the bracketed term represents the number of consumers who would stop searching even without a match, a proportion $\Lambda(\mathcal{B}_1)\,\mathrm{M}(\mathcal{B}_2)$ of which have a match only with products from $\mathcal{B}_2$.

If $\mathcal{B}_2$ offers highly unpopular products (i.e., $\bar{\lambda}(\mathcal{B}_2) \sim 1$), the benefit of inducing active consumers to inspect these products is negligible: $B(\mathcal{B}_1, \mathcal{B}_2) \propto \mathrm{M}(\mathcal{B}_2) \simeq 0$; by contrast, the cost of reduced participation, $C(\mathcal{B}_1, \mathcal{B}_2)$, is non-negligible, as the participation threshold drops from $\hat{c}(\mathcal{B}_1) = \left(1 - \bar{\lambda}(\mathcal{B}_1)\right) s^m$ to $\hat{c}(\mathcal{B}) = \left(1 - \bar{\lambda}(\mathcal{B}_1)^{n_1/(n_1+n_2)}\right) s^m < \hat{c}(\mathcal{B}_1)$. It follows that combining the two blocks cannot be profitable. Likewise, when $\mathcal{B}_1$ offers highly popular products (i.e., $\bar{\lambda}(\mathcal{B}_1) \sim 1$), almost all active consumers obtain a match in $\mathcal{B}_1$; the benefit of inducing the few unlucky ones to keep searching is thus again negligible compared to the cost of reduced participation.[45]

Interestingly, Proposition 6 shows that combining the two blocks is not profitable either when they offer similar products (i.e., $\bar{\lambda}(\mathcal{B}_1) \sim \bar{\lambda}(\mathcal{B}_2)$), implying that the participation thresholds are also similar: $\hat{c}(\mathcal{B}_1) \sim \hat{c}(\mathcal{B}) \sim \hat{c}(\mathcal{B}_2)$; it follows that the cost and benefit of combining the two blocks are both small. If the blocks are of equal size (i.e., $n_1 = n_2 = n$), the number of consumers discouraged from participating is similar to the number of consumers willing to inspect $\mathcal{B}_1$ but not $\mathcal{B}_2$.[46] However, the per-consumer cost of reduced participation is given by the probability of having a match in the first

[44] $\bar{\lambda}(\mathcal{B}_1)$ and $\bar{\lambda}(\mathcal{B}_2)$ can be close to each other only if the two blocks include a single product (i.e., $n_1 = n_2 = 1$) or sufficiently similar ones, as $\bar{\lambda}(\mathcal{B}_1) \leq \lambda_k < \lambda_{k+1} < \lambda_{h+k} \leq \bar{\lambda}(\mathcal{B}_2)$, where the first (resp., second) weak inequality is strict whenever $n_1 > 1$ (resp., $n_2 > 1$).

[45] Specifically, using $\lambda_i = \bar{\lambda}(\mathcal{B}_i)$, we then have $B(\mathcal{B}_1, \mathcal{B}_2) \propto \Lambda(\mathcal{B}_1) = \lambda_1^{n_1}$, and $C(\mathcal{B}_1, \mathcal{B}_2) \propto G((1 - \lambda_1)s^m) - G((1 - \lambda_1^{\frac{n_1}{n}} \lambda_2^{\frac{n_2}{n}})s^m) \simeq g(s^m)\frac{n_1}{n}\lambda_1^{\frac{n_1}{n}-1}\lambda_2^{\frac{n_2}{n}} s^m \pi^m$, and so $\frac{B(\mathcal{B}_1,\mathcal{B}_2)}{C(\mathcal{B}_1,\mathcal{B}_2)} \propto \lambda_1^{1+n_1\frac{n_1+n_2-1}{n_1+n_2}} \simeq 0$.

[46] Indeed, if $n_1 = n_2 = n$ and $\varepsilon = (\lambda_2 - \lambda_1)/\lambda_1$ is small, then $G(\hat{c}(\mathcal{B}_1)) - G(\hat{c}(\mathcal{B})) \simeq G(\hat{c}(\mathcal{B})) - G(\hat{c}(\mathcal{B}_2)) \simeq g((1 - \lambda_1)s^m)\,\lambda_1 s^m \varepsilon/2$.

block, $\mathrm{M}(\mathcal{B}_1)$, whereas the per-consumer benefit of intensive search is given by the probability of having a match only in the second block, $\Lambda(\mathcal{B}_1)\mathrm{M}(\mathcal{B}_2)$; as the first block offers more popular products, it follows that the per-consumer cost exceeds the per-consumer benefit of combining the two blocks: $\mathrm{M}(\mathcal{B}_1) > \Lambda(\mathcal{B}_1)\mathrm{M}(\mathcal{B}_2)$.[47] Hence, combining the two blocks is not profitable. As shown in Appendix L, the reasoning extends to blocks of unequal size.

We now turn to the impact of the distribution of search costs. The benefit from combining blocks stems from consumers with relatively low search costs (namely, from $\hat{c}(\mathcal{B}_1)$ to $\hat{c}(\mathcal{B})$), who now search more intensively; the cost stems instead from consumers with relatively high search costs (from $\hat{c}(\mathcal{B})$ to $\hat{c}(\mathcal{B}_2)$), who stop participating. Hence, the comparison between the cost and benefit of combining two blocks hinges on the relative numbers of consumers with search costs lying in these two ranges.

To gain further insight, we now focus on distributions characterized by a constant hazard rate $g(c)/[1-G(c)] = \gamma > 0$, that is:

$$G_\gamma(c) = 1 - \exp(-\gamma c).$$

For any successive blocks $\mathcal{B}_1$ and $\mathcal{B}_2$, let $\Delta_\gamma(\mathcal{B}_1, \mathcal{B}_2)$ denote the change in profit from combining them when search costs are distributed according to $G_\gamma(\cdot)$. This family of distributions satisfies the monotone likelihood ratio property: for any $c'' > c'$, $g_\gamma(c'')/g_\gamma(c')$ is strictly decreasing in γ. It follows that raising γ also increases the cost-benefit ratio—which, from the above, is given by $[G_\gamma(\hat{c}(\mathcal{B}_2)) - G_\gamma(\hat{c}(\mathcal{B}))]/[G_\gamma(\hat{c}(\mathcal{B})) - G_\gamma(\hat{c}(\mathcal{B}_1))]$—thus making it more profitable to combine the two blocks. Building on this leads to:

Proposition 7 (cost distribution) *For any $\mathcal{B}_1$ and $\mathcal{B}_2$ there exists $\hat{\gamma}(\mathcal{B}_1, \mathcal{B}_2) \geq 0$ such that $\Delta_\gamma(\mathcal{B}_1, \mathcal{B}_2) > 0$ if and only if $\gamma > \hat{\gamma}(\mathcal{B}_1, \mathcal{B}_2)$.*

Proof. See Appendix M. ■

It is shown in Appendix M that $\Delta_\gamma(\mathcal{B}_1, \mathcal{B}_2)$ strictly increases with γ. Hence, for any given successive blocks $\mathcal{B}_1$ and $\mathcal{B}_2$, there exists a threshold $\hat{\gamma}(\mathcal{B}_1, \mathcal{B}_2)$ such that combining the two blocks is profitable if γ exceeds the threshold, and is otherwise unprofitable.

It follows from Proposition 7 that an increase in γ, which generates a first-order stochastically dominant reduction in search costs, can prompt the most profitable

[47] Indeed, $\mathrm{M}(\mathcal{B}_1) - \Lambda(\mathcal{B}_1)\mathrm{M}(\mathcal{B}_2) = 1 - \Lambda(\mathcal{B}_1) - \Lambda(\mathcal{B}_1)[1 - \Lambda(\mathcal{B}_2)] = 1 - 2\Lambda(\mathcal{B}_1) + \Lambda(\mathcal{B}_1)\Lambda(\mathcal{B}_2) > [1 - \Lambda(\mathcal{B}_1)]^2 > 0$, where the inequalities stem from $1 > \Lambda(\mathcal{B}_2) > \Lambda(\mathcal{B}_1)$.

equilibrium to switch from pure to noisy positioning. As a result, such a reduction in search costs can actually harm consumers:[48]

Corollary 4 (reduction in search costs) *Suppose that the most profitable equilibrium is played. Then a reduction in every consumer' search cost can harm all participating consumers.*

Proof. See Appendix N. ■

It also follows from Proposition 7 that, for γ large enough, combining *any* two successive blocks is *profitable*; the most profitable steering equilibrium is thus when the firm adopts the best noisy positioning strategy, so as to induce all active consumers to keep searching until finding a match. Conversely, for γ low enough, combining *any* two successive blocks is *unprofitable*; the most profitable steering equilibrium then entails pure positioning. That is, letting $\mathcal{P}_N$ denote the set of pairs of successive blocks from $\mathcal{I}_N$,[49] and defining

$$\underline{\gamma} \equiv \min \{\hat{\gamma}(\mathcal{B}_1, \mathcal{B}_2) \mid (\mathcal{B}_1, \mathcal{B}_2) \in \mathcal{P}_N\} \text{ and } \bar{\gamma} \equiv \max \{\hat{\gamma}(\mathcal{B}_1, \mathcal{B}_2) \mid (\mathcal{B}_1, \mathcal{B}_2) \in \mathcal{P}_N\},$$

we have:

Corollary 5 (scope for noisy positioning) *If $\gamma < \underline{\gamma}$, then pure positioning maximizes the firm's profit among all steering continuation equilibria. By contrast, if $\gamma > \bar{\gamma}$, then for any given product portfolio size $n \in \mathcal{I}_N$ adopted in the first stage, the most profitable steering continuation equilibrium entails $\boldsymbol{\mathcal{B}}^S = \{\mathcal{I}_n\}$.*

Proof. Follows immediately from Proposition 7 and the definition of $\bar{\gamma}$. ■

Note that $\underline{\gamma}$ may be zero, in which case some noisy positioning is always optimal.

Remark 7 (choice overload) *Our analysis emphasizes that the scope for choice overload critically depends on the firm's positioning strategy, which in turn depends on its incentive to introduce noise in the information available to consumers. Suppose for instance that the firm adds newly available products to its assortment. Under pure positioning, choice overload simply does not arise, as consumers then have perfect information and can therefore focus on the most popular products if they wish so. By contrast, noisy positioning limits the information available to consumers; this reduces*

[48]For example, if γ is slightly below $\hat{\gamma}(\mathcal{B}_1, \mathcal{B}_2)$, then an arbitrarily small decrease in consumers' search costs would have a negligible direct impact on their surplus but expand the scope for noisy positioning.

[49]That is, $\mathcal{P} \equiv \{(\mathcal{B}_1, \mathcal{B}_2) \mid \mathcal{B}_1 = \mathcal{I}_k \setminus \mathcal{I}_h \text{ and } \mathcal{B}_2 = \mathcal{I}_\ell \setminus \mathcal{I}_{k+1} \text{ for some } h \leq k \leq \ell - 1 \leq N\}$.

consumer participation but may nevertheless be profitable if it induces more intensive search. By characterizing the firm's incentives to limit the information available to consumers, Propositions 6 and 7 shed light on when choice overload is more likely to arise—namely, when products are neither too different nor too similar, and the distribution of search costs has a large hazard rate.

4.4 Welfare implications

The above analysis highlights two types of distortions—in addition to the classic distortion arising from monopoly pricing. First, for any given assortment size, balancing consumer participation and search intensity can lead the firm to limit the information available to consumers. This harms consumers, by preventing them to inspect first the best products.[50]

Second, the firm may offer too many products. Indeed, the firm never finds it optimal to limit the size of its assortment: from Proposition 5, the most profitable equilibrium entails $n = N$. Yet, noisy positioning creates choice overload, which harms those consumers who are uninterested in the least popular products. As a result, the firm may provide excessive variety.

To see this, suppose for instance that, for any assortment size chosen in stage 1, (i) the most profitable continuation equilibrium is played and (ii) this entails placing the products in one single block (e.g., because the search cost distribution has a sufficiently large constant hazard rate, $\gamma > \bar{\gamma}$). Compare now two of the continuation equilibria, one ($\mathcal{E}_N$) in which the firm chose in stage 1 to offer all N products ($\mathcal{I} = \mathcal{I}_N$), and one ($\mathcal{E}_n$) in which it chose instead to offer only $n < N$ of them ($\mathcal{I} = \mathcal{I}_n$). By construction, $\mathcal{E}_N$ is more profitable than $\mathcal{E}_n$: as $\gamma > \bar{\gamma}$, $\mathcal{E}_N$ is more profitable than the equilibrium $\mathcal{E}'_N$ in which the firm adopts the best strategy inducing intensive search in $\mathcal{I}_n$ and assigns each product in $\mathcal{I}_N \diagdown \mathcal{I}_n$ to a designated slot; and as noted above, $\mathcal{E}'_N$ is more profitable than $\mathcal{E}_n$, as consumers are given more options (and full information about the additional options). For consumers, however, the comparison between $\mathcal{E}_N$ and $\mathcal{E}_n$ is more nuanced. To be sure, consumers with negligible search costs would also favor $\mathcal{E}_N$, which offers more choice. But consumers with higher costs would instead favor $\mathcal{E}_n$, which offers a higher expected probability of a match in every slot. In particular, consumers with a search cost $c \in (\hat{c}(\mathcal{I}_N), \hat{c}(\mathcal{I}_n))$ are active in $\mathcal{E}_n$, and inactive in $\mathcal{E}_N$. Furthermore, if $c^P_{n+1} < \hat{c}(\mathcal{I}_N)$ (which can be the case when n is sufficiently close to N, or the products in $(\mathcal{I}_n)$ are sufficiently popular), then consumers with $c \in \left(c^P_{n+1}, \hat{c}(\mathcal{I}_n)\right)$ favor $\mathcal{E}_n$ over $\mathcal{E}_N$, even though they are active in both equilibria.[51] It is actually easy

[50] It even induces some consumers to inspect products that they would never inspect if they knew their identity.

[51] To see this, note that (i) all consumers prefer $\mathcal{E}'_N$ to $\mathcal{E}_N$, as they are better informed in the former

to find examples in which *all* consumers would benefit from downsizing the product line.[52]

This discussion suggests that, even though the interest of the firm appears aligned with consumers' own interests, in that they both seek to maximize the probability of a sale, not internalizing consumers' search costs can induce the firm to adopt strategies that harm consumers: first, it has an incentive to introduce some noise when positioning its product portfolio; conversely, conditional on noisy positioning, it has an incentive to offer too many products. Among possible remedies, imposing a ban on noisy positioning would appear natural; it would improve the information available to consumers, enabling them to maximize their utility, and would moreover preserve the firm's incentives to offer all available products—indeed, under pure positioning, it is always profitable for the firm to expand its product line. Short of achieving this, an alternative approach may consist in putting a cap on the size of the product line—determining the appropriate cap would however require substantial information on product characteristics.

5 Extensions

5.1 Conditioning on search costs

Observing past behavior can sometimes enable firms—e.g., online retailers—to collect information about consumers' characteristics, and use that information to tailor their offering. In our setting, the firm could for instance learn some information about consumers' search costs, and may be able to offer different recommendations based on that information. To study the implications of such personalized offering, suppose that the firm perfectly observes a consumer's search cost, c, and chooses the assortment size n and product positioning accordingly.[53]

The trade-off between participation (calling for pure positioning) and search intensity (calling for noisy positioning) now takes a different form, as the firm has an incentive to maximize the number of products that a consumer with this particular search cost c is willing to inspect. That is, it is best for the firm to offer the largest assortment that "meets" the consumer's participation constraint and position the products so as to induce the consumer to keep searching until finding a match. It follows that pure positioning is always dominated: the most profitable equilibrium

equilibrium, and (ii) consumers with $c \in \left(c_{n+1}^P, \hat{c}\left(\mathcal{I}_n\right)\right)$ would not inspect the products $\mathcal{I}_N \diagdown \mathcal{I}_n$ if positioned in designated slots, and so obtain the same surplus in $\mathcal{E}_n$ and $\mathcal{E}'_N$.

[52] See Appendix O for a formal derivation.

[53] Alternatively, this corresponds to the case of homogeneous search costs.

consists in choosing the largest assortment size n satisfying

$$\left[1-\bar{\lambda}\left(\mathcal{I}_n\right)\right] s^m \geq c, \tag{9}$$

and to position the n most popular products using a noisy positioning strategy, such as $\hat{\boldsymbol{\rho}}\left(\mathcal{I}_n\right)$, that induces the consumer to keep searching but minimizes the inspected number of inspections. That is, we have:

Proposition 8 (conditioning on search costs) *Suppose that the firm can condition its strategy on consumers' search costs. For any given search cost c, the most profitable steering equilibria are such that:*

(i) in stage 1, the firm chooses

$$n^C(c) \equiv \max\left\{n \mid (9) \ \textit{holds}\right\},$$

which is decreasing in c;

(ii) if $n^C(c) > 1$, then in stage 2 the firm selects the $n^C(c)$ most popular products and introduces noise in their positioning so as to induce the consumer to inspect all selected products until finding a match;

(iii) ignoring integer constraints, (9) is binding and the consumer thus obtains zero surplus.[54]

Proof. This follows directly from the above observations. ■

The conflict of interest between the firm and the consumers is therefore particularly crisp when the firm can tailor its strategy to consumers' search costs. While efficiency and consumers' interest would call for pure positioning and no restriction on the product line, it is instead most profitable for the firm to introduce noise in its positioning strategy, and for limiting the product line to the maximal number of products that consumers are willing to inspect given that positioning. Absent integer constraints on the number of products, consumers would therefore obtain zero surplus.

5.2 Aggregate uncertainty

We have so far assumed—in line with the consumer search literature—that consumers have a well-identified prior about the range of products that the firm may choose from. This is a reasonable assumption when, through past experience or feedback

[54] When integer constraints are taken into account, an upper bound on consumer surplus is $[\bar{\lambda}(\mathcal{I}_{n^C(c)+1}) - \bar{\lambda}(\mathcal{I}_{n^C(c)})]s^m$, which decreases as c decreases; in particular, if infinitely many products are available (i.e., $N = \infty$), then it becomes arbitrarily small as c tends to 0 .

from friends and family, consumers have an idea of "what to expect", both in terms of available products (e.g., hotels of well-known cities) and in terms of the quality of the firm in providing access to these products (e.g., a platform such as *Expedia* or *Booking*). In other situations, however, there may be some uncertainty about the set of available products (e.g., hotels in a remote location) or about the intrinsic quality of the firm (e.g., the level of assistance in case of trouble). In such situations, the absence of a valuable match on the first inspections may lead consumers to revise their beliefs about the overall quality of the offerings and discourage them from inspecting further options.[55]

To study this issue, we consider in Appendix P an extended version of our setting that allows for aggregate uncertainty about the general desirability of the products. For ease of exposition, we suppose that there are two available products, both offered by the firm (i.e., $n = N = 2$), and two states of the world, a "good" one and a "bad" one in which both products have lower match probabilities. It is first shown that, under *random positioning*, the absence of a match on a first inspection induces consumers to become more optimistic about the uninspected product as long as there is sufficiently more heterogeneity across products than across states. Specifically, this increased optimism feature persists as long as

$$\mathrm{E}\left[\mu_1\right] - \mathrm{E}\left[\mu_2\right] > 2\sqrt{\mathrm{Cov}\left(\mu_1, \mu_2\right)},$$

where the left-hand side reflects product uncertainty (measured by the difference in expected popularity), whereas the right-hand side reflects state uncertainty (measured by $\mathrm{Cov}\left(\mu_1, \mu_2\right) = \mathrm{E}\left[\left(\mu_1 - \mathrm{E}\left[\mu_1\right]\right)\left(\mu_2 - \mathrm{E}\left[\mu_2\right]\right)\right]$). Intuitively, whether a match occurs then conveys more information about the inspected product than about the underlying state. As a result, random positioning still creates a search addition pattern—consumers who start searching are eager to keep searching until finding a match.

It is also shown that, as long as this condition holds, the rest of the previous analysis carries over: there exists a variety of equilibria, in which the firm either opts for pure positioning (which maximizes participation) or for sufficiently noisy positioning (which maximizes search intensity among active consumers). Interestingly, the introduction of aggregate uncertainty can alter this trade-off in either direction. Two examples are provided in which, keeping $\mathrm{E}\left[\mu_1\right]$ and $\mathrm{E}\left[\mu_2\right]$ constant, increasing $\mathrm{Cov}\left(\mu_1, \mu_2\right)$ can tilt the balance in favor of pure or noisy positioning.

[55] In addition, as noted by Kamenica (2008), if the uninformed consumers' uncertainty is not about the individual products, but rather about the general desirability of the good, then expanding the product line can constitute a positive signal that increases demand.

5.3 Increasing or decreasing search costs

We have so far assumed—as almost all the literature on sequential consumer search—that the cost of inspecting an additional product does not depend on the number of previous inspections. We consider in Appendix Q variants in which search costs are either decreasing in the number of inspections (e.g., thanks to learning effects, or because visiting the firm inflates the cost of the first inspection), or increasing (e.g., because of diseconomies of scale or time constraints).[56]

In the case of decreasing costs, consumers are even more prone to *search addiction*. The firm can thus reduce the amount of noise needed to induce consumers to keep searching until finding a match, which fosters participation.[57] Decreasing search costs thus enhance the extensive margin under noisy positioning, by encouraging consumers to start searching. However, it also enhances the intensive margin under pure positioning, by encouraging consumers to keep searching in the absence of a match with the first product. As a result, it can tilt the balance in favor of either pure or noisy positioning—see Appendix Q for illustrative examples.

Conversely, if search costs increase with the number of inspections, consumers are less prone to search addiction—and indeed, inducing an intensive search pattern may no longer be possible if search costs are sufficiently increasing, in which case noisy positioning is unprofitable. Furthermore, even if search addiction still arises, inducing consumers to keep searching requires noisier positioning, which discourages participation (the extensive search margin). However, compared to constant search costs, increasing costs also reduce the intensive margin under pure positioning, by discouraging consumers to keep searching. As a result, an increasing cost pattern can also tilt the balance in favor of either noisy or pure positioning—see Appendix Q for illustrative examples.

5.4 Disclosure policies

We have so far assumed that an inspection reveals product characteristics such as prices and match valuations, but does not convey any additional information about the identity of the inspected products – and, thus, of the uninspected ones. Suppose now that the firm has a "disclosure technology" at its disposal, which enables it to communicate, upon inspection, the identity of the product. We first note that the possibility of disclosure eliminates any scope for pure positioning:

[56] The latter case also arises when search costs are slot-specific (e.g., because some slots are easier to reach than others), as the firm has then an incentive to induce consumers to start with the slots that are the least costly to inspect.

[57] If costs are sufficiently decreasing, search addiction may arise even under pure positioning, in which case introducing noise cannot be profitable.

Proposition 9 (disclosure) *Suppose that the disclosure technology is available. For any given product portfolio size $n \in \mathcal{I}_N$ adopted in the first stage, any steering continuation equilibrium has the block structure $\boldsymbol{\mathcal{B}} = (\mathcal{I}_n)$; conversely, the steering equilibria featuring intensive search over $\mathcal{I}_n$ all survive.*

Proof. See Appendix R. ■

The intuition is straightforward. If the firm expects some consumers to stop searching after the k^{th} inspection, it has an incentive to place one of the least popular products in that slot *and* disclose its identity, so as to induce consumers to keep inspecting the remaining, more popular products.[58] As a result, consumers would not follow the recommended order but instead inspect first the last slots.

The possibility of disclosure also eliminates the non-monopoly pricing equilibria discussed in Remark 4, in which prices are used to reveal the identity of the products. These equilibria have the feature that popular products are offered at the monopoly price, whereas less popular products are offered at sub-optimal prices; yet, the firm is prevented from deviating to the monopoly price, as this would "signal" a popular product and prompt some consumers to stop searching. Thanks to the disclosure technology, the firm can instead deviate to the monopoly price *and* reassure consumers that they inspected a not-so-popular product, thus encouraging consumers to keep searching in the absence of a match.[59]

5.5 Platforms

We have so far considered a setting in which the firm directly sells its products to consumers. We now study how our analysis applies when the firm acts instead as a platform on which third-party suppliers sell their products to consumers. For ease of exposition, we suppose that each product is supplied by an independent firm, which only offers that product; we will thus identify a product with its supplier.

5.5.1 Non-linear tariffs based on sales

We first suppose that the platform charges a sales-based non-linear tariff to each hosted supplier i. The timing is as follows:

Stage 1 The platform publicly chooses the number n of suppliers that it is willing to host.

[58] This argument relies on the assumption that consumers, in the spirit of passive beliefs, keep expecting the firm to offer the equilibrium product portfolio—see footnote 81.

[59] Conversely, the possibility of disclosure gives rise to additional monopoly pricing equilibria. Nocke and Rey (2021) provide a complete characterization of disclosure equilibria for the case $N = 2$.

Stage 2 a. The platform makes take-it-or-leave-it offers to n suppliers, which the suppliers accept or reject;[60] the number of suppliers eventually hosted is publicly observed.

b. Hosted suppliers set their prices; the platform then observes suppliers' prices and positions them.

c. Consumers observe their search costs and sequentially decide which slots to inspect (if any).

The steering equilibria considered in our baseline setting all have an equivalent here:

Proposition 10 (platforms) *Fix any given product portfolio size $n \in \mathcal{I}_N$ and any continuation steering equilibrium of the baseline setting, in which the firm offers the n most popular products at the monopoly price p^m, positions them according to a strategy $\boldsymbol{\rho}$ that induces a search pattern characterized by a block structure $\boldsymbol{\mathcal{B}} = (\mathcal{B}_1, \ldots, \mathcal{B}_{|\boldsymbol{\mathcal{B}}|})$ and search cost thresholds $\hat{\mathbf{c}} = \left(\hat{c}_1, \ldots, \hat{c}_{|\boldsymbol{\mathcal{B}}|}\right)$, and obtains the profit $\Pi\left(\boldsymbol{\mathcal{B}}, \hat{\mathbf{c}}\right)$ given by (4). When the platform charges sales-based non-linear tariffs, there exists an analogous equilibrium in which the platform selects the n most popular suppliers, who sell at the monopoly price p^m, positions them according to the same block structure $\boldsymbol{\mathcal{B}}$ and the same strategy $\boldsymbol{\rho}$, induces the same search behavior from consumers, and obtains the same profit.*

Proof. See Appendix S. ■

Remark 8 (bargaining power and linear tariffs) *The analysis carries over when suppliers have some bargaining power in their negotiations with the platform, e.g., through "Nash-in-Nash" bilateral bargaining. Suppose indeed that the platform and every supplier $i \in \mathcal{I}$ engage in Nash bargaining, holding fixed the equilibrium agreements between the platform and the other suppliers in $\mathcal{I}$, and let $\omega \in (0, 1]$ denote the platform's bargaining power. As suppliers obtain zero profit in the absence of an agreement and the platform derives its profit solely through the fixed fees, the platform*

[60] We assume that the platform and the suppliers observe the identities of the products at the beginning of stage 2, so that it still constitutes a proper subgame.

Whether the offers and acceptance decisions are private or public does not affect the analysis. We maintain the assumption of passive beliefs and extend it to suppliers as well. For the sake of exposition, we also assume that, if the platform were to approach supplier $i > n$, that supplier would expect to be positioned according to the equilibrium positioning strategy designed for supplier n.

then obtains a fraction ω of the profit generated by each supplier on the platform;[61] *hence, it still has an incentive to select the most popular suppliers.*

A similar analysis applies when the platform charges a uniform fee per transaction (no discrimination), with the caveat that this creates a double marginalization problem and thus induces suppliers to charge higher prices.[62]

5.5.2 From sales to clicks

Conditioning the remuneration of the platform on *clicks* (i.e., on inspections), rather than sales, would create a drastic conflict of interests, as consumers want to *minimize* the expected number of inspections, whereas the platform would instead seek to foster consumer attention and thus *maximize* the expected number of inspections. This, in turn, would have implications for positioning and product selection.[63]

First, this would eliminate any scope for steering: indeed, if all consumers were to start with the same slot 1, the platform would have an incentive to position the least popular product in that slot; but then, consumers would not inspect this slot first. There would however exist an equilibrium with random positioning, in which consumers would also uniformly randomize (or randomize in such a way that slots are equally likely to be inspected first, second, and so forth).

Second, for any given portfolio size n adopted in stage 1, the platform would actually have an incentive to choose in stage 2 the least popular suppliers. However, if the platform could commit itself to a given portfolio, it may still select the most popular products.[64]

6 Conclusion

In this paper, we study a firm's ability and incentive to steer consumers. A monopolist chooses which products to offer and positions them in distinguishable slots. Consumers

[61]Let π_i denote the profit obtained by seller $i \in \mathcal{I}$ if it joins the platform. If the negotiation breaks down, seller $i \in \mathcal{I}$ obtains $\underline{\pi}_i = 0$ and the platform obtains $\underline{\Pi} = \sum_{j \in (\mathcal{I} \setminus \{i\})} \phi_j$; if instead the negotiation is successful, their profits are respectively $\hat{\pi}_i = \pi_i - \phi_i$ and $\hat{\Pi} = \underline{\Pi} + \phi_i$. The Nash bargaining rule then yields $\phi_i = \hat{\Pi} - \underline{\Pi} = \omega \left(\hat{\Pi} + \hat{\pi} - \underline{\Pi} - \underline{\pi} \right) = \omega \pi_i$.

[62]Letting w denote the wholesale price charged by the platform for each transaction, each supplier i then chooses $p^m(w) = \arg\max_p \{(p-w)[1-F(p)]\}$, regardless of the popularity of its product.

[63]In a similar environment, Eliaz and Spiegler (2011) study the problem of a monopoly platform that has to decide on the price-per-click, assuming that search on the platform is necessarily random.

[64]This would be the case, for instance, when search costs are uniformly distributed, as the platform would then seek to maximize

$$G\left(c^R(\mathcal{I})\right)\Gamma(\mathcal{I}) = G\left(\frac{\mathrm{M}(\mathcal{I})}{\Gamma(\mathcal{I})} s^m\right)\Gamma(\mathcal{B}) = \frac{\mathrm{M}(\mathcal{B})}{\Gamma(\mathcal{B})} s^m \times \Gamma(\mathcal{B}) = \mathrm{M}(\mathcal{B})\, s^m.$$

observe the size of the product line and can engage in sequential and directed search to learn about prices, whether or not they have a match and, if so, their conditional valuations. While searching, consumers update their beliefs about the not-yet-inspected slots.

Whether or not the firm has the incentive to steer consumers towards the most popular products first is shown to be governed by a trade-off between extensive and intensive search margins. Pure positioning enables consumers to infer products' locations and thus encourages them to start searching; it therefore maximizes the extensive search margin. The flip side is that consumers with higher search costs stop searching if they do not have a match with the most popular products. Introducing noise in the allocation of products to slots encourages instead consumers who start searching to keep searching until finding a match, thereby maximizing the intensive search margin. The flip side is that it reduces consumers' incentive to start searching in the first place. The most profitable equilibrium may feature noisy positioning of some products and pure positioning of others. We provide conditions under which the most profitable equilibrium involves pure or noisy positioning over *all* offered products.

These findings relate to the on-going policy debate on steering. Our analysis suggests that firms may not steer consumers to the best products first, even in the absence of concerns such as self-preferencing or distorted competition arising from position auctions. Indeed, in our setting the interest of the firm is seemingly aligned with those of consumers: conditional on a match, all products generate the same expected profit for the firm and the same expected surplus for consumers. Both parties thus have an incentive to maximize the probability of a match, and there is no bias in favor of any "in-house" product or of any particular third-party seller. Yet, the firm does not internalize consumers' search costs and, as a result, may benefit from limiting the information available to them so as to induce more intensive—and more costly—search patterns.

We also show that the firm may offer too much variety—a form of choice overload. Our analysis moreover reveals that this can happen precisely when the firm introduces noise in its positioning. Hence, studying the incentives to limit the information available to consumers helps identify the types of situations in which choice overload is more likely to occur. In particular, we find that noisy positioning and choice overload are particularly likely to occur when the firm has information on consumers' characteristics that enables it to tailor its strategy to consumers' search costs.

An exciting avenue for future work consists in studying the extent to which competition between firms or platforms may act as a disciplining device and further align firms' interests with those of consumers.

References

Anderson, Simon P., and Régis Renault (2021). Search Direction: Position Externalities and Position Auction Bias, CEPR Discussion Paper 16724, Centre for Economic Policy Research.

Armstrong, Mark (2017). Ordered Consumer Search, *Journal of the European Economic Association*, vol. 15(5), 989-1024.

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

Athey, Susan, and Glenn Ellison (2011). Position Auctions with Consumer Search, *Quarterly Journal of Economics*, vol. 126 (3): 1213-1270

Boatwright, Peter and Joseph C. Nunes (2001). Reducing Assortment: An Attribute-Based Approach, *Journal of Marketing*, vol. 65(3): 50–63.

Chen, Yongmin, and Chuan He (2011). Paid Placement: Advertising and Search on the Internet, *Economic Journal*, vol. 121: F309-F328.

Eliaz, Kfir, and Ran Spiegler (2011). A Simple Model of Search Engine Pricing, *Economic Journal*, vol. 121: F329-F339.

Gamp, Tobias (2019). Guided Search. *Mimeo.*

Hagiu, Andrei, and Bruno Jullien (2011). Why Do Intermediaries Divert Search?, *RAND Journal of Economics*, vol. 42(2): 337-362.

Hämäläinen, Saara (2022). Multiproduct Search Obfuscation, *International Journal of Industrial Organization*, vol. 85, 102863.

Hannak, Aniko, Gary Soeller, David Lazer, Alan Misolove, and Christo Wilson (2014). Measuring Price Discrimination and Steering on E-commerce Web Sites, *Proceedings of the 2014 Conference on Internet Measurement*, 305–318.

Hart, Oliver, and Jean Tirole (1990). Vertical Integration and Market Foreclosure, *Brookings Papers on Economic Activity. Microeconomics*, 1990: 205-276.

Hunold, Matthias, Ulrich Laitenberger, and Guillaume Thébaudin (2022). Bye-box: An Analysis of Non-Promotion on the Amazon Marketplace, CRED Working Paper No. 2022-4.

Iyengar, Sheena S., and Mark R. Lepper (2000). When Choice is Demotivating: Can One Desire Too Much of a Good Thing?, *Journal of Personality and Sociological Psychology*, vol. 79(6), 995-1006.

Kahn, Barbara E. and Brian Wansink (2004). The Influence of Assortment Structure on Perceived Variety and Consumption Quantities, *Journal of Consumer Research*, vol. 30(4): 519–33.

Kamenica, Emir (2008). Contextual Inference in Markets: On the Informational Content of Product Lines, *American Economic Review*, vol. 98(5): 2127-2149.

Kreps, David M., and Robert Wilson (1982). Sequential Equilibria, *Econometrica*, vol. 50(4): 863-894.

Kuksov, Dmitri, and J. Miguel Villas-Boas (2010). When More Alternatives Lead to Less Choice, *Marketing Science*, vol. 29(3): 507-524.

Lee, Kwok Hao, and Leon Musolff (2021). Entry Into Two-Sided Markets Shaped By Platform-Guided Search, Mimeo.

McAfee, R. Preston, and Marius Schwartz (1994). Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity, *American Economic Review*, vol. 84(1): 210-230.

McDevitt, Ryan C. (2014). "A" Business by Any Other Name: Firm Name Choice as a Signal of Firm Quality, *Journal of Political Economy*, vol. 122(4): 909-944.

Moraga-González, José Luis, Zsolt Sándor, and Matthijs R. Wildenbeest (2017). Prices and Heterogeneous Search Costs, *RAND Journal of Economics*, vol. 48(1): 125-146.

Natan, Olivia R. (2022). Choice Frictions in Large Assortments, Mimeo.

Nocke, Volker, and Patrick Rey (2021). Consumer Search and Choice Overload. CEPR Discussion Paper 16440, Centre for Economic Policy Research.

O'Brien, Dan P., and Greg Shaffer (1992). Vertical Control with Bilateral Contracts, *RAND Journal of Economics*, vol. 23(3): 299-308.

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

Petroski, Henry (2003). *Small Things Considered: Why There Is No Perfect Design*, New York: Knopf.

Rayo, Luis, and Ilya Segal (2010). Optimal Information Disclosure, *Journal of Political Economy*, vol. 118(5): 949-987.

Rhodes, Andrew (2015). Multiproduct Retailing, *Review of Economic Studies*, vol. 82(1): 360-390.

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

Scheibehenne, Benjamin, Rainer Greifeneder, and Peter M. Todd (2010). Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload, *Journal of Consumer Research*, vol. 37: 409-425.

Shih, Willy C., Stephen P. Kaufman, and David Spinola (2007). Netflix, Harvard Business School Case 607-138. (Revised April 2009.)

Spence, A. Michael (1975). Monopoly, Quality, and Regulation, *Bell Journal of Economics*, vol. 6(2): 417-429.

Stahl, Dale O. (1989). Oligopolistic Pricing with Sequential Consumer Search, *American Economic Review*, vol. 79(4): 700-712.

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

Villas-Boas, J. Miguel (2009). Product Variety and Endogenous Pricing with Evaluation Costs, *Management Science*, vol. 55(8), 1338-1346.

Wolinsky, Asher (1986). True Monopolistic Competition as a Result of Imperfect Information, *Quarterly Journal of Economics*, vol. 101(3): 493-512.

Zhou, Jidong (2014). Multiproduct Search and the Joint Search Effect, *American Economic Review*, vol. 104(9): 2918-2939

Appendix

A Proof of Lemma 1 (monopoly pricing)

Fix the product portfolio size $n \in \mathcal{I}_N$ adopted in the first stage, and consider a candidate equilibrium in which, with probability 1, the firm selects the product portfolio $\mathcal{I}$ (satisfying $|\mathcal{I}| = n$), allocates product $[k] \in \mathcal{I}$ to slot $k \in \mathcal{I}_n \equiv \{1, \ldots, n\}$, and charges the price $p_{[k]}$ for that product.

Upon having a first match, a consumer learns his valuation but may however keep searching in order to obtain a lower price. The following lemma shows that the consumer then inspects slots by increasing order of their *cost-adjusted* prices, and stops searching when these adjusted prices exceed the maximal value from obtaining a lower price. Specifically, for $k \in \mathcal{I}_n$, let

$$\tilde{p}_{[k]}(c) \equiv p_{[k]} + \frac{c}{\mu_{[k]}}$$

denote the price in slot k, properly adjusted for the cost required to obtain it, and

$$\tilde{v}(v; p) \equiv \min\{v, p\}$$

denote the consumer's maximal value from inspecting additional slots after a first match delivering a valuation v at price p. We have:[65]

Claim A.1 *Consider a consumer with search cost c and valuation v having a first match upon inspecting slot $\ell \in \mathcal{I}_n$, and let $\mathcal{U}$ denote the set of yet-uninspected slots, and $\boldsymbol{\sigma} \equiv (\sigma_1, \ldots, \sigma_k) \in \mathcal{S}(\mathcal{I}_k)$, for some $k \in \mathcal{I}_{|\mathcal{U}|}$, denote the consumer's equilibrium search sequence among these slots in the absence of any further match. This search sequence satisfies (with the convention $\tilde{p}_{[|\mathcal{U}|+1]}(c) = p_{[\ell]}$):*

$$\tilde{p}_{[\sigma_1]} \le \cdots \le \tilde{p}_{[\sigma_k]} \le \tilde{v}\left(v, \tilde{p}_{[\ell]}\right) \le \tilde{p}_{[\sigma_k+1]}.$$

Proof. Consider a consumer with search cost c and valuation v having a first match upon inspecting slot $\ell \in \mathcal{I}_n$, and let $\boldsymbol{\sigma} \equiv (\sigma_1, \ldots, \sigma_k) \in \mathcal{S}(\mathcal{I}_k)$, for some $k \in \mathcal{I}_{|\mathcal{U}|}$, denote the consumer's equilibrium search sequence in the absence of any further match.

Step 1. We first note that, if the consumer inspects additional slots after his first match, it does so at most until the next match. To see this, suppose otherwise that the

[65] The proof of Claim A.1 relies on the consumer knowing his valuation, the only motivation for inspecting additional slots then being the quest for a lower price. By contrast, the choice of inspecting previous slots may be motivated by the desire of learning the valuation.

consumer, having inspected some additional slots after slot ℓ, obtains his second match upon inspecting slot $h \in \mathcal{U}$ and nevertheless chooses to inspect one more additional slot $k \in \mathcal{U} \setminus \{h\}$. The optimality of this search strategy requires $p_{[k]} < p_{[h]} < p_{[\ell]}$. Consider now an alternative search strategy in which the consumer inspects slot k first, and then inspects slot h only in the absence of a match. The two search strategies perform equally well if none of the slots produces a match, or if only slot h produces a match. By contrast, whenever slot k produces a match, the alternative strategy dominates the original one, as it leads to the same lowest obtained price (namely, $p_{[k]}$) and saves the cost of inspecting slot h.

The consumer therefore stops searching upon finding a second match and otherwise keeps searching according to the search sequence σ. We can thus write the consumer's value from this search sequence as (with the convention $\mu_{[\sigma_0]} = 0$):

$$V(\boldsymbol{\sigma}; v, c) \equiv \sum_{j=1}^{k} \prod_{i=0}^{j-1} (1 - \mu_{[\sigma_i]}) \mu_{[\sigma_j]} \left[\tilde{v}\left(v, \tilde{p}_{[\ell]}\right) - \tilde{p}_{[\sigma_j]}(c) \right],$$

Step 2. We now show that the consumer inspects slots by increasing order of their cost-adjusted prices, and stops searching when these adjusted prices exceed the maximal value from obtaining a lower price. To see this, suppose otherwise that $\tilde{p}_{[\sigma_j]} > \tilde{p}_{[\sigma_{j+1}]}$ for some $j \in \mathcal{I}_{k-1}$. Letting $\sigma' = (\sigma_1, \ldots, \sigma_{j+1}, \sigma_j, \ldots, \sigma_k)$ denote the alternative search sequence that differs from σ only in the order in which slots σ_j and σ_{j+1} are inspected, we have:

$$V(\boldsymbol{\sigma}'; v, c) - V(\boldsymbol{\sigma}; v, c) = \left[\prod_{i=0}^{j-1} (1 - \mu_{[\sigma_i]}) \right] \mu_{[\sigma_j]} \mu_{[\sigma_{j+1}]} \left[\tilde{p}_{[\sigma_j]}(c) - \tilde{p}_{[\sigma_{j+1}]}(c) \right] > 0.$$

Hence, the consumer would be strictly better off adopting search sequence σ' rather than σ, a contradiction.

Suppose now that $\tilde{v}(v) < \tilde{p}_{[\sigma_k]}$. Letting $\sigma' = (\sigma_1, ..., \sigma_{k-1})$ denote the alternative search sequence that differs from σ only in that the consumer stops searching after inspecting slot σ_{k-1}, we have:

$$V(\boldsymbol{\sigma}'; v, c) - V(\boldsymbol{\sigma}; v, c) = -\prod_{i=0}^{k-1} (1 - \mu_{[\sigma_i]}) \mu_{[\sigma_k]} \left[\tilde{v}\left(v, \tilde{p}_{[\ell]}\right) - \tilde{p}_{[\sigma_k]} \right] > 0.$$

Hence, the consumer would be strictly better off adopting search sequence σ' rather than σ, a contradiction. ■

Building on this, we now show that the firm cannot do better than charging the monopoly price:

Claim A.2 *Generically over c, the equilibrium expected profit generated by the consumer is at most π^m.*

Proof. Generically over c, the thresholds $\left\{\tilde{p}_{[s_i]}\right\}_{s_i \in \mathcal{U}}$ are strictly ranked. It then follows from Claim A.1 that, after a first match: (i) the consumer inspects the slots in $\mathcal{U}$ in the same order, regardless of the realized valuation v, and (ii) the number of inspected slots weakly increases with v. Hence, if the consumer buys at price p when drawing a valuation v, he buys at a weakly lower price $p' \leq p$ whenever drawing a higher valuation $v' > v$, as the higher valuation induces the consumer to inspect at least as many products. For any realized match sequence $\mathbf{m} = (m_{[1]}, ..., m_{|\mathcal{U}|})$, where $m_{[i]} = 1$ if the consumer has a match with the product in slot i, and $m_{[i]} = 0$ otherwise, let $\hat{v}(\mathbf{m})$ denote the infimum of the valuations for which a consumer buys, and $p(v; \mathbf{m})$ denote the price at which a consumer with valuation $v > \hat{v}(\mathbf{m})$ buys. From the above observations, $p(v; \mathbf{m})$ is weakly decreasing in v; hence, $\hat{p}(\mathbf{m}) \equiv \lim_{v \longrightarrow \hat{v}^+} \{p(v; \mathbf{m})\}$ exists and satisfies, for every $v > \hat{v}(\mathbf{m})$:

$$p(v; \mathbf{m}) \leq \hat{p}(\mathbf{m}) \leq \hat{v}(\mathbf{m}),$$

where the first inequality stems from the monotonicity of $p(v; \mathbf{m})$ in v, and the second inequality from the fact that, by construction, $p(v; \mathbf{m}) \leq v$ for any v (above but) arbitrarily close to $\hat{v}$. Hence, the equilibrium expected profit generated by the consumer is at most:

$$\int_{\hat{v}(\mathbf{m})}^{+\infty} p(v; \mathbf{m})\, dF(v) \leq \int_{\hat{v}(\mathbf{m})}^{+\infty} \hat{v}(\mathbf{m})\, dF(v) = [1 - F(\hat{v}(\mathbf{m}))]\, \hat{v}(\mathbf{m}) \leq \pi^m,$$

where the last inequality follows from $\pi^m = \max_p \{[1 - F(p)]\, p\}$. ■

The next claim concludes the proof:

Claim A.3 *Under passive beliefs, in any pure positioning equilibrium, the firm charges the monopoly price p^m on all offered products.*

Proof. Suppose otherwise. By deviating to an out-of-equilibrium price arbitrarily close to the monopoly price, under passive beliefs the firm does not affect consumers' search behavior before finding a match, and thus does not affect either their probabilities of having at least one match, but makes arbitrary close to the monopoly profit on those who do have at least one match, a contradiction. ■

B Proof of Proposition 1 (pure positioning)

Fix the product portfolio size $n \in \mathcal{I}_N$ adopted in the first stage, and consider a candidate equilibrium in which, with probability 1, the firm selects the product portfolio $\mathcal{I} = \{i_1, \ldots, i_n\}$, with the convention that $\mu_{i_1} > \cdots > \mu_{i_n}$, and allocates product $[k] \in \mathcal{I}$ to slot $k \in \mathcal{I}_n \equiv \{1, \ldots, n\}$. From Lemma 1, it offers the selected products at the monopoly price p^m. Hence, the slots differ in the probability of delivering a match but, conditional on a match, they all offer the same value, $v - p^m$, where v is the consumer's realized match-valuation. It follows that active consumers inspect slots by decreasing order of the popularity of their products (i.e., they inspect first the slot holding product i_1, and so forth) and stop searching once they have a first match (as any further search would at best yield the same net value); furthermore, in the absence of a prior match, consumers inspect slot k only if their search cost does not exceed $\mu_{[k]} s^m$.

The firm's expected profit is thus equal to (with the convention $\lambda_{i_0} = 1$):

$$\sum_{k \in I_n} \left(\prod_{h=0}^{k-1} \lambda_{i_h} \right) G\left(\mu_{i_k} s^m\right) \mu_{i_k} \pi^m.$$

It is therefore clearly optimal for the firm to select the n most popular products and to position the more popular ones in the first slots inspected by consumers. Indeed, as consumers never observe the identity of the products, replacing any selected product $i \notin \mathcal{I}_n$ with a product $j \in \mathcal{I}_n$ does not affect consumers' search and purchasing behavior, and strictly increases the probability of a match for those consumers who inspect the slot holding product i. Likewise, as the first inspected slots attract more consumers, positioning the most popular product in the first inspected slot, and so forth, maximizes the overall probability of a match. The resulting profit is $\Pi^P(n)$, as described in the proposition.

C Proof of Corollary 1 (efficiency)

Consider a candidate equilibrium in which a pure-strategy equilibrium is played in every subgame of stage 2. From Lemma 1, for any given product portfolio size n adopted in stage 1, the expected profit from the continuation subgame is given by $\Pi^P(n)$, where by construction:

$$\Pi^P(n+1) - \Pi^P(n) = \Lambda_n G\left(c_{n+1}^P\right) \mu_{n+1} \pi^m > 0.$$

It follows that the firm offers all available products, which, together with Proposition 1, establishes uniqueness. Conversely, offering all available products, and playing the unique pure-strategy equilibrium in every continuation subgame, does constitute an equilibrium of the overall game.

D Proof of Lemma 2 (intensive search)

Fix the product portfolio size $n \in \mathcal{I}_N \diagdown \{1\}$ adopted in the first stage, and consider a candidate equilibrium in which the firm selects the product portfolio $\mathcal{I} = \{i_1, \ldots, i_n\}$, with the convention that $\mu_{i_1} > \cdots > \mu_{i_n}$, offers the selected products at the monopoly price p^m, and uniformly randomizes over their positions. Along the equilibrium path, prices convey no information about the identity of the inspected products; by contrast, whether a match occurs leads consumers to revise their beliefs about the inspected products and, thus, about the remaining ones. For $k \in \mathcal{I}_n$, let $m_k \in \{0, 1\}$ denote whether a match occurs ($m_k = 1$) or not ($m_k = 0$) at the k^{th} inspection, $M_k \equiv \Sigma_{h=1}^{h=k} m_h$ denote the number of matches at the first k inspections (with the convention $M_0 = 0$) and

$$\alpha_k \equiv \Pr\left[m_k = 1 \mid M_{k-1} = 0\right]$$

denote the probability of a match at the k^{th} inspection, conditional on not having had any match at the previous $k-1$ inspections. Furthermore, for $k \in \mathcal{I}_n \diagdown \{1\}$, let

$$\hat{\alpha}_k \equiv \Pr\left[m_k = 1 \mid M_{k-2} = 0 \text{ and } m_{k-1} = 1\right]$$

denote the probability of a match at the k^{th} inspection, conditional on not having had any match at the previous $k-2$ inspections and having had a match at the $(k-1)^{\text{th}}$ inspection.

We first show that observing a match makes consumers more pessimistic about future matches:

Lemma D.1 (a match brings bad news about future matches) *For $k \in \mathcal{I}_n \diagdown \{1\}$, $\alpha_k > \hat{\alpha}_k$.*

Proof. For $k \in \mathcal{I}_n$, let P_k denote the sub-portfolio of products visited at the first k inspections, and

$$\mathcal{P}_k \equiv \{\mathcal{J} \subseteq \mathcal{I} \mid |\mathcal{J}| = k\}$$

denote the set of such sub-portfolios (with the convention $P_0 = \mathcal{P}_0 = \varnothing$). For $k \in$

$\mathcal{I}_n \diagdown \{1\}$, we have:

$$\alpha_k = \sum_{\mathcal{J}_{k-2} \in \mathcal{P}_{k-2}} \Pr\left[P_{k-2} = \mathcal{J}_{k-2}\right] \beta_k\left(\mathcal{J}_{k-2}\right),$$
$$\hat{\alpha}_k = \sum_{\mathcal{J}_{k-2} \in \mathcal{P}_{k-2}} \Pr\left[P_{k-2} = \mathcal{J}_{k-2}\right] \hat{\beta}_k\left(\mathcal{J}_{k-2}\right),$$

where

$$\begin{aligned}
\beta_k\left(\mathcal{J}_{k-2}\right) &= \Pr\left[m_k = 1 \mid M_{k-1} = 0 \text{ and } P_{k-2} = \mathcal{J}_{k-2}\right], \\
\hat{\beta}_k\left(\mathcal{J}_{k-2}\right) &= \Pr\left[m_k = 1 \mid M_{k-2} = 0,\ m_{k-1} = 1 \text{ and } P_{k-2} = \mathcal{J}_{k-2}\right].
\end{aligned}$$

Fix $\mathcal{J}_{k-2}$, and let $\hat{n} \equiv \left|\mathcal{I} \diagdown \mathcal{J}_{k-2}\right| (= n + 2 - k)$ denote the size of the complement sub-portfolio $\mathcal{I} \diagdown \mathcal{J}_{k-2}$, $\hat{\mu}_i$ denote the popularity of its i^{th} most popular product, for $i \in I_{\hat{n}} \equiv \{1, ..., \hat{n}\}$, and $\hat{\mu} \equiv (\sum_{j \in I_{\hat{n}}} \hat{\mu}_j)/\hat{n}$ denote the average popularity. We have:

$$\begin{aligned}
\beta_k\left(\mathcal{J}_{k-2}\right) - \hat{\beta}_k\left(\mathcal{J}_{k-2}\right) &= \frac{\sum_{i \in I_{\hat{n}}} \sum_{j \in I_{\hat{n}} \setminus \{i\}} \left(1 - \hat{\mu}_j\right) \hat{\mu}_i}{\sum_{h \in I_{\hat{n}}} \sum_{j \in I_{\hat{n}} \setminus \{h\}} \left(1 - \hat{\mu}_j\right)} - \frac{\sum_{i \in I_{\hat{n}}} \sum_{j \in I_{\hat{n}} \setminus \{i\}} \hat{\mu}_j \hat{\mu}_i}{\sum_{h \in I_{\hat{n}}} \sum_{j \in I_{\hat{n}} \setminus \{h\}} \hat{\mu}_j} \\
&= \sum_{i \in I_{\hat{n}}} \left[\frac{\hat{n}\left(1 - \hat{\mu}\right) - \left(1 - \hat{\mu}_i\right)}{\left(\hat{n} - 1\right) \hat{n}\left(1 - \hat{\mu}\right)} - \frac{\hat{n}\hat{\mu} - \hat{\mu}_i}{\left(\hat{n} - 1\right) \hat{n}\hat{\mu}}\right] \hat{\mu}_i \\
&= \sum_{i \in I_{\hat{n}}} \frac{\left(\mu_i - \hat{\mu}\right) \hat{\mu}_i}{\hat{n}\left(\hat{n} - 1\right) \hat{\mu}\left(1 - \hat{\mu}\right)}.
\end{aligned}$$

As $\sum_{i \in I_{\hat{n}}} \left(\mu_i - \hat{\mu}\right) \mu_i = \sum_{i \in I_{\hat{n}}} \left(\mu_i - \hat{\mu}\right)^2 > 0$, it follows that $\beta_k\left(\mathcal{J}_{k-2}\right) > \hat{\beta}_k\left(\mathcal{J}_{k-2}\right)$ for any $\mathcal{J}_{k-2} \in \mathcal{P}_{k-2}$, which in turn implies $\alpha_k > \hat{\alpha}_k$. ■

Given the information available after the first $k - 1$ inspections, the expected probability of a match is the same for all future inspections. Hence:

$$\begin{aligned}
\alpha_k &= \Pr\left[m_k = 1 \mid M_{k-1} = 0\right] \\
&= \Pr\left[m_{k+1} = 1 \mid M_{k-1} = 0\right] \\
&= \Pr\left[m_k = 1 \mid M_{k-1} = 0\right] \Pr\left[m_{k+1} = 1 \mid M_{k-1} = 0 \text{ and } m_k = 1\right] \\
&\quad + \Pr\left[m_k = 0 \mid M_{k-1} = 0\right] \Pr\left[m_{k+1} = 1 \mid M_{k-1} = 0 \text{ and } m_k = 0\right] \\
&= \alpha_k \hat{\alpha}_{k+1} + \left(1 - \alpha_k\right) \alpha_{k+1} \\
&= \alpha_{k+1} + \alpha_k\left(\hat{\alpha}_{k+1} - \alpha_{k+1}\right) \\
&< \alpha_{k+1},
\end{aligned}$$

where the inequality follows from Lemma D.1.

Hence, if a consumer finds it optimal to make the first inspection, he will optimally continue to inspect all products until he has found a match (and learned his match-conditional valuation). To see this, suppose by way of contradiction, that the consumer finds it optimal to stop searching after $k < n$ inspections. This requires $\alpha_{k+1} \leq c$, as otherwise the consumer would have an incentive to search at least one more time. But the consumer must also find it optimal to conduct the k^{th} inspection, *knowing* that it will be the last one, which in turn requires $\alpha_k \geq c$. We thus have $\alpha_k \geq c \geq \alpha_{k+1}$, a contradiction.

E Proof of Proposition 2 random positioning)

Fix the product portfolio size $n \in \mathcal{I}_N$ adopted in the first stage, and consider a candidate equilibrium in which the firm selects the product portfolio $\mathcal{I} = \{i_1, \ldots, i_n\}$, with the convention that $\mu_{i_1} > \cdots > \mu_{i_n}$, offers the selected products at the monopoly price p^m, and uniformly randomizes over their positions. From Lemma 2, all active consumers keep searching until finding a match. It follows that, in equilibrium, the firm necessarily offers the n most popular products: starting from any $\mathcal{I} \neq \mathcal{I}_n$, replacing $\mathcal{I}$ with $\mathcal{I}_n$ would neither affect the number of active consumers nor their search behavior, as consumers never observe the identity of the products, but would strictly increase their probability of having a match, making the deviation strictly profitable. Thus, in equilibrium, $\mathcal{I} = \mathcal{I}_n$. A consumer therefore engages in search if the expected probability of a match times the expected surplus conditional on a match, $\mathrm{M}_n^R s^m$, (weakly) exceeds the expected search cost, which is equal to the expected number of inspections times the consumer's cost per search, $\Gamma_n c$, i.e., if and only if $c \leq c_n^R = \mathrm{M}_n^R s^m / \Gamma_n$. Note that all active consumers are indifferent between all possible search sequences in $\mathcal{S}(\mathcal{I}_n)$.

We now show that, conversely, offering the n most popular products at the monopoly price and randomly positioning them, together with any search behavior such that (i) consumers start searching if and only $c \leq c_n^R$, and (ii) active consumers keep searching until finding a match, constitutes a PBE with passive beliefs. We have already checked that the described search behavior constitutes a best response for consumers. Suppose now that the firm deviates and selects the product portfolio $\mathcal{I}^d$, positions the selected products (possibly randomly) and, for any $(h, k) \in \mathcal{I}^d \times \mathcal{I}_n$ charges a price p_k^h for product h whenever it is assigned to slot k. As consumers cannot observe this deviation before their initial search decision, they still start searching if and only if $c \leq c_n^R$. Furthermore, under passive beliefs, the deviation does not affect active consumers' search behavior in the absence of a match: they still expect the k^{th} inspection to deliver a match with the equilibrium probability α_k, which increases with k, and expect to encounter the monopoly price in the remaining slots; hence, they keep searching until finding a match. Consider now a consumer with search cost who has a first match on the k^{th} inspection, where he learns his valuation v and faces the price p_k (equal to p_k^h for some product h). If $p_k \leq p^m$, then the consumer stops searching and buys if and only if $v \geq p_k$. If instead $p_k > p^m$, the consumer may choose to keep searching in the hope of encountering the monopoly price. Specifically, the consumer keeps searching if $\hat{\Gamma}_k c < \hat{M}_k \left[\min\{v, p_k\} - p^m\right]$, or

$$\hat{v}_k(v) \equiv \min\{v, p_k\} > \hat{p}_k(c) \equiv p^m + \frac{\hat{\Gamma}_k}{\hat{M}_k} c,$$

where $\hat{M}_k$ and $\hat{\Gamma}_k$, respectively, denote the probability of having at least one match in the remaining $n - k$ uninspected slots, and the expected number of inspections until a first match in these slots, conditional on having had a first match on the k^{th} inspection. As $\hat{v}(v)$ and $\hat{p}(c)$ are respectively increasing in v and in c, it follows that consumers with greater value (i.e., a higher v or a lower c) are more likely to keep searching. Furthermore, among those who do keep searching, they do so until finding another match (due to increasing optimism) and, upon having such a match on the $(k+h)^{\text{th}}$ inspection and facing a price p_{k+h}, they will keep searching if

$$\hat{v}_{k,h}(v) \equiv \min\{v, p_k, p_{k+h}\} > \hat{p}_{k,h}(c) \equiv p^m + \frac{\hat{\Gamma}_{k,h}}{\hat{M}_{k,h}} c,$$

where $\hat{M}_{k,h}$ and $\hat{\Gamma}_{k,h}$, respectively, denote the probability of having at least one match in the remaining $n - k - h$ uninspected slots, and the expected number of inspections until a first match in these slots, conditional on having had a first match on the k^{th} inspection and a second one on the $(k+h)^{\text{th}}$ inspection. Hence, consumers with a higher v or a lower c are again more likely to keep searching; and so on. It follows that the same logic as that in the proof of Lemma 1 applies here: as consumers with greater values are more prone to keep searching, they are also more likely to encounter lower prices (even among unexpected ones), which rules out any scope for profitable price discrimination. As a result, the firm cannot obtain more profit than by charging the price p^m for all products. It readily follows that it cannot benefit either by offering a different product portfolio than $\mathcal{I}_n$ and, given consumers' search behavior, cannot benefit either from adopting a non-uniformly random positioning policy.

F Proof of Corollary 2 (choice overload)

The expected surplus of a consumer with search cost c when facing a product line of size n can be expressed as

$$S(\mathcal{I}_n; c) \equiv \frac{\sum_{i=1}^{n} \mu_i}{n} s^m - c + \sum_{i=1}^{n} \left\{ \frac{1-\mu_i}{n} \left[\mathrm{M}(\mathcal{I}_n \backslash \{i\}) s^m - \Gamma(\mathcal{I}_n \backslash \{i\}) c \right] \right\},$$

where the first two terms correspond to the expected surplus from the first inspection, and each term in curly brackets corresponds to the expected surplus from subsequent inspections, conditional on the first inspected product having been product i.

As $\partial S(c; n)/\partial c < 0$, the marginal type c_n^R is uniquely determined by $S(\mathcal{I}_n; c_n^R) = 0$ and, to prove that $c_n^R > c_{n+1}^R$, we only need to show that $S(\mathcal{I}_{n+1}; c_n^R) < 0$. We have:

$$\begin{aligned} S(c_n^R; n+1) &= \frac{\sum_{i=1}^{n+1} \mu_i}{n+1} s^m - c_n^R + \frac{1-\mu_{n+1}}{n+1} \left[\mathrm{M}(\mathcal{I}_n) s^m - \Gamma(\mathcal{I}_n) c_n^R \right] \\ &\quad + \sum_{i=1}^{n} \left\{ \frac{1-\mu_i}{n+1} \left[\mathrm{M}(\mathcal{I}_{n+1} \backslash \{i\}) s^m - \Gamma(\mathcal{I}_{n+1} \backslash \{i\}) c_n^R \right] \right\} \\ &= \frac{\sum_{i=1}^{n+1} \mu_i}{n+1} s^m - c_n^R + \sum_{i=1}^{n} \left\{ \frac{1-\mu_i}{n+1} \left[\mathrm{M}(\mathcal{I}_{n+1} \backslash \{i\}) s^m - \Gamma(\mathcal{I}_{n+1} \backslash \{i\}) c_n^R \right] \right\}, \end{aligned} \tag{10}$$

where the last equality follows from the definition of c_n^R. To conclude the proof, it suffices to note that, on the RHS of (10): (a) the sum of the first two terms is strictly negative; and (b) each term in curly brackets is strictly negative as well.

To see (a), recall first that $(\sum_{i=1}^{n} \mu_i)/n$ is strictly decreasing in n, implying that

$$\frac{\sum_{i=1}^{n+1} \mu_i}{n+1} s^m - c_n^R < \frac{\sum_{i=1}^{n} \mu_i}{n} s^m - c_n^R.$$

Furthermore, c_n^R is such that

$$\left[\alpha_1(\mathcal{I}_n) s^m - c_n^R\right] + [1-\alpha_1(\mathcal{I}_n)][\alpha_2(\mathcal{I}_n) s^m - \underline{c}_n] + \ldots + \prod_{i=1}^{i=n-1} [1-\alpha_i(\mathcal{I}_n)][\alpha_n(\mathcal{I}_n) s^m - c_n^R] = 0,$$

where $\alpha_k(\mathcal{I}_n)$ is the probability of a match at the k^{th} inspection, conditional on no prior match and, from the proof of Lemma 2, is strictly increasing in k. Hence, the first term (in brackets) on the LHS of the last equation is non-positive (otherwise, the sum of the terms would be strictly positive). The conclusion then follows from the observation that $\alpha_1(\mathcal{I}_n) = (\sum_{i=1}^{n} \mu_i)/n$.

Finally, (b) follows from the definition of c_n^R, in conjunction with the observation that, for any $i \leq n$, $\mathrm{M}(\mathcal{I}_{n+1} \backslash \{i\}) < \mathrm{M}(\mathcal{I}_n)$ and $\Gamma(\mathcal{I}_{n+1} \backslash \{i\}) > \Gamma(\mathcal{I}_n)$: this can be seen from (1) and (2), using $\mu_i > \mu_{n+1}$.

G Proof of Lemma 3 (steering)

Consider a candidate equilibrium in which the firm chooses to offer n products at the monopoly price, and all consumers inspect the slots in the same order (some consumers may however remain inactive, and some active consumers may stop searching before others, even in the absence of a match). We stick to the convention that slots are labelled according to consumers' search sequence.

Along the equilibrium path, some consumers (including all those with a cost $c < \mu_N s^m$) start searching and are willing to inspect all slots until finding a match; furthermore, all active consumers stop searching whenever they have a match, as they expect all other products to offer the same value at the same price. Consider now the behavior of active consumers in the absence of any match, and let $\mathcal{N} = \left(n_1, \ldots, n_{|\mathcal{N}|}\right) \subseteq \mathcal{I}_N$ denote the set of slots after which some active consumers stop searching. That is, in the absence of a match, all active consumers inspect the first n_1 slots; some consumers then stop searching, whereas all those who inspect slot $n_1 + 1$ are willing to inspect all slots up to slot n_2; and so forth. By construction, $n_{|\mathcal{N}|} = n$. For any $k \in \mathcal{I}_{|\mathcal{N}|}$, define $\mathcal{B}_k \equiv \{n_{k-1} + 1, \ldots, n_k\}$ (with the convention $n_0 = 0$) and let $\mathcal{C}_k$ denote the set of consumers willing to visit the slots in block $\mathcal{B}_k$. Finally, define $\boldsymbol{\mathcal{B}} \equiv (\mathcal{B}_1, \ldots, \mathcal{B}_{|\mathcal{N}|})$; by construction, $|\boldsymbol{\mathcal{B}}| = |\mathcal{N}|$ and $\mathcal{C}_1 \supset \cdots \supset \mathcal{C}_{|\boldsymbol{\mathcal{B}}|}$.

Obviously, the firm has a strict incentive to offer the n most popular products. To see this, suppose that it picks with positive probability a product placement $\boldsymbol{\sigma} \in \mathcal{S}(\mathcal{I}_n)$ that does not offer product $i \leq n$; it must therefore be the case that some product $j > n$ is assigned to one of the slots, say ℓ. But then, replacing product j with product i would (i) be undetected by consumers, and thus have no impact on their behavior, and (ii) increase the firm's profit by increasing the probability of a match for all consumers inspecting slot ℓ. Hence, without loss of generality, we can assume that the firm chooses to offer the product portfolio $\mathcal{I}_n$.

Likewise, for any $k \in \mathcal{I}_{|\boldsymbol{\mathcal{B}}|}$ the firm assigns the products from $\mathcal{B}_k$ exclusively to the slots in $\mathcal{B}_k$. To see this, suppose for instance that the firm picks with positive probability a product placement $\boldsymbol{\sigma} \in \mathcal{S}(\mathcal{I}_n)$ that assigns product $i \leq n_1$ to slot $h > n_1$; it must therefore be the case that some product $j > n_1$ is assigned to some slot $\ell \leq n_1$. But then, swapping the positions of products i and j would again be undetected by consumers and increase the firm's profit, by increasing the probability of a match for the consumers in $\mathcal{C}_1 \diagdown \mathcal{C}_2$, who are only willing to inspect the first n_1

slots. Thus, without loss of generality, we can assume that the firm assigns the n_1 most popular products to the first n_1 slots; that is, it assigns all products from $\mathcal{B}_1$ to slots in $\mathcal{B}_1$. Iterating the argument shows that, more generally, the firm assigns the products in $\mathcal{B}_k$ exclusively to the slots in $\mathcal{B}_k$.

By construction, for any $k \in \mathcal{I}_{|\boldsymbol{\mathcal{B}}|}$, a consumer who inspects the first slot in $\mathcal{B}_k$ had no match when inspecting the previous slots, which occurs with probability $\Lambda\left(\bar{\mathcal{B}}_{k-1}\right)$; furthermore, the consumer is then willing to inspect all products in $\mathcal{B}_k$ until finding a match, and thus obtains a match with one of these products with probability $\mathrm{M}\left(\mathcal{B}_k\right)$. The firm's profit can therefore be expressed as:

$$\sum_{k=1}^{|\boldsymbol{\mathcal{B}}|} \mathcal{C}_k \Lambda(\bar{\mathcal{B}}_{k-1}) \mathrm{M}(\mathcal{B}_k) \pi^m$$

Finally, as a consumer's expected net surplus is strictly decreasing in the search cost c, there exists $\hat{c}_k$ such that all consumers with $c \leq \hat{c}_k$ are willing to inspect the block $\mathcal{B}_k$; furthermore, by construction $\hat{c}_1 > \cdots > \hat{c}_{|\boldsymbol{\mathcal{B}}|}$, as some consumers stop searching after having inspected the slots in any given block $\mathcal{B}_k$. The number of consumers willing to inspect the slots is therefore given by $\mathcal{C}_k = G\left(\hat{c}_k\right)$.

To conclude the proof, consider a given block $\mathcal{B}_k$ such that $|\mathcal{B}_k| > 1$. If in equilibrium the firm were to adopt a deterministic positioning strategy for that block, consumers would inspect products in decreasing order of popularity; consumers with $c \in \left(\mu_{n_{k-1}+2}, \mu_{n_{k-1}+1}\right)$ would then be willing to inspect slot $n_{k-1}+1$ but would stop searching afterwards, even in the absence of a match, contradicting the working assumption that, within $\mathcal{B}_k$, active consumers keep inspecting all slots until finding a match. It follows that the firm must adopt a random positioning strategy for $\mathcal{B}_k$.

H Proof of Lemma 4 (worst intensive search equilibrium)

Consider a given block $\mathcal{B}_k$ and suppose first that the firm randomizes uniformly its positioning strategy over that block. From Lemma 2, a consumer who starts inspecting $\mathcal{B}_k$ keeps doing so until finding a match. The expected surplus generated by $\mathcal{B}_k$ is thus equal to $S^R\left(\mathcal{B}_k; c\right) \equiv \mathrm{M}\left(\mathcal{B}_k\right) s^m - \Gamma\left(\mathcal{B}_k\right) c$, where the superscript R stands for *Random positioning*, c is the consumer's search cost, and $\mathrm{M}\left(\cdot\right)$ and $\Gamma\left(\cdot\right)$ are respectively given by (1) and (2).

Consider now a candidate equilibrium in which the firm adopts a different positioning strategy over $\mathcal{B}_k$, but consumers who start inspecting $\mathcal{B}_k$ still keep doing so until finding a match. If a consumer with search cost c were to uniformly randomize over his search sequence (and keep searching in $\mathcal{B}_k$ until finding a match), his expected surplus would be independent of positioning, and thus equal to $S^R(\mathcal{B}_k; c)$. It follows that, in any equilibrium in which consumers who inspect $\mathcal{B}_k$ never stop in the absence of a match, consumers can always secure the maximal expected surplus achievable under random positioning. By revealed preferences, if consumers favor a different search sequence, they must obtain a higher expected surplus. This, in turn, boosts consumer participation and, thus, the firm's profit.

I Proof of Proposition 3 (controlled search)

We suppose here that slots can only be inspected in increasing order—consumers remaining free to start or stop searching. We fix $\mathcal{B} = \left\{h_1, \ldots, h_{|\mathcal{B}|}\right\} \subseteq \mathcal{I}_N$ such that $|\mathcal{B}| > 1$ and $\mu_{h_1} > \cdots > \mu_{h_{|\mathcal{B}|}}$, and seek to maximize consumer participation among intensive search equilibria—i.e., equilibria in which the firm charges the monopoly price for all products it offers, and active consumers keep inspecting slots (in increasing order) until finding a match.

For any such equilibrium and any $k \in \mathcal{I}_{|\mathcal{B}|}$, let

$$\beta_k \equiv \Pr\left[m_1 = \cdots = m_k = 0\right]$$

denote the probability that the first k slots do not deliver a match; by construction, $\beta_{|\mathcal{B}|} = \Lambda(\mathcal{B})$. If a consumer with search cost c inspects up to k slots until finding a match, his net expected payoff is equal to $V_k(c)$, given by (5). As $V_k(c)$ is decreasing in c, participating consumers are those with $c \leq \hat{c}$, where $\hat{c}$ is the largest cost satisfying (IR) (with the convention $\beta_0 = 1$). Furthermore, a consumer with cost c is indeed willing to inspect all slots until finding a match if, for every $k \in \mathcal{I}_{|\mathcal{B}|-1}$, $V_{|\mathcal{B}|}(c) \geq V_k(c)$, which amounts to:

$$\left[\beta_k - \Lambda(\mathcal{B})\right] s^m \geq \left(\sum_{i=k}^{|\mathcal{B}|-1} \beta_i\right) c.$$

As the right-hand side is increasing in c, whereas the left-hand side does not depend on c, this condition is most stringent for the marginal consumer, for which it amounts to (IC_k).

Summing-up, for any equilibrium featuring intensive search, there exists $\boldsymbol{\beta} = (\beta_k)_{k \in \mathcal{I}_{|\mathcal{B}|-1}}$ and a marginal consumer characterized by a search cost $\hat{c}$, satisfying (IR) and $(IC_k)_{k \in \mathcal{I}_{|\mathcal{B}|-1}}$. All these equilibria yield the same probability of a match, $\mathrm{M}(\mathcal{B})$, but differ in consumer participation; among them, the best one maximizes participation. To provide an upper bound on achievable participation, we now consider a relaxed problem, where any $\boldsymbol{\beta} = (\beta_k)_{k \in \mathcal{I}_{|\mathcal{B}|-1}}$ and $\hat{c}$ satisfying (IR) and $(IC_k)_{k \in \mathcal{I}_{|\mathcal{B}|-1}}$

are supposed to be available—and where no sign or boundary condition on either $\hat{c}$ or $\boldsymbol{\beta}$ is imposed:

$$\max_{(\hat{c},\boldsymbol{\beta})\in\mathbb{R}^{|\mathcal{B}|}} \quad \hat{c}, \qquad \text{s.t.} \quad (IR) \ \text{ and } \ (IC_k)_{k\in\mathcal{I}_{|\mathcal{B}|-1}}. \tag{$\mathcal{P}$}$$

The following lemma characterizes the solution to this relaxed problem by showing that all constraints are binding:

Lemma I.1 (maximal participation) *The solution to the relaxed problem* ($\mathcal{P}$) *is such that*

$$\hat{c} = \hat{c}(\mathcal{B}) \equiv \left[1 - \bar{\lambda}(\mathcal{B})\right] s^m$$

and, for $k \in \mathcal{I}_{|\mathcal{B}|-1}$*:*

$$\beta_k = \hat{\beta}_k(\mathcal{B}) \equiv \left[\bar{\lambda}(\mathcal{B})\right]^k.$$

Proof. Let $\nu_0 \geq 0$ denote the Lagrangian multiplier associated with (IR), $\nu_k \geq 0$ the multiplier associated with (IC_k), for $k \in \mathcal{I}_{|\mathcal{B}|-1}$, and define $\boldsymbol{\nu} = \left(\nu_0, \nu_1, \ldots, \nu_{|\mathcal{B}|-1}\right)$. The Lagrangian is equal to:

$$L(\hat{c}, \boldsymbol{\beta}; \boldsymbol{\nu}) \equiv \hat{c} + \nu_0 \left[\mathrm{M}(\mathcal{B}) s^m - \left(\sum_{i=0}^{|\mathcal{B}|-1} \beta_i\right) \hat{c}\right] + \sum_{k=1}^{|\mathcal{B}|-1} \nu_k \left[(\beta_k - \Lambda(\mathcal{B})) s^m - \left(\sum_{i=k}^{|\mathcal{B}|-1} \beta_i\right) \hat{c}\right]$$

The first-order condition with respect to $\hat{c}$ yields:

$$1 = \nu_0 \sum_{i=0}^{|\mathcal{B}|-1} \beta_i + \sum_{k=1}^{|\mathcal{B}|-1} \left[\nu_k \left(\sum_{i=k}^{|\mathcal{B}|-1} \beta_i\right)\right] \tag{FC_0}$$

whereas the first-order condition with respect to β_k yields:

$$\nu_k (s^m - \hat{c}) = \left(\nu_0 + \sum_{i=1}^{k-1} \nu_i\right) \hat{c}. \tag{FC_k}$$

Any consumer with $c > \mu_N s^m$ is always active, whereas no consumer with $c > \mu_1 s^m$ is ever active; it follows that $0 < \hat{c} < s^m$. If $\nu_0 = 0$, iteratively applying $(FC_k)_{k\in\mathcal{I}_{|\mathcal{B}|-1}}$ then yields $\nu_1 = \cdots \nu_{|\mathcal{B}|-1} = 0$,[66] contradicting (FC_0); hence, $\nu_0 > 0$. This, in turn, implies that, for every $k \in \mathcal{I}_{|\mathcal{B}|-1}$, the right-hand side of (FC_k) is positive, as $\hat{c} > 0$ and $\nu_0 + \sum_{i=1}^{k-1} \nu_i \geq \nu_0 > 0$; hence, $\nu_k (s^m - \hat{c}) > 0$, which, together with $s^m > \hat{c}$, yields $\nu_k > 0$. All constraints are therefore binding, that is: $V_1(\hat{c}) = \cdots = V_{|\mathcal{B}|}(\hat{c}) = 0$. In

[66] Together with $s^m > \hat{c}$, (FC_1) yields $\nu_1 = 0$; (FC_2) then yields $\nu_2 = 0$, and so on.

particular, $V_1(\hat{c}) = 0$ yields

$$\hat{c} = (1 - \beta_1) s^m \iff \beta_1 = 1 - \frac{\hat{c}}{s^m},$$

and $V_k(\hat{c}) = V_{k+1}(\hat{c})$ yields, for $k \in \mathcal{I}_{|\mathcal{B}|-1}$:

$$\beta_k \hat{c} = (\beta_k - \beta_{k+1}) s^m \iff \beta_{k+1} = \left(1 - \frac{\hat{c}}{s^m}\right) \beta_k.$$

Combining these conditions leads to, for $k \in \mathcal{I}_{|\mathcal{B}|}$:

$$\beta_k = \left(1 - \frac{\hat{c}}{s^m}\right)^k. \tag{11}$$

Finally, plugging these expressions in $V_{|\mathcal{B}|}(\hat{c}) = 0$ then yields (using $\mathrm{M}(\mathcal{B}) = 1 - \Lambda(\mathcal{B})$):

$$\begin{aligned}
[1 - \Lambda(\mathcal{B})] s^m &= \left(\sum_{i=0}^{|\mathcal{B}|-1} \beta_i\right) \hat{c} \\
\iff & \Lambda(\mathcal{B}) = 1 - \frac{\hat{c}}{s^m} \sum_{i=0}^{|\mathcal{B}|-1} \left(1 - \frac{\hat{c}}{s^m}\right)^i = \left(1 - \frac{\hat{c}}{s^m}\right)^{|\mathcal{B}|},
\end{aligned}$$

leading to (using $\sqrt[|\mathcal{B}|]{\Lambda(\mathcal{B})} = \bar{\lambda}(\mathcal{B})$):

$$\hat{c} = \left[1 - \bar{\lambda}(\mathcal{B})\right] s^m$$

and, using (11):

$$\beta_k = \left[\bar{\lambda}(\mathcal{B})\right]^k.$$

■

J Proof of Proposition 4 (best intensive search)

We now show that the positioning strategy $\hat{\boldsymbol{\rho}}(\mathcal{B}) = (\hat{\rho}_{\boldsymbol{\sigma}}(\mathcal{B}))_{\boldsymbol{\sigma} \in S(\mathcal{B})}$, given by (6), implements the desired no-match probabilities characterized by Lemma I.1:

Lemma J.2 (implementation) *If the firm assigns products to slots according to $\hat{\boldsymbol{\rho}}(\mathcal{B})$, then $\beta_k = \hat{\beta}_k(\mathcal{B})$ for every $k \in I_{|\mathcal{B}|}$.*

Proof. We have:

$$\begin{aligned}
\beta_k &= \sum_{\boldsymbol{\sigma}\in S(\mathcal{B})} \hat{\rho}_{\boldsymbol{\sigma}}(\mathcal{B}) \prod_{i\in\mathcal{I}_k} \lambda_{\sigma_i} \\
&= \sum_{\boldsymbol{\sigma}\in S(\mathcal{B})} \frac{\prod_{i\in I_{|\mathcal{B}|}} \lambda_{\sigma_i}^{\frac{i-1}{|\mathcal{B}|}}}{\sum_{\tilde{\boldsymbol{\sigma}}\in S(\mathcal{B})} \prod_{i\in I_{|\mathcal{B}|}} \lambda_{\tilde{\sigma}_i}^{\frac{i-1}{|\mathcal{B}|}}} \prod_{i\in\mathcal{I}_k} \lambda_{\sigma_i} \\
&= \prod_{i\in I_{|\mathcal{B}|}} \lambda_{\sigma_i}^{\frac{k}{|\mathcal{B}|}} \sum_{\boldsymbol{\sigma}\in S(\mathcal{B})} \frac{\prod_{i=1}^{i=k} \lambda_{\sigma_i}^{\frac{(i+|\mathcal{B}|-k)-1}{|\mathcal{B}|}} \prod_{i=k+1}^{i=|\mathcal{B}|} \lambda_{\sigma_i}^{\frac{(i-k)-1}{|\mathcal{B}|}}}{\sum_{\tilde{\boldsymbol{\sigma}}\in S(\mathcal{B})} \prod_{i\in I_{|\mathcal{B}|}} \lambda_{\tilde{\sigma}_i}^{\frac{i-1}{|\mathcal{B}|}}} \\
&= \left[\bar{\lambda}(\mathcal{B})\right]^k \frac{\sum_{\boldsymbol{\sigma}\in S(\mathcal{B})} \prod_{i\in I_{|\mathcal{B}|}} \lambda_{\sigma_{g_k(i)}}^{\frac{i-1}{|\mathcal{B}|}}}{\sum_{\tilde{\boldsymbol{\sigma}}\in S(\mathcal{B})} \prod_{i\in I_{|\mathcal{B}|}} \lambda_{\tilde{\sigma}_i}^{\frac{i-1}{|\mathcal{B}|}}},
\end{aligned}$$

where g_k is the permutation of $\mathcal{I}_{|\mathcal{B}|}$ such that:

$$g_k(i) \equiv \begin{cases} i+|\mathcal{B}|-k & \text{if} \quad i \leq k, \\ i-k & \text{if} \quad i > k. \end{cases}$$

As g_k is a bijection from $S\left(\mathcal{I}_{|\mathcal{B}|}\right)$ to $S\left(\mathcal{I}_{|\mathcal{B}|}\right)$, the function $\boldsymbol{\sigma} \longrightarrow \boldsymbol{\sigma} \circ g_k$ (transforming $\boldsymbol{\sigma} = (\sigma_i)_{i\in\mathcal{I}_{|\mathcal{B}|}}$ into $\boldsymbol{\sigma}\circ g_k = \left(\sigma_{g_k(i)}\right)_{i\in\mathcal{I}_{|\mathcal{B}|}}$) is a bijection from $S(\mathcal{B})$ to $S(\mathcal{B})$; it follows that $\beta_k = \left[\bar{\lambda}(\mathcal{B})\right]^k = \hat{\beta}_k(\mathcal{B})$. ■

J.1 Existence

We now show that the positioning strategy $\hat{\boldsymbol{\rho}}(\mathcal{B})$ induces active consumers to inspect all slots in increasing order until finding a match, and therefore sustains the most profitable equilibrium featuring intensive search:

Lemma J.3 (existence) *There exists an equilibrium in which the firm assigns products to slots according to $\hat{\boldsymbol{\rho}}(\mathcal{B})$ and consumers with $c \leq \hat{c}(\mathcal{B})$ inspect all slots in increasing order until finding a match; the firm's profit is therefore given by $\hat{\Pi}(\mathcal{B})$.*

Proof. By construction, if active consumers inspect all slots until finding a match, the firm is indifferent among all positioning strategies, which all yield the same expected probability of a match (namely, $\mathrm{M}(\mathcal{B}) = 1 - \Lambda(\mathcal{B})$); it is thus willing to adopt the positioning strategy $\hat{\boldsymbol{\rho}}(\mathcal{B})$. Furthermore, if the active consumers are those with $c \leq \hat{c}(\mathcal{B})$, then the firm's profit is given by $\hat{\Pi}(\mathcal{B})$.

Conversely, if the firm adopts the positioning strategy $\hat{\boldsymbol{\rho}}(\mathcal{B})$ *and* active consumers inspect slots in increasing order, then from Lemma J.2 all consumers with $c \leq \hat{c}(\mathcal{B})$ are willing to participate and keep searching until finding a match. To conclude the proof, it remains to show that active consumers are indeed willing to inspect slots in increasing order.

For any $k \in I_{|\mathcal{B}|}$, and any search sequence $\mathbf{s} \in S\left(\mathcal{I}_{|\mathcal{B}|}\right)$, where s_k denotes the k^{th} inspected slot, let

$$\hat{\beta}_k\left(\mathbf{s}\right) \equiv \Pr\left[m_{s_1} = \cdots = m_{s_k} = 0\right]$$

denote the probability that the slots $(s_1, \ldots, s_k)$ yield no match. By construction, $\hat{\beta}_{|\mathcal{B}|}\left(\mathbf{s}\right) = 1 - \mathrm{M}\left(\mathcal{B}\right)$ for any $\mathbf{s} \in S\left(\mathcal{I}_{|\mathcal{B}|}\right)$, and $\hat{\beta}_k\left((1, \ldots, |\mathcal{B}|)\right) = \beta_k$ for any $k \in I_{|\mathcal{B}|-1}$. If a consumer with search cost c inspects slots according to the sequence $\mathbf{s} \in S\left(\mathcal{B}\right)$, then inspecting at most k slots yields an expected utility given by:

$$\hat{V}_k\left(c; \mathbf{s}\right) \equiv \left[1 - \hat{\beta}_k\left(\mathbf{s}\right)\right] s^m - \sum_{h=1}^{k} \hat{\beta}_{h-1}\left(\mathbf{s}\right) c.$$

By construction, $\hat{\beta}_k\left(\cdot\right)$ is independent of the order of its first k arguments, as the probability that the first k visited slots do not produce a match does not depend on the order in which these slots are visited. We now show that $\hat{\beta}_k\left(\cdot\right)$ is moreover strictly increasing in each of its first k arguments. To see this, consider two search sequences $\mathbf{s}, \tilde{\mathbf{s}} \in S\left(\mathcal{B}\right)$ that differ only in the $\bar{h}^{\text{th}}$ inspected slot among the first k ones, with $s_{\bar{h}} = \bar{s}$ being replaced by $\tilde{s}_{\bar{h}} = \bar{s} + 1$; that is:[67]

$$\begin{cases} \tilde{s}_i = s_i & \text{if} \quad i \in \mathcal{I}_k \diagdown \left\{\bar{h}\right\}, \\ s_i = \bar{s} \text{ and } \tilde{s}_i = \bar{s} + 1 & \text{if} \quad i = \bar{h}. \end{cases}$$

It follows from the definition of $\hat{\beta}_k\left(\cdot\right)$ that $\hat{\beta}_k\left(\tilde{\mathbf{s}}\right) - \hat{\beta}_k\left(\mathbf{s}\right) = N/D$, where $D \equiv \sum_{\tilde{\boldsymbol{\sigma}} \in S(\mathcal{B})} \prod_{i \in \mathcal{I}_{|\mathcal{B}|}} \lambda_{\tilde{\sigma}_i}^{\frac{i-1}{|\mathcal{B}|}} > 0$ and

$$N \equiv \sum_{\boldsymbol{\sigma} \in S(\mathcal{B})} \left\{ \prod_{i \in \mathcal{I}_{|\mathcal{B}|}} \lambda_{\sigma_i}^{\frac{i-1}{|\mathcal{B}|}} \left(\prod_{i \in \mathcal{I}_k} \lambda_{\sigma_{\tilde{s}_i}} - \prod_{i \in \mathcal{I}_k} \lambda_{\sigma_{s_i}} \right) \right\}.$$

Let g denote the permutation of $\mathcal{I}_{|\mathcal{B}|}$ swapping $\bar{s}$ and $\bar{s}+1$:

$$g\left(i\right) \equiv \begin{cases} i & \text{if} \quad i \neq \bar{I}, \\ \bar{s}+1 & \text{if} \quad i = \bar{s}, \\ \bar{s} & \text{if} \quad i = \bar{s}+1. \end{cases}$$

[67] To be feasible, this requires $\bar{s} + 1 \notin \{s_1, \ldots, s_k\}$.

As g is a bijection from $S\left(\mathcal{I}_{|\mathcal{B}|}\right)$ to $S\left(\mathcal{I}_{|\mathcal{B}|}\right)$, N can be expressed as:

$$\begin{aligned}
N &= \frac{1}{2}\left\{\begin{array}{c} \sum_{\boldsymbol{\sigma}\in S(\mathcal{B})}\left\{\left(\prod_{i\in\mathcal{I}_{|\mathcal{B}|}}\lambda_{\sigma_i}^{\frac{i-1}{|\mathcal{B}|}}\right)\left(\prod_{i\in\bar{S}}\lambda_{\sigma_i}\right)\left(\lambda_{\sigma_{\bar{s}+1}}-\lambda_{\sigma_{\bar{s}}}\right)\right\} \\ +\sum_{\boldsymbol{\sigma}\in S(\mathcal{B})}\left\{\left(\prod_{i\in\mathcal{I}_{|\mathcal{B}|}}\lambda_{\sigma_{g(i)}}^{\frac{i-1}{|\mathcal{B}|}}\right)\left(\prod_{i\in\bar{S}}\lambda_{\sigma_{g(i)}}\right)\left(\lambda_{\sigma_{g(\bar{s}+1)}}-\lambda_{\sigma_{g(\bar{s})}}\right)\right\} \end{array}\right\} \\
&= \frac{1}{2}\sum_{\boldsymbol{\sigma}\in S(\mathcal{B})}\left(\prod_{i\in\bar{I}}\lambda_{\sigma_i}^{\frac{i-1}{|\mathcal{B}|}}\right)\left(\prod_{i\in\bar{S}}\lambda_{\sigma_i}\right)\lambda_{\sigma_{\bar{s}}}^{\frac{\bar{s}-1}{|\mathcal{B}|}}\lambda_{\sigma_{\bar{s}+1}}^{\frac{\bar{s}-1}{|\mathcal{B}|}}\left(\lambda_{\sigma_{\bar{s}+1}}^{\frac{1}{|\mathcal{B}|}}-\lambda_{\sigma_{\bar{s}}}^{\frac{1}{|\mathcal{B}|}}\right)\left(\lambda_{\sigma_{\bar{s}+1}}-\lambda_{\sigma_{\bar{s}}}\right) \\
&> 0,
\end{aligned}$$

where $\bar{I}\equiv\mathcal{I}_{|\mathcal{B}|}\diagdown\{\bar{s},\bar{s}+1\}$ and $\bar{S}\equiv\{1,\ldots,k\}\diagdown\{\bar{s}\}$, the second equality stems from the definition of g, and the inequality follows from $\lambda_i>0$ for $i\in I_{|\mathcal{B}|}$ and, for any $x,y>0$:

$$\left(x^{\frac{1}{|\mathcal{B}|}}-y^{\frac{1}{|\mathcal{B}|}}\right)(x-y)=\left(x^{\frac{1}{|\mathcal{B}|}}-y^{\frac{1}{|\mathcal{B}|}}\right)^2\sum_{k=0}^{|\mathcal{B}|-1}x^{\frac{k}{|\mathcal{B}|}}y^{\frac{|\mathcal{B}|-k}{|\mathcal{B}|}}>0.$$

The monotonicity of $\hat{\beta}_k(\mathbf{s})$ in its first k arguments ensures in turn that, for any $k\in I_{|\mathcal{B}|}$, any consumer willing to inspect up to k slots strictly prefers to inspect them in increasing order. To see this, it suffices to note that: (i) $\hat{V}_k(c;\mathbf{s})$ depends on $\mathbf{s}$ only through $(\beta_h(\mathbf{s}))_{h\in\mathcal{I}_k}$ and, for every $h\in\mathcal{I}_k$, is strictly decreasing in $\beta_h(\cdot)$; and (ii) for each $h\in\mathcal{I}_k$, $\beta_h(\mathbf{s})$ is minimal when and only when $\{s_1,\ldots,s_h\}=\mathcal{I}_h$. It follows that $\hat{V}_k(c;\mathbf{s})$ is maximal when and only when $(s_1,\ldots,s_k)=(1,...,k)$. ■

Remark 9 (monotonicity of expected match probabilities) *The monotonicity of $\hat{\beta}_k(\mathbf{s})$ in its first k arguments also implies that, conditional on not having a match in the first k slots, the expected match probabilities offered by the remaining slots are decreasing in their ranks. To see this, it suffices to note that the conditional probability of a match in slot $h>k$ is given by:*

$$\begin{aligned}
\Pr\left[m_h=1\mid m_1=\cdots=m_k=0\right] &= 1-\Pr\left[m_h=0\mid m_1=\cdots=m_k=0\right] \\
&= 1-\frac{\Pr\left[m_h=0 \textit{ and } m_1=\cdots=m_k=0\right]}{\Pr\left[m_1=\cdots=m_k=0\right]} \\
& \quad 1-\frac{\beta_{k+1}(1,\ldots,k,h)}{\beta_k(1,\ldots,k)}.
\end{aligned}$$

Therefore, for any $k\in\mathcal{I}_{|\mathcal{B}|-2}$, any $h\in\mathcal{I}_{|\mathcal{B}|-1}\setminus\mathcal{I}_k$ and any $\ell\in\mathcal{I}_{|\mathcal{B}|}\setminus\mathcal{I}_h$:

$$\begin{aligned}
&\Pr\left[m_\ell=1\mid m_1=\cdots=m_k=0\right]-\Pr\left[m_h=1\mid m_1=\cdots=m_k=0\right] \\
&\quad=\frac{\beta_{k+1}(1,\ldots,k,h)-\beta_{k+1}(1,\ldots,k,\ell)}{\beta_k(1,\ldots,k)} \\
&\quad>0,
\end{aligned}$$

where the inequality follows from the monotonicity of $\beta_{k+1}(\cdot)$ and $\ell > h$.

K Proof of Proposition 5 (most profitable steering)

We know from Lemma 3 that any steering equilibrium is characterized by a block structure $\boldsymbol{\mathcal{B}} = (\mathcal{B}_1, \ldots, \mathcal{B}_{|\boldsymbol{\mathcal{B}}|})$ and $\hat{\mathbf{c}} = (\hat{c}_1, \ldots, \hat{c}_{|\boldsymbol{\mathcal{B}}|})$, where $\hat{c}_k$ denotes the search cost of the marginal consumer willing to inspect the block $\mathcal{B}_k$; the firm's equilibrium profit is thus given by:

$$\sum_{k=1}^{|\boldsymbol{\mathcal{B}}|} G(\hat{c}_k)\, \Lambda(\bar{\mathcal{B}}_{k-1}) \mathrm{M}(\mathcal{B}_k) \pi^m.$$

Given the block structure $\boldsymbol{\mathcal{B}}^S$ of the most profitable equilibrium, maximizing the profit of the firm among the steering equilibria amounts to maximizing participation. It then follows from Proposition 4 that, in the most profitable steering equilibrium, $\hat{c}_k = \hat{c}\left(\mathcal{B}_k^S\right)$ for each block $\mathcal{B}_k^S$.

L Proof of Proposition 6 (product popularities)

Using $\mathrm{M}(\mathcal{B}) = \mathrm{M}(\mathcal{B}_1) + \Lambda(\mathcal{B}_1)\mathrm{M}(\mathcal{B}_2)$, $\Delta(\mathcal{B}_1, \mathcal{B}_2)$ can be expressed as:

$$\begin{aligned}
\Delta(\mathcal{B}_1, \mathcal{B}_2) &= G(\hat{c}(\mathcal{B}))\left[\mathrm{M}(\mathcal{B}_1) + \Lambda(\mathcal{B}_1)\mathrm{M}(\mathcal{B}_2)\right] - \left[G(\hat{c}(\mathcal{B}_1))\mathrm{M}(\mathcal{B}_1) + \Lambda(\mathcal{B}_1) G(\hat{c}(\mathcal{B}_2))\mathrm{M}(\mathcal{B}_2)\right] \\
&= \Lambda(\mathcal{B}_1)\mathrm{M}(\mathcal{B}_2)\left[G(\hat{c}(\mathcal{B})) - G(\hat{c}(\mathcal{B}_2))\right] - \left[G(\hat{c}(\mathcal{B}_1)) - G(\hat{c}(\mathcal{B}))\right]\mathrm{M}(\mathcal{B}_1) \\
&= B(\mathcal{B}_1, \mathcal{B}_2) - C(\mathcal{B}_1, \mathcal{B}_2),
\end{aligned}$$

where $C(\mathcal{B}_1, \mathcal{B}_2)$ and $B(\mathcal{B}_1, \mathcal{B}_2)$ are respectively given by (7) and (8).

In what follows, we fix n_1, n_2 and restrict attention to $\mathcal{B}_1, \mathcal{B}_2 \subseteq \mathcal{I}_N$ satisfying $|\mathcal{B}_i| = n_i$ for $i = 1, 2$, $\bar{\lambda}(\mathcal{B}_1) = \lambda_1 \in (0,1)$ and $\bar{\lambda}(\mathcal{B}_2) = \lambda_2 \in (\lambda_1, 1)$. We thus have $\Lambda(\mathcal{B}_i) = \lambda_i^{n_i}$ and:

$$\begin{aligned}
\Delta(\mathcal{B}_1, \mathcal{B}_2) &= G\left(\left[1 - \bar{\lambda}(B_1 \cup \mathcal{B}_2)\right] s^m\right)\left[1 - \Lambda(B_1 \cup \mathcal{B}_2)\right]\pi^m \\
&\quad - G\left(\left[1 - \bar{\lambda}(B_1)\right] s^m\right)\left[1 - \Lambda(B_1)\right]\pi^m - \Lambda(B_1)\, G\left(\left[1 - \bar{\lambda}(B_2)\right] s^m\right)\left[1 - \Lambda(B_2)\right]\pi^m \\
&= G\left(\left(1 - \lambda_1^{\frac{n_1}{n_1+n_2}} \lambda_2^{\frac{n_2}{n_1+n_2}}\right) s^m\right)\left(1 - \lambda_1^{n_1}\lambda_2^{n_2}\right)\pi^m \\
&\quad - G\left((1 - \lambda_1) s^m\right)\left(1 - \lambda_1^{n_1}\right)\pi^m - \lambda_1^{n_1} G\left((1 - \lambda_2) s^m\right)\left(1 - \lambda_2^{n_2}\right)\pi^m \\
&\equiv \hat{\Delta}(\lambda_1, \lambda_2),
\end{aligned}$$

where $\hat{\Delta}(\cdot)$ satisfies:

$$\hat{\Delta}(\lambda, \lambda) = G\left((1 - \lambda) s^m\right)\left[\left(1 - \lambda^{n_1+n_2}\right) - \left(1 - \lambda^{n_1}\right) - \lambda^{n_1}\left(1 - \lambda^{n_2}\right)\right]\pi^m = 0. \quad (12)$$

Part (i). For $\lambda_2 = 1$, we have:

$$\hat{\Delta}(\lambda_1, 1) = \left[G\left(\left(1 - \lambda_1^{\frac{n_1}{n_1+n_2}}\right) s^m\right) - G\left((1-\lambda_1) s^m\right)\right]\left(1 - \lambda_1^{n_1}\right)\pi^m < 0,$$

where the inequality stems from $\lambda_1 \in (0,1)$ and $0 < n_1 < n_1 + n_2$. It follows that, for any λ_1, $\hat{\Delta}(\lambda_1, \lambda_2) < 0$ for λ_2 close enough to 1.

For $\lambda_1 = 0$, we have:

$$\hat{\Delta}(0, \lambda_2) = G(s^m)\pi^m - G(s^m)\pi^m - 0 = 0.$$

Furthermore, the derivative of $\hat{\Delta}$ with respect to λ_1 is given by:

$$\begin{aligned}
\frac{\partial \hat{\Delta}}{\partial \lambda_1}(\lambda_1, \lambda_2) \quad = \quad & G\left(\left(1 - \lambda_1^{\frac{n_1}{n_1+n_2}}\lambda_2^{\frac{n_2}{n_1+n_2}}\right) s^m\right)\left(-n_1\lambda_1^{n_1-1}\lambda_2^{n_2}\right)\pi^m \\
& + g\left(\left(1 - \lambda_1^{\frac{n_1}{n_1+n_2}}\lambda_2^{\frac{n_2}{n_1+n_2}}\right) s^m\right)\left(-\frac{n_1}{n_1+n_2}\lambda_1^{\frac{n_1}{n_1+n_2}-1}\lambda_2^{\frac{n_2}{n_1+n_2}}\right) s^m\left(1 - \lambda_1^{n_1}\lambda_2^{n_2}\right)\pi^m \\
& - G\left((1-\lambda_1) s^m\right)\left(-n_1\lambda_1^{n_1-1}\right)\pi^m - g\left((1-\lambda_1) s^m\right)\left(-s^m\right)\left(1 - \lambda_1^{n_1}\right)\pi^m \\
& - n_1\lambda_1^{n_1-1} G\left((1-\lambda_2) s^m\right)\left(1 - \lambda_2^{n_2}\right)\pi^m.
\end{aligned}$$

This derivative has the same sign as $\phi(\lambda_1, \lambda_2) \equiv \lambda_1 \frac{\partial \hat{\Delta}}{\partial \lambda_1}(\lambda_1, \lambda_2)$, which for $\lambda_1 = 0$, boils down to:

$$\phi(0, \lambda_2) = g(s^m) s^m \pi^m > 0.$$

Hence, for any $\lambda_2 \in (0,1)$, we have $\hat{\Delta}(0, \lambda_2) = 0$ and $\lim_{\lambda_1 \to 0} \frac{\partial \hat{\Delta}}{\partial \lambda_1}(\lambda_1, \lambda_2) = +\infty$. It follows that $\hat{\Delta}(\lambda_1, \lambda_2) < 0$ for λ_1 close enough to 0.

Part (ii). For $\lambda_1 = \lambda_2$, the derivative of $\hat{\Delta}$ with respect to λ_1 boils down to:

$$\frac{\partial \hat{\Delta}}{\partial \lambda_1}(\lambda_2, \lambda_2) = g\left((1-\lambda_2) s^m\right) s^m \pi^m \frac{1-\lambda_2}{n_1+n_2} \times A,$$

where:

$$\begin{aligned}
A \quad &= \quad (n_1+n_2)\frac{1-\lambda_2^{n_1}}{1-\lambda_2} - n_1\frac{1-\lambda_2^{n_1+n_2}}{1-\lambda_2} \\
&= \quad (n_1+n_2)\sum_{k=0}^{n_1-1}\lambda_2^k - n_1\sum_{k=0}^{n_1+n_2-1}\lambda_2^k \\
&= \quad n_2\sum_{k=0}^{n_1-1}\lambda_2^k - n_1\sum_{k=n_1}^{n_1+n_2-1}\lambda_2^k \\
&\geq \quad n_1 n_2 \lambda_2^{n_1-1}(1-\lambda_2) \\
&> \quad 0,
\end{aligned}$$

where the first inequality stems from all n_1 terms in the first sum being at least equal to $\lambda_1^{n_1-1}$ and all n_2 terms in the second sum being at most equal to $\lambda_1^{n_1}$. It follows that

$$\frac{\partial\hat{\Delta}}{\partial\lambda_1}(\lambda_2,\lambda_2) > 0.$$

Combined with (12) for $\lambda = \lambda_2$, this implies that, for any λ_2, $\hat{\Delta}(\lambda_1,\lambda_2) < 0$ for λ_1 (lower than and) close enough to λ_2.

Likewise, the derivative of $\hat{\Delta}(\cdot)$ with respect to λ_2 is given by:

$$\begin{aligned}\frac{\partial\hat{\Delta}}{\partial\lambda_2}(\lambda_1,\lambda_2) &= G\left(\left(1-\lambda_1^{\frac{n_1}{n_1+n_2}}\lambda_2^{\frac{n_2}{n_1+n_2}}\right)s^m\right)\left(-n_2\lambda_1^{n_1}\lambda_2^{n_2-1}\right)\pi^m \\ &\quad +g\left(\left(1-\lambda_1^{\frac{n_1}{n_1+n_2}}\lambda_2^{\frac{n_2}{n_1+n_2}}\right)s^m\right)\left(-\frac{n_2}{n_1+n_2}\lambda_1^{\frac{n_1}{n_1+n_2}}\lambda_2^{\frac{n_2}{n_1+n_2}-1}\right)s^m\left(1-\lambda_1^{n_1}\lambda_2^{n_2}\right)\pi^m \\ &\quad -\lambda_1^{n_1}G\left((1-\lambda_2)s^m\right)\left(-n_2\lambda_2^{n_2-1}\right)\pi^m-\lambda_1^{n_1}g\left((1-\lambda_2)s^m\right)\left(-s^m\right)\left(1-\lambda_2^{n_2}\right)\pi^m,\end{aligned}$$

which, for $\lambda_2 = \lambda_1$, boils down to:

$$\frac{\partial\hat{\Delta}}{\partial\lambda_2}(\lambda_1,\lambda_1) = -g\left((1-\lambda_1)s^m\right)s^m\pi^m\frac{1-\lambda_1}{n_1+n_2}\times B,$$

where:

$$\begin{aligned}B &= n_2\frac{1-\lambda_1^{n_1+n_2}}{1-\lambda_1}-(n_1+n_2)\lambda_1^{n_1}\frac{1-\lambda_1^{n_2}}{1-\lambda_1} \\ &= n_2\sum_{k=0}^{n_1+n_2-1}\lambda_1^k-(n_1+n_2)\lambda_1^{n_1}\sum_{k=0}^{n_2-1}\lambda_1^k \\ &= n_2\sum_{k=0}^{n_1-1}\lambda_1^k-n_1\sum_{k=n_1}^{n_1+n_2-1}\lambda_1^k \\ &\geq n_2n_1\lambda_1^{n_1-1}(1-\lambda_1) \\ &> 0,\end{aligned}$$

where the first inequality stems from all n_1 terms in the first sum being at least equal to $\lambda_1^{n_1-1}$ and all n_2 terms in the second sum being at most equal to $\lambda_1^{n_1}$. It follows that

$$\frac{\partial\hat{\Delta}}{\partial\lambda_2}(\lambda_1,\lambda_1) < 0.$$

Combined with (12) for $\lambda = \lambda_1$, this implies that, for any λ_1, $\hat{\Delta}(\lambda_1,\lambda_2) < 0$ for λ_2 (higher than and) close enough to λ_1.

M Proof of Proposition 7 (cost distribution)

Fix two successive blocks $\mathcal{B}_1$ and $\mathcal{B}_2$ and let $\lambda_i \equiv \bar{\lambda}(\mathcal{B}_i)$ and $n_i \equiv |\mathcal{B}_i|$ denote the associated no-match probabilities and sizes. By assumption, we have:

$$\lambda_2 > \lambda_1. \tag{13}$$

From the proof of Proposition 6, $\Delta_\gamma(\mathcal{B}_1, \mathcal{B}_2) \gtrless 0$ if and only if

$$\Lambda(\mathcal{B}_1)\,\mathrm{M}(\mathcal{B}_2)\left[G_\gamma(\hat{c}(\mathcal{B})) - G_\gamma(\hat{c}(\mathcal{B}_2))\right] \gtrless \left[G_\gamma(\hat{c}(\mathcal{B}_1)) - G_\gamma(\hat{c}(\mathcal{B}))\right]\mathrm{M}(\mathcal{B}_1),$$

which amounts to
$$\frac{G_\gamma(\hat{c}(\mathcal{B})) - G_\gamma(\hat{c}(\mathcal{B}_2))}{G_\gamma(\hat{c}(\mathcal{B}_1)) - G_\gamma(\hat{c}(\mathcal{B}))} \gtrless \frac{1-\lambda_1^{n_1}}{\lambda_1^{n_1}\left(1-\lambda_2^{n_2}\right)},$$
where the right-hand side is independent of γ and the left-hand side is equal to (using $\lambda = \bar{\lambda}(\mathcal{B})$, the definition of G_γ and some manipulation):

$$r(\gamma) \equiv \frac{\exp(\gamma s^m(\lambda_2 - \lambda)) - 1}{1 - \exp(-\gamma s^m(\lambda - \lambda_1))} = \frac{\exp(\gamma s^m b) - 1}{1 - \exp(-\gamma s^m a)},$$

where

$$a \equiv \lambda - \lambda_1 > 0 \text{ and } b \equiv \lambda_2 - \lambda > 0.$$

The derivative of $r(\cdot)$ is given by:

$$r'(\gamma) = \frac{\psi(\gamma)}{\left[1 - \exp(-\gamma s^m a)\right]^2} s^m,$$

where

$$\psi(\gamma) \equiv a\exp(-\gamma s^m a) + b\exp(\gamma s^m b) - (a+b)\exp(-\gamma s^m(a-b)).$$

It follows that $r(\gamma)$ is strictly increasing in γ if and only if $\psi(\gamma) > 0$. We have:

$$\begin{aligned}
\psi'(\gamma) &= s^m\left[-a^2\exp(-\gamma s^m a) + b^2\exp(\gamma s^m b) + \left(a^2 - b^2\right)\exp(-\gamma s^m(a-b))\right],\\
\psi''(\gamma) &= (s^m)^2\left[a^3\exp(-\gamma s^m a) + b^3\exp(\gamma s^m b) - \left(a^2 - b^2\right)(a-b)\exp(-\gamma s^m(a-b))\right],
\end{aligned}$$

and so:

$$\begin{aligned}
\psi(0) &= a + b - (a+b) = 0,\\
\psi'(0) &= s^m\left[-a^2 + b^2 + \left(a^2 - b^2\right)\right] = 0,\\
\psi''(0) &= (s^m)^2\left[a^3 + b^3 - \left(a^2 - b^2\right)(a-b)\right] = (s^m)^2\,ab\,(a+b) > 0.
\end{aligned}$$

Hence, $\psi(\gamma) > 0$ for γ close enough to 0. Furthermore, for any candidate value of γ for which $\psi(\gamma) = 0$, which amounts to:

$$\exp(-\gamma s^m (a-b)) = \frac{a \exp(-\gamma s^m a) + b \exp(\gamma s^m b)}{a+b},$$

we have:

$$\psi'(\gamma) = abs^m \left[\exp(\gamma s^m b) - \exp(-\gamma s^m a)\right] > 0,$$

where the inequality stems from $\gamma s^m b, \gamma s^m a > 0$. It follows that $\psi(\cdot)$ remains positive for any $\gamma > 0$. Hence, $r(\gamma)$ is strictly increasing in γ for $\gamma \geq 0$.

Furthermore, $\lim_{\gamma \longrightarrow +\infty} r(\gamma) = +\infty$; hence, there exists $\hat{\gamma}(\mathcal{B}_1, \mathcal{B}_2) < +\infty$ such that $\Delta_\gamma(\mathcal{B}_1, \mathcal{B}_2) \gtreqless 0$ if and only if $\gamma \gtreqless \hat{\gamma}(\mathcal{B}_1, \mathcal{B}_2)$.

N Proof of Corollary 4 (reduction in search costs)

Consider for example the case where $N = 2$. Fix the product characteristics λ_1 and λ_2, and suppose that the distribution of search costs has a constant hazard rate γ. Let $\bar{\lambda} \equiv \bar{\lambda}(\mathcal{I}_2)$, $\hat{c} \equiv \hat{c}(\mathcal{I}_2)$ and $\hat{\gamma} \equiv \hat{\gamma}(\{1\}, \{2\})$. An increase in γ from $\hat{\gamma} - \epsilon$ to $\hat{\gamma} + \epsilon$, for some $\epsilon > 0$, corresponds to a first-order stochastically dominant enhancement, consistent with a decrease in every consumer's search cost, and triggers a switch in the most profitable equilibrium, from pure to fully random positioning. Recall that every consumer's surplus is maximal under pure positioning; hence, for ϵ small enough, all participating consumers are worse-off.

Specifically, initially the active consumers are those with a search cost $c \leq c_1^P$. Among these consumers, those with $c < c_2^P$ face a jump in their expected number of inspections, from $1+\lambda_1$ to $1+\bar{\lambda}$,[68] those with $c_2^P < c < \hat{c}$ face a discrete drop in surplus, from $(1-\lambda_1)s^m - c$ to $(1-\bar{\lambda}^2)s^m - (1+\bar{\lambda})c$, at least equal to $\Delta \equiv (\bar{\lambda}-\lambda_1)(1-\lambda_2)s^m > 0$,[69] and those with $\hat{c} < c < c_2^P$ drop out and thus lose their initial surplus, equal to $\mu_1 s^m - c = c_1^P - c$. It follows that, for any $\underline{c}$ and $\bar{c}$ such that $0 < \underline{c} < \bar{c} < c_1^P$, there exists ϵ small enough that all consumers with a search cost $c \in (\underline{c}, \bar{c})$ are strictly worse-off.[70]

[68]See Appendix I.

[69]The drop in surplus is equal to $\bar{\lambda}c - \lambda_1(1-\lambda_2)s^m \geq (\bar{\lambda} - \lambda_1)(1-\lambda_2)s^m$, where the inequality stems from $c \geq c_2^P = (1-\lambda_2)s^m$.

[70]It suffices to take ε small enough to ensure that any individual consumer' search cost decreases by less than $\min\{(\bar{\lambda} - \lambda_1)\underline{c}, \Delta, c_1^P - \bar{c}\}$.

O Welfare analysis

Suppose that the firm chose to offer n products and consider the best equilibrium featuring intensive search. A consumer with search cost $c \leq \hat{c}(\mathcal{I}_n) = \left(1 - \bar{\lambda}_n\right) s^m$ obtains a surplus equal to

$$\hat{s}_n(c) \equiv (1 - \Lambda_n) s^m - \frac{1 - \Lambda_n}{1 - \bar{\lambda}_n} c = \frac{1 - \Lambda_n}{1 - \bar{\lambda}_n} \left[\hat{c}(\mathcal{I}_n) - c\right].$$

Therefore, total consumer surplus is given by:

$$\hat{S}_n \equiv \int_0^{\hat{c}_n} \hat{s}_n(c)\, dG(c) = \frac{1 - \Lambda_n}{1 - \bar{\lambda}_n} \int_0^{\hat{c}_n} (\hat{c}_n - c)\, dG(c) = \frac{1 - \Lambda_n}{1 - \bar{\lambda}_n} \bar{G}(\hat{c}_n),$$

where

$$\bar{G}(c) \equiv \int_0^c G(\tilde{c})\, d\tilde{c}$$

denotes the primitive of the distribution $G(c)$.

From now on, we focus on the case $N = 2$. Consumers prefer a limited selection if and only if $\hat{s}_1(c) > \hat{s}_2(c)$, that is:

$$(1 - \lambda_1) s^m - c > \frac{1 - \lambda_1 \lambda_2}{1 - \sqrt{\lambda_1 \lambda_2}} \left[\left(1 - \sqrt{\lambda_1 \lambda_2}\right) s^m - c\right],$$

which amounts to

$$c > \tilde{c} \equiv (1 - \lambda_2) \sqrt{\frac{\lambda_1}{\lambda_2}} s^m \left(< c_2^P < s^m\right).$$

It follows that if all consumers have a cost $c \in \left(c_2^P, \hat{c}_2\right)$: (i) the firm finds it profitable to (offer all products and) introduce noise in its positioning, as consumers would never inspect product 2 under pure positioning, and yet they do so under noisy positioning; and (ii) *all* consumers would favor a limited selection.

P Aggregate uncertainty

We consider here an extended version of our setting, in which consumers are uncertain about the overall quality of the available products. Specifically, there are two states of the world, a *good* one ($\omega = G$) that occurs with probability p^G, and a *bad* one ($\omega = B$) in which all products are less popular, occurring with probability $p^B = 1 - p^G$. To simplify the exposition, we focus on the case where there are two available products, both offered by the firm: $n = N = 2$.

Let μ_i^ω denote product i's popularity in state $\omega \in \{G, B\}$, with $\mu_i^B < \mu_i^G$, $\mathrm{E}[\mu_i] = p^G \mu_i^G + \left(1 - p^G\right) \mu_i^B$ denote its expected value, and $\mathrm{Cov}(\mu_1, \mu_2) = \mathrm{E}\left[(\mu_1 - \mathrm{E}[\mu_1])(\mu_2 - \mathrm{E}[\mu_2])\right]$

measure the amount of uncertainty about the state of world. As $\mu_i^B \leq \mu_i^G$ for $i = 1, 2$, $\mathrm{Cov}(\mu_1, \mu_2) \geq 0$; furthermore, using

$$\mathrm{Cov}(\mu_1, \mu_2) = \mathrm{E}[\mu_1 \mu_2] - \mathrm{E}[\mu_1] \mathrm{E}[\mu_2] \tag{14}$$

and $\mu_1^\omega \leq 1$, we have:

$$\mathrm{Cov}(\mu_1, \mu_2) \leq (1 - \mathrm{E}[\mu_1]) \mathrm{E}[\mu_2]. \tag{15}$$

P.1 Search Addiction

Under random positioning, the probability of a match on a first inspection is:

$$\Pr[m_1 = 1] = \frac{\mathrm{E}[\mu_1] + \mathrm{E}[\mu_2]}{2}.$$

In the absence of a match on the first inspection, the probability of a match on a second one is given by:

$$\begin{aligned}
\Pr[m_2 = 1 \mid m_1 = 0] &= \frac{\Pr[m_2 = 1 \text{ and } m_1 = 0]}{\Pr[m_1 = 0]} \\
&= \frac{\begin{array}{c} \frac{p^G}{2}\left(1 - \mu_1^G\right)\mu_2^G + \frac{p^G}{2}\left(1 - \mu_2^G\right)\mu_1^G \\ + \frac{p^B}{2}\left(1 - \mu_1^B\right)\mu_2^B + \frac{p^B}{2}\left(1 - \mu_2^B\right)\mu_1^B \end{array}}{1 - \frac{\mathrm{E}[\mu_1] + \mathrm{E}[\mu_2]}{2}} \\
&= \frac{\frac{\mathrm{E}[\mu_1] + \mathrm{E}[\mu_2]}{2} - \mathrm{E}[\mu_1 \mu_2]}{1 - \frac{\mathrm{E}[\mu_1] + \mathrm{E}[\mu_2]}{2}}
\end{aligned}$$

There is search addiction if and only if $\Pr[m_2 = 1 \mid m_1 = 0] > \Pr[m_1 = 1]$. Using (14), this amounts to:

$$\frac{\frac{\mathrm{E}[\mu_1] + \mathrm{E}[\mu_2]}{2} - \mathrm{Cov}(\mu_1, \mu_2) - \mathrm{E}[\mu_1] \mathrm{E}[\mu_2]}{1 - \frac{\mathrm{E}[\mu_1] + \mathrm{E}[\mu_2]}{2}} > \frac{\mathrm{E}[\mu_1] + \mathrm{E}[\mu_2]}{2},$$

or, multiplying by $1 - (\mathrm{E}[\mu_1] + \mathrm{E}[\mu_2])/2$ and simplifying:

$$\mathrm{E}[\mu_1] - \mathrm{E}[\mu_2] > 2\sqrt{\mathrm{Cov}(\mu_1, \mu_2)}, \tag{16}$$

In particular, search addition arises whenever $\mathrm{E}[\mu_2]$ is not too close to $\mathrm{E}[\mu_1]$: using the upper bound given by (15), condition (16) holds whenever $\mathrm{E}[\mu_2] < \phi(\mathrm{E}[\mu_1])$, where

$$\phi(\mu) \equiv 2 - \mu - 2\sqrt{1 - \mu}$$

lies below μ and increases with μ in the relevant range $\mu \in (0, 1)$.

P.2 Pure Positioning

Consider now a pure positioning equilibrium in which product i is assigned to slot i, and consumers inspect slot 1 first. The probability of state ω when there is no match in the first slot is given by

$$\Pr(\text{state is } \omega \mid m_1 = 0) = \frac{\Pr[\text{state is } \omega \text{ and } m_1 = 0]}{\Pr[m_1 = 0]} = \frac{p^\omega(1-\mu_1^\omega)}{1-\mathrm{E}[\mu_1]}.$$

It follows that the probability of having a match at the second inspection, conditional on having no match at the first is

$$\begin{aligned}\Pr(m_2 = 1 | m_1 = 0) &= \frac{p^G(1-\mu_1^G)\mu_2^G + p^B(1-\mu_1^B)\mu_2^B}{1-\mathrm{E}[\mu_1]} \\ &= \frac{\mathrm{E}[\mu_2] - \mathrm{E}[\mu_1\mu_2]}{1-\mathrm{E}[\mu_1]} \\ &= \frac{(1-\mathrm{E}[\mu_1])\,\mathrm{E}[\mu_2] - \mathrm{Cov}(\mu_1,\mu_2)}{1-\mathrm{E}[\mu_1]},\end{aligned}$$

where the last equality follows from (14). The marginal searchers are thus given by $\hat{c}_1^P = \mathrm{E}[\mu_1]\, s^m$ and

$$\hat{c}_2^P = \left[\mathrm{E}[\mu_2] - \frac{\mathrm{Cov}(\mu_1,\mu_2)}{1-\mathrm{E}[\mu_1]}\right] s^m$$

The firm's expected profit is

$$\begin{aligned}\pi^P &= G(\mathrm{E}[\mu_1]\, s^m)\mu_1\pi^m \\ &\quad + (1-\mathrm{E}[\mu_1])G\left(\left[\mathrm{E}[\mu_2] - \frac{\mathrm{Cov}(\mu_1,\mu_2)}{1-\mathrm{E}[\mu_1]}\right] s^m\right)\left[\mathrm{E}[\mu_2] - \frac{\mathrm{Cov}(\mu_1,\mu_2)}{1-\mathrm{E}[\mu_1]}\right]\pi^m.\end{aligned}$$

Holding $\mathrm{E}[\mu_1]$ and $\mathrm{E}[\mu_2]$ fixed, π^P is strictly decreasing in $\mathrm{Cov}(\mu_1,\mu_2)$: as uncertainty about the state of the world increases, not having a match on the first inspection brings worse news, which reduces the expected probability of a match with the second product and discourages consumers from inspecting it.

P.3 Best Intensive Search

If product 1 is assigned to slot 1 with probability r, the first slot produces a match with probability:

$$\Pr(m_1 = 1) = r\,\mathrm{E}[\mu_1] + (1-r)\,\mathrm{E}[\mu_2].$$

Hence, the probability of a match in slot 2, conditional on no match in slot 1, can be written as

$$\begin{aligned}
\Pr(m_2 = 1|m_1 = 0) &= \frac{\sum_{\omega\in\{G,B\}} \Pr\left[\text{state is } \omega, m_2 = 1 \text{ and } m_1 = 0\right]}{\Pr\left[m_1 = 0\right]} \\
&= \frac{\begin{array}{c} p^G\left[r(1-\mu_1^G)\mu_2^G + (1-r)(1-\mu_2^G)\mu_1^G\right] \\ +p^B\left[r(1-\mu_1^B)\mu_2^B + (1-r)(1-\mu_2^B)\mu_1^B\right]\end{array}}{1 - r\,\mathrm{E}\left[\mu_1\right] - (1-r)\,\mathrm{E}\left[\mu_2\right]} \\
&= \frac{r\,\mathrm{E}\left[\mu_2\right] + (1-r)\,\mathrm{E}\left[\mu_1\right] - \mathrm{E}\left[\mu_1\mu_2\right]}{1 - r\,\mathrm{E}\left[\mu_1\right] - (1-r)\,\mathrm{E}\left[\mu_2\right]} \\
&= \frac{r\,\mathrm{E}\left[\mu_2\right] + (1-r)\,\mathrm{E}\left[\mu_1\right] - \mathrm{E}\left[\mu_1\right]\mathrm{E}\left[\mu_2\right] - \mathrm{Cov}\left(\mu_1,\mu_2\right)}{1 - r\,\mathrm{E}\left[\mu_1\right] - (1-r)\,\mathrm{E}\left[\mu_2\right]}.
\end{aligned}$$

The best equilibrium featuring intensive search is such that $\Pr(m_1 = 1) = \Pr(m_2 = 1|m_1 = 0)$, which amounts to (using (14)):

$$\begin{aligned}
&r\,\mathrm{E}\left[\mu_2\right] + (1-r)\,\mathrm{E}\left[\mu_1\right] - \mathrm{E}\left[\mu_1\right]\mathrm{E}\left[\mu_2\right] - \mathrm{Cov}\left(\mu_1,\mu_2\right) \\
&= \left\{1 - r\,\mathrm{E}\left[\mu_1\right] - (1-r)\,\mathrm{E}\left[\mu_2\right]\right\}\left\{r\,\mathrm{E}\left[\mu_1\right] + (1-r)\,\mathrm{E}\left[\mu_2\right]\right\}.
\end{aligned}$$

Solving for r yields:

$$\hat{r} \equiv \frac{1 - \mathrm{E}\left[\mu_2\right] - \sqrt{\left(1 - \mathrm{E}\left[\mu_1\right]\right)\left(1 - \mathrm{E}\left[\mu_2\right]\right) + \mathrm{Cov}\left(\mu_1,\mu_2\right)}}{\mathrm{E}\left[\mu_1\right] - \mathrm{E}\left[\mu_2\right]},$$

which exceeds $1/2$ (implying that consumers are indeed willing to inspect slot 1 before slot 2) whenever (16) holds.[71] For $r = \hat{r}$, we have $\Pr(m_1 = 1) = \Pr(m_2 = 1|m_1 = 0) = \hat{\mu}$, where:

$$\hat{\mu} \equiv 1 - \sqrt{\left(1 - E[\mu_1]\right)\left(1 - E[\mu_2]\right) + \mathrm{Cov}\left(\mu_1,\mu_2\right)}. \tag{17}$$

The marginal searcher is given by $\hat{c}^I = \hat{\mu}s^m$ and the probability of an active consumer having at least one match is

$$p^G\left[1 - (1-\mu_1^G)(1-\mu_2^G)\right] + p^B\left[1 - (1-\mu_1^B)(1-\mu_2^B)\right] = \mathrm{E}\left[\mu_1\right] + \mathrm{E}\left[\mu_2\right] - \mathrm{E}\left[\mu_1\mu_2\right],$$

which, using (14) and (17), can be expressed as $1 - (1-\hat{\mu})^2$. The firm's expected profit in this equilibrium is thus equal to

$$\begin{aligned}
\pi^I = G\left(\left[1 - \sqrt{\left(1 - E[\mu_1]\right)\left(1 - E[\mu_2]\right) + \mathrm{Cov}\left(\mu_1,\mu_2\right)}\right]s^m\right) \\
\times\left[1 - \left(1 - E[\mu_1]\right)\left(1 - E[\mu_2]\right) - \mathrm{Cov}\left(\mu_1,\mu_2\right)\right]\pi^m.
\end{aligned}$$

[71] When $\mathrm{Cov}\left(\mu_1,\mu_2\right) = 0$, the expression for $\hat{r}$ boils down to $\hat{\rho}_{(1,2)}\left(\{1,2\}\right) = \sqrt{\lambda_2}/\left(\sqrt{\lambda_1} + \sqrt{\lambda_2}\right)$.

Holding $\mathrm{E}[\mu_1]$ and $\mathrm{E}[\mu_2]$ fixed, π^I is strictly decreasing in the amount of uncertainty about the state of the world, as measured by $\mathrm{Cov}(\mu_1, \mu_2)$. This is because, as already noted, not having a match on the first inspection tends to bring worse news when the state of the world becomes more uncertain; it follows that, to induce consumers to keep searching, the firm must assign the more popular product to the second slot with greater probability. This, in turn, increases the expected number of inspections needed to obtain a match, which discourages consumers from participating.

P.4 Equilibrium comparison

As π^I and π^P are both decreasing in $\mathrm{Cov}(\mu_1, \mu_2)$, the impact of aggregate uncertainty is a priori ambiguous. To show that it can indeed tilt the balance in either direction, we now provide two contrasting examples. Both examples have in common that:

- expected consumer surplus and profit per match is normalized to unity: $\pi^m = s^m = 1$;
- both states are equally likely: $p^G = p^B = 1/2$.

We also fix the expected values of the product popularities to $\mathrm{E}[\mu_1] = 0.8$ and $\mathrm{E}[\mu_2] = 0.3$. The corresponding range for the covariance is:

$$0 \leq \mathrm{Cov}(\mu_1, \mu_2) \leq (1 - \mathrm{E}[\mu_1])\,\mathrm{E}[\mu_2] = 0.06,$$

where the upper bound lies below $(\mathrm{E}[\mu_1] - \mathrm{E}[\mu_2])^2/4 = 0.0625$; it follows that (16) holds.

In both examples, we contrast the benchmark case of no aggregate uncertainty (i.e., $\mathrm{Cov}(\mu_1, \mu_2) = 0$) with the case where $\mathrm{Cov}(\mu_1, \mu_2) = 0.04$. We have:

$$\begin{aligned}&\hat{c}_2^P\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0.04} = 0.1 < \hat{c}_2^P\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0} = 0.3 \\ &< \hat{c}^I\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0.04} \simeq 0.58 < \hat{c}^I\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0} \simeq 0.63 < \hat{c}_1^P = 0.8.\end{aligned}$$

The two examples differ in terms of search cost distributions.

Example 1: aggregate uncertainty tilts the balance against noisy positioning.

Suppose first that consumers' search costs are all equal to $c = 0.6$. For the best equilibria featuring intensive search, we have:

$$\hat{c}^I\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0.04} < c < \hat{c}^I\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0}.$$

Hence, all consumers inspect both slots in the benchmark case, and none of them participates in case of aggregate uncertainty. For pure positioning, we have instead:

$$\hat{c}_2^P\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0.04}, \hat{c}_2^P\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0} < c < \hat{c}_1^P.$$

Hence, in both cases, all consumers inspect slot 1, and only that one.

It follows that noisy positioning is more profitable in the benchmark case (as it then maximizes both participation and search intensity), whereas it is unprofitable in case of aggregate uncertainty. Introducing aggregate uncertainty thus tilts the balance against noisy positioning.

Example 2: aggregate uncertainty tilts the balance in favor of noisy positioning.

Suppose now that the search cost takes two values: it is equal to $\underline{c} = 0.2$ for a fraction $\eta \in (0,1)$ of consumers have, and to $\bar{c} = 0.7$ otherwise. For pure positioning, we have $\bar{c} < \hat{c}_1^P$ and:

$$\hat{c}_2^P\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0.04} < \underline{c} < \hat{c}_2^P\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0}.$$

Hence, all consumers inspect the first slot, and high-cost consumers never inspect the second slot; low-cost consumers are however willing to keep searching in the absence of aggregate uncertainty (and only in that case).

For the best equilibrium featuring intensive search, we have instead:

$$\underline{c} < \hat{c}^I\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0.04}, \hat{c}^I\big|_{\mathrm{Cov}(\mu_1,\mu_2)=0} < \bar{c}.$$

Hence, in both cases only the low-cost consumers participate.

It follows that, in the benchmark case, pure positioning is more profitable: low-cost consumers are willing to inspect both slots in both types of equilibria, but high-cost consumers participate (and inspect slot 1) only under pure positioning. In case of aggregate uncertainty, however, a trade-off arises: high-cost consumers participate only under pure positioning (in which case they only inspect the first slot), but low-cost consumers are willing to inspect both slots only under noisy positioning. If the fraction η of low-cost consumers is large enough, the latter effect dominates,[72] in which case introducing aggregate uncertainty tilts the balance in favor of noisy positioning.

[72] For $\mathrm{Cov}(\mu_1,\mu_2) = 0.04$, $\pi^I - \pi^P = 0.82\eta - 0.8$, which is positive for $\eta > 0.98$.

Q Increasing or decreasing search costs

We consider here a variant of our setting, in which search costs vary with the number of inspections. As in Appendix P, to simplify the exposition we assume that there are two available products, both offered by the firm: $n = N = 2$.

Q.1 Decreasing costs

We start with the case where search costs are decreasing in the number of inspections (e.g., due to learning benefits). Specifically, we assume that the cost of a second inspection is a fraction $1 - \eta$ of the cost of the first inspection, for some $\eta \in (0, 1)$.

Q.1.1 Search Addiction

Under random positioning, a first inspection produces a match with probability:

$$\Pr[m_1 = 1] = \mu^e \equiv \frac{\mu_1 + \mu_2}{2},$$

and, in the absence of a match, a second inspection produces one with probability:

$$\Pr[m_2 = 1 \mid m_1 = 0] = \frac{\Pr[m_2 = 1 \text{ and } m_1 = 0]}{\Pr[m_1 = 0]} = \frac{\frac{(1-\mu_1)\mu_2}{2} + \frac{(1-\mu_2)\mu_1}{2}}{1 - \frac{\mu_1+\mu_2}{2}} = \mu^e + \Delta,$$

where

$$\Delta \equiv \frac{(\mu_1 - \mu_2)^2}{2(2 - \mu_2 - \mu_1)} > 0. \tag{18}$$

There is therefore search addiction, as a second inspection is both more promising and less costly.

Q.1.2 Pure Positioning

Consider now a pure positioning equilibrium in which product i is assigned to slot i, and consumers thus inspect slot 1 first. The marginal searchers are thus characterized by $c_1^P = \mu_1 s^m$ and

$$c_2^P(\eta) = \frac{\mu_2 s^m}{1 - \eta},$$

which increases with η and coincides with c_1^P for $\eta = \bar{\eta}$, where

$$\bar{\eta} \equiv \frac{\mu_1 - \mu_2}{\mu_1} (\in (0, 1)).$$

It follows that, if $\eta \geq \bar{\eta}$, then pure positioning induces intensive search, and therefore sustains the best possible equilibrium—in particular, this equilibrium generates the largest profit, as the intensive and extensive search margins are both maximal.

From now on, we focus on the case where $\eta < \bar{\eta}$, so that pure positioning does not induce intensive search. The firm's expected profit is then given by:

$$\Pi^P(\mathcal{I}_2;\eta) \equiv G(\mu_1 s^m)\mu_1\pi^m + (1-\mu_1)G\left(\frac{\mu_2 s^m}{1-\eta}\right)\mu_2\pi^m.$$

Q.1.3 Best Intensive Search

If product 1 is assigned to slot 1 with probability r, inspecting that slot produces a match with probability:

$$\Pr(m_1 = 1; r) = r\mu_1 + (1-r)\mu_2,$$

which increases with r. In the absence of a match, inspecting the second slot produces a match with probability:

$$\Pr(m_2 = 1|m_1 = 0; r) = \frac{\Pr[m_2 = 1 \text{ and } m_1 = 0]}{\Pr[m_1 = 0]} = \frac{r(1-\mu_1)\mu_2 + (1-r)(1-\mu_2)\mu_1}{1 - r\mu_1 - (1-r)\mu_2},$$

with is decreasing in r:

$$\frac{\partial \Pr(m_2 = 1|m_1 = 0; r)}{\partial r} = -\frac{(1-\mu_1)(1-\mu_2)(\mu_1-\mu_2)}{[1 - r\mu_1 - (1-r)\mu_2]^2} < 0.$$

Hence, as r increases, the incentive to participate (i.e., inspect slot 1) increases, but the incentive to keep searching (i.e., inspect slot 2) decreases. It follows that the best equilibrium featuring intensive search is such that the marginal consumer is willing to keep searching, that is:

$$\Pr(m_1 = 1; r) = \frac{\Pr(m_2 = 1|m_1 = 0; r)}{1-\eta}, \tag{19}$$

which amounts to $r = \hat{r}(\eta)$, where

$$\hat{r}(\eta) \equiv \frac{2 - \eta - 2\mu_2 + 2\eta\mu_2 - \sqrt{4(1-\eta)(1-\mu_1)(1-\mu_2) + \eta^2}}{2(1-\eta)(\mu_1-\mu_2)}$$

increases from $\hat{\rho}(\mathcal{I}_2) > 1/2$ (implying that consumers are indeed willing to inspect slot 1 first) to 1 (pure positioning) as η increases from 0 to $\bar{\eta}$.[73] It follows that, as η increases, the best equilibrium featuring intensive search involves less and less noise

[73] To see that $\hat{r}(\eta)$ increases with η, it suffice to note that (19) can be expressed as

$$\frac{\Pr(m_2 = 1|m_1 = 0; r)}{\Pr(m_1 = 1; r)} = 1 - \eta,$$

in positioning, which fosters participation, thereby increasing consumer surplus and profit. The participation threshold becomes:[74]

$$\hat{c}\left(\mathcal{I}_2;\eta\right)\equiv\left\{\hat{r}\left(\eta\right)\mu_1+\left[1-\hat{r}\left(\eta\right)\right]\mu_2\right\}s^m=\frac{2-\eta-\sqrt{4\left(1-\eta\right)\lambda_1\lambda_2+\eta^2}}{2\left(1-\eta\right)}s^m,$$

which increases[75] from $\hat{c}\left(\mathcal{I}_2\right)=\left(1-\sqrt{\lambda_1\lambda_2}\right)s^m$ to $c_1^P=\left(1-\lambda_1\right)s^m$ as η increases from 0 to $\bar{\eta}$. It follows by revealed preferences that consumer surplus also increases with η—not only because second inspections become less costly, but also because product positioning provides more accurate information. The firm's profit is given by:

$$\hat{\Pi}\left(\mathcal{I}_2;\eta\right)\equiv G(\hat{c}\left(\mathcal{I}_2;\eta\right))\left(1-\lambda_1\lambda_2\right)\pi^m,$$

which also increases with η.

Q.1.4 Comparison

An increase in η enhances the firm's profit under both pure and noisy positioning: in the former case, it fosters the intensive margin, by encouraging consumers to keep searching in the absence of a match with the first product (i.e., $c_2^P\left(\eta\right)$ increases with η); in the latter case, it fosters instead the extensive margin, by encouraging consumers to start searching (i.e., $\hat{c}\left(\mathcal{I}_2;\eta\right)$ increases with η). We now provide two examples showing that, as a result, an increase in η can tilt the balance in favor of either noisy or pure positioning.

Example 1: decreasing costs tilt the balance in favor of noisy positioning. Suppose first that all consumers have a search cost $\hat{c}\in\left(\hat{c}\left(\mathcal{I}_2;0\right),c_1^P\right)$, and let $\hat{\eta}\in\left(0,\bar{\eta}\right)$ and $\tilde{\eta}\in\left(\hat{\eta},\bar{\eta}\right)$ denote the values of η for which $\hat{c}\left(\mathcal{I}_2;\hat{\eta}\right)=\hat{c}$ and $c_2^P\left(\tilde{\eta}\right)=\hat{c}$. It is straightforward to check that pure positioning then dominates noisy positioning for $\eta\in\left[0,\hat{\eta}\right)$, whereas the opposite holds for $\eta\in\left(\hat{\eta},\tilde{\eta}\right)$:

- Under pure positioning, for $\eta<\tilde{\eta}$ consumers inspect product 1, and only that one; the profit is thus equal to $\mu_1\pi^m$.
- Under best noisy positioning:
  - for $\eta<\hat{\eta}$, consumers do not participate, and the firm thus obtains zero profit;

where the left-hand side is decreasing in r, whereas the right-hand side is decreasing in η.

[74] As the marginal consumer is indifferent between inspecting or not the second slot in the absence of a match at the first inspection, the expected value of starting a search is $\left\{\hat{r}\left(\eta\right)\mu_1+\left[1-\hat{r}\left(\eta\right)\right]\mu_2\right\}s^m-c$.

[75] Recall that $\hat{r}\left(\cdot\right)$ is increasing and $\mu_1>\mu_2$.

- for $\eta > \hat{\eta}$, consumers participate, and the firm obtains the maximal profit, $[\mu_1 + (1 - \mu_1)\mu_2]\pi^m > \mu_1\pi^m$.

Example 2: decreasing costs tilt the balance against noisy positioning. Suppose now that a fraction φ of consumers have a search cost $\tilde{c} \in \left(c_2^P(0), \hat{c}(\mathcal{I}_2; 0)\right)$, whereas the remaining consumers have a higher search cost $\hat{c} \in \left(\hat{c}(\mathcal{I}_2; 0), c_1^P\right)$. Let $\tilde{\eta}$ and $\hat{\eta}$ denote the values of η for which $c_2^P(\tilde{\eta}) = \tilde{c}$ and $\hat{c}(\mathcal{I}_2; \hat{\eta}) = \hat{c}$. By construction, $\tilde{\eta}$ and $\hat{\eta}$ both lie between 0 and $\bar{\eta}$. Suppose further that[76]

$$\tilde{\eta} < \hat{\eta},$$

and

$$\varphi > \frac{\mu_1}{\mu_1 + (1 - \mu_1)\mu_2}. \tag{20}$$

It is straightforward to check that noisy positioning then dominates pure positioning for $\eta \in [0, \tilde{\eta})$, and that the opposite holds for $\eta \in (\tilde{\eta}, \hat{\eta})$:

- Under best noisy positioning, for $\eta < \hat{\eta}$ only the fraction φ of low-cost consumers participate; the firm thus obtains a profit

$$\Pi^G \equiv \varphi[\mu_1 + (1 - \mu_1)\mu_2]\pi^m.$$

- Under pure positioning:
  - for $\eta < \tilde{\eta}$, all consumers inspect product 1, and only that one; hence, the profit is equal to $\mu_1\pi^m$, which, under (20), is strictly lower than Π^G;
  - by contrast, for $\eta > \tilde{\eta}$, the fraction φ of low-cost consumers inspect product 2 in the absence of a match with product 1, and the profit thus becomes (at least)[77] $\mu_1\pi^m + \varphi(1 - \mu_1)\mu_2\pi^m = \Pi^G + (1 - \varphi)\mu_1\pi^m$, and thus strictly exceeds Π^G.[78]

Q.2 Increasing costs

We now consider the case where search costs are increasing in the number of inspections (e.g., due to decreasing returns to scale); specifically, we assume a second inspection

[76] The first assumption holds for $\tilde{c}$ sufficiently close to $c_2^P(0)$ and/or $\hat{c}$ sufficiently close to c_1^P. The second assumption holds when the proportion of low-cost consumers is sufficiently large.

[77] For η large enough, all consumers are willing to inspect product 2.

[78] In this example, a further increase in η (namely, for $\eta \in (\hat{\eta}, \check{\eta})$, where $\check{\eta}$ is such that $c_2^P(\check{\eta}) = \hat{c}$) reverses again the balance in favor of noisy positioning, as it then induces all consumers to inspect both products until finding a match, whereas the fraction $1 - \varphi$ of high-cost consumers never inspect product 2—for $\eta \in (\check{\eta}, \bar{\eta})$, pure and noisy positioning both enable the firm to achieve the maximal profit.

costs $c + \delta s^m$, for some $\delta > 0$. As the analysis builds on the previous one, we only sketch the main steps here.

Q.2.1 Search Addiction

Under random positioning, a consumer with search cost c starts searching if

$$c \leq \hat{c}_1^R \equiv \mu^e s^m,$$

and keeps searching in the absence of a match if:

$$c \leq \hat{c}_2^R(\delta) \equiv (\mu^e + \Delta) s^m - \delta s^m.$$

There is search addition as long as $\hat{c}_2^R(\delta) \leq \hat{c}_1^R$, which amounts to:

$$\delta \leq \Delta,$$

where Δ is given by (18). That is, search addiction arises as long as search costs do not increase too much with the number of inspections.

Q.2.2 Pure Positioning

Under pure positioning, the costs of the marginal searchers are given by $\hat{c}_1^P = \mu_1 s^m$ and

$$\hat{c}_2^P(\delta) = \mu_2 s^m - \delta s^m.$$

The firm's expected profit is therefore

$$\pi^P = G(\mu_1 s^m)\mu_1 \pi^m + (1 - \mu_1) G((\mu_2 - \delta) s^m) \mu_2 \pi^m.$$

Q.2.3 Best Intensive Search

The best equilibrium featuring intensive search is again such that the marginal consumer is willing to keep searching, which now amounts to $\Pr(m_1 = 1; r) = \Pr(m_2 = 1 | m_1 = 0; r) - \delta$, or:

$$r = \hat{r}(\delta) \equiv \frac{2 - 2\mu_2 - \delta - \sqrt{4(1 - \mu_1)(1 - \mu_2) + \delta^2}}{2(\mu_1 - \mu_2)},$$

which decreases with δ and exceeds $1/2$ for $\delta < \Delta$. It follows that, as long as there is search addiction, introducing noise in its positioning enables the firm to induce intensive search. However, as δ increases, more noise is required, which discourages participation and decreases profit as well as consumer surplus. The participation threshold becomes:

$$\begin{aligned} \hat{c}\left(\mathcal{I}_2;\delta\right) &\equiv \left\{\hat{r}\left(\delta\right)\mu_1+\left[1-\hat{r}\left(\delta\right)\right]\mu_2\right\}s^m \\ &= \frac{2-\delta-\sqrt{4\left(1-\mu_1\right)\left(1-\mu_2\right)+\delta^2}}{2}s^m, \end{aligned}$$

which decreases[79] from $\hat{c}\left(\mathcal{I}_2\right)=\left(1-\sqrt{\lambda_1\lambda_2}\right)s^m$ to $\hat{c}_2^R\left(\Delta\right)=\mu^e s^m$ as δ increases from 0 to Δ. It follows by revealed preferences that consumer surplus is also decreasing in δ—not only because second inspections become more costly, but also because consumers have less accurate information about which product is in which slot. The firm's profit is given by:

$$\hat{\Pi}\left(\mathcal{I}_2;\delta\right)\equiv G(\hat{c}\left(\mathcal{I}_2;\delta\right))\left(1-\lambda_1\lambda_2\right)\pi^m,$$

and is also decreasing in δ.

Q.2.4 Comparison

An increase in δ reduces the firm's profit under both pure and noisy positioning: it reduces the intensive margin in the former case, and the extensive margin in the latter case. As a result, it can again tilt the balance either way.

Example 1: increasing costs tilt the balance against noisy positioning. Suppose first that all consumers have a search cost $\hat{c}\in\left(c_2^P\left(0\right),\hat{c}\left(\mathcal{I}_2;0\right)\right)$, and let $\hat{\delta}\in(0,\Delta)$ denote the value of δ for which $\hat{c}(\mathcal{I}_2;\hat{\delta})=\hat{c}$. It is straightforward to check that noisy positioning then dominates pure positioning for $\delta<\hat{\delta}$, whereas the opposite holds for $\delta>\hat{\delta}$:

- Under pure positioning, consumers inspect product 1, and only that one; the profit is thus equal to $\mu_1\pi^m$.
- Under best positioning: for $\delta<\hat{\delta}$, consumers participate, and the firm thus obtains the maximal profit, $\left[\mu_1+\left(1-\mu_1\right)\mu_2\right]\pi^m>\mu_1\pi^m$; for $\delta>\hat{\delta}$, consumers do not participate, and the firm obtains zero profit.

Example 2: increasing costs tilt the balance in favor of noisy positioning. Suppose now that a fraction φ of consumers have a search cost $\tilde{c}\in\left(0,c_2^P\left(0\right)\right)$, whereas the remaining consumers have a higher search cost $\hat{c}\in\left(\hat{c}\left(\mathcal{I}_2;0\right),c_1^P\right)$. Let $\tilde{\delta}$ and $\hat{\delta}$ denote the values of δ for which $c_2^P(\tilde{\delta})=\tilde{c}$ and $\hat{c}(\mathcal{I}_2;\hat{\delta})=\tilde{c}$. By construction, $0<\tilde{\delta}<\hat{\delta}<\Delta$. Suppose further that (20) holds. It is straightforward to check that pure

[79] Recall that $\hat{r}\left(\cdot\right)$ is increasing and $\mu_1>\mu_2$.

positioning then dominates noisy positioning for $\delta \in [0, \tilde{\delta})$, and that the opposite holds for $\delta \in (\tilde{\delta}, \hat{\delta})$:

- Under best noisy positioning, for $\delta < \hat{\delta}$ only the fraction φ of low-cost consumers participate; the profit of the firm is thus Π^G.
- Under pure positioning:
  - for $\delta < \tilde{\delta}$, all consumers inspect product 1 and low-cost consumers are willing to inspect product 2 as well; the profit is thus equal to $\mu_1 \pi^m + \varphi (1 - \mu_1) \mu_2 \pi^m = \Pi^G + (1 - \varphi) \mu_1 \pi^m > \Pi^G$;
  - for $\delta > \tilde{\delta}$, consumers inspect only product 1; the profit therefore drops to $\mu_1 \pi^m$, which, under (20), is strictly lower than Π^G.[80]

R Proof of Proposition 9

The same logic as in the baseline model implies that any steering equilibrium is characterized by a block structure $\boldsymbol{\mathcal{B}} = (\mathcal{B}_1, \ldots, \mathcal{B}_k)$, such that, at the end each block, some active consumers stop searching even in the absence of a match. Consider now a candidate steering equilibrium with a block structure of the form $\boldsymbol{\mathcal{B}} = (\mathcal{B}_1, \ldots)$, where $\mathcal{B}_1 \subsetneq \mathcal{I}_n$, implying that some of the active consumers stop searching after $|\mathcal{B}_1| < n$ inspections. The firm would then have an incentive to position the least popular of the selected products in the slots $\mathcal{B}_1$, and disclose the identities of these products:[81] this would ensure that all consumers $c < \hat{c}_1$ would keep searching in the absence of a match, and thus increase profit.

Conversely, in any of the steering equilibria in which consumers keep searching until finding a match, the firm has no incentive to deviate and disclose the identity of inspected products, as this could only induce some consumers to stop searching – namely, upon learning that they already inspected the most popular product(s). It follows that the possibility of disclosure does not disrupt the steering equilibria that feature intensive search.

[80] A further increase in δ (namely, for $\delta \in \left(\hat{\delta}, \check{\delta}\right)$, where $\check{\delta}$ is such that $\hat{c}(\mathcal{I}_2; \check{\delta}) = \hat{c}$) reverses again the balance in favor of pure positioning, as noisy positioning then discourages all consumers from participating.

[81] This calls for a definition of consumers' beliefs following such a deviation. In the spirit of passive beliefs, we will assume that, upon discovering a product expected to be assigned to a subsequent block, a consumer believes that the firm maintained the equilibrium selection of products and simply swapped the positioning of the least popular product, among the yet undisclosed ones expected in that block, with that of the disclosed but unexpected product, keeping unchanged the positioning of the remaining products.

S Proof of Proposition 10 (platforms)

Fix n, $\boldsymbol{\rho}$, $\boldsymbol{\mathcal{B}}$ and $\hat{\mathbf{c}}$, and consider a candidate equilibrium in which (i) the platform selects the n most popular suppliers, charges them a fixed fee appropriating their expected equilibrium profit,[82] and positions them according to the strategy $\boldsymbol{\rho}$, and (ii) suppliers charge the monopoly price. Obviously, in stage 2c, consumers are then willing to stick to the search pattern characterized by $\boldsymbol{\mathcal{B}}$ and $\hat{\mathbf{c}}$. Furthermore, in stage 2b, the platform is indifferent about its positioning and is thus willing to stick to the strategy $\boldsymbol{\rho}$—*regardless* of the prices set by suppliers; conversely, if prices cannot affect their positions, suppliers find it optimal to charge the monopoly price. Finally, in stage 2a, the platform can indeed charge each selected supplier a fee equal to its expected profit; it thus has an incentive to select the most popular ones, as they generate more profit.

[82]The use of a fixed fee follows from our zero cost assumption. If the platform were incurring a cost per sale, then a two-part tariff should be used, with a wholesale price covering the cost of a sale.

Online Appendix for Consumer Search, Steering and Choice Overload

In this Online Appendix, we construct an equilibrium for the case $n = 2$ in which the firm offers its products at different prices. For simplicity, we assume that all consumers who start searching face the same search cost, c. All other assumptions are as in the baseline setting of Section 2.

A Candidate Equilibrium

We consider a candidate equilibrium in which the firm offers product 1 at the monopoly price ($p_1^* = p^m$) and product 2 at a moderately lower price ($p_2^* < p^m$), in such a way that product 1 still generates the greater expected surplus: $\mu_1 s^m > \mu_2 s_2^*$, where $s_2^* \equiv s\left(p_2^*\right)$. In what follows, we fix $F\left(\cdot\right)$, μ_1, and the equilibrium price p_2^*, and characterize the relevant ranges of c and μ_2 for which there exists such an equilibrium; the above condition amounts to:

$$\mu_1 s^m > \mu_2 s_2^* \Longleftrightarrow \mu_2 < \bar{\mu}_2 \equiv \mu_1 \frac{s^m}{s_2^*}. \tag{21}$$

We further focus on an equilibrium with passive beliefs in which consumers make a first inspection, and then inspect the remaining product only if they observe p_2^* and have no match. To ensure that consumers start searching, we must have:

$$\begin{aligned} c &< \frac{1}{2}\mu_1 s^m + \frac{1}{2}\left[\mu_2 s_2^* + \left(1-\mu_2\right)\left(\mu_1 s^m - c\right)\right] \\ &\Longleftrightarrow \quad c < \bar{c}\left(\mu_2\right) \equiv \frac{\left(2-\mu_2\right)\mu_1 s^m + \mu_2 s_2^*}{3-\mu_2}. \end{aligned} \tag{22}$$

For the second inspection decision we must have:

- In the absence of a match at the first inspection, upon observing an unexpected price $p \notin \{p^m, p_2^*\}$, consumers stop searching:

$$c > \hat{c}\left(\mu_2\right) \equiv \frac{\left(1-\mu_2\right)\mu_1 s^m + \left(1-\mu_1\right)\mu_2 s_2^*}{2-\mu_1-\mu_2}. \tag{23}$$

To ensure that this condition is compatible with (22), we will assume that[1]

$$\mu_1 < \frac{1}{2}.$$

Altogether, the above conditions yield:

$$\mu_2 s^m < \hat{c}(\mu_2) < \bar{c}(\mu_2) < \mu_1 s^m.$$

It follows that consumers stop searching if they observe $p = p^m$ (as $\mu_1 s^m > (\bar{c}(\mu_2) >) c$), and keep searching if instead they observe $p = p_2^*$ (as $\mu_2 s_2^* < (\hat{c}(\mu_2) <) c$).

- After a first match, consumers stop searching. This is obviously the case if they observed $p = p_2^* < p^m$; if instead they observed $p \notin \{p^m, p_2^*\}$, this holds if they stop searching even when p is prohibitive (i.e., such that $s(p) = 0$), which amounts to:
$$c > \frac{\mu_2 \mu_1 s^m + \mu_1 \mu_2 s_2^*}{\mu_1 + \mu_2},$$
and is implied by (23).[2] Finally, if they observed $p = p^m$, this requires:
$$c > \tilde{c}(\mu_2) \equiv \mu_2 (p^m - p_2^*).$$

It follows that all the above conditions are satisfied if $\mu_2 < \bar{\mu}_2$ and $c \in (\underline{c}(\mu_2), \bar{c}(\mu_2))$, where:

$$\underline{c}(\mu_2) \equiv \max\{\tilde{c}(\mu_2), \hat{c}(\mu_2)\}.$$

As μ_2 goes to zero, $\underline{c}(\mu_2)$ tends to $\hat{c}(0) = \mu_1 s^m / (2 - \mu_1)$, whereas $\bar{c}(\mu_2)$ tends to $2\mu_1 s^m / 3$, which exceeds $\hat{c}(0)$ for $\mu_1 < 1/2$; hence, for any given $\mu_1 < 1/2$, there exists a non-empty search cost range for μ_2 small enough.

[1] We have:
$$\bar{c}(\mu_2) - \hat{c}(\mu_2) = \frac{1 - 2\mu_1 + \mu_1 \mu_2}{(3 - \mu_2)(2 - \mu_1 - \mu_2)} (\mu_1 s^m - \mu_2 s_2^*) > 0,$$
where the inequality follows from (21) and $\mu_1 < 1/2$.

[2] We have:
$$\hat{c}(\mu_2) - \frac{\mu_2 \mu_1 s^m + \mu_1 \mu_2 s_2^*}{\mu_1 + \mu_2} = \frac{(\mu_1 - \mu_2)(\mu_1 s^m - \mu_2 s_2^*)}{(2 - \mu_1 - \mu_2)(\mu_1 + \mu_2)} > 0.$$

B Possible Deviations

In the above candidate equilibrium, the firm obtains an expected profit equal to:

$$\begin{aligned}\Pi^* &\equiv \frac{1}{2}\mu_1\pi^m + \frac{1}{2}\left[\mu_2\pi_2^* + (1-\mu_2)\,\mu_1\pi^m\right] \\ &= \frac{(2-\mu_2)\,\mu_1\pi^m + \mu_2\pi_2^*}{2},\end{aligned}$$

where $\pi_2^* \equiv \pi(p_2^*)$. As μ_2 goes to zero, this expected profit tends to

$$\Pi_0^* \equiv \mu_1\pi^m.$$

We can distinguish three types of deviations, depending on which prices are affected.

B.1 Single deviation on p_1

A single deviation on the price of the product 1 from p^m to $p_1 \notin \{p^m, p_2\}$ does not affect consumers' search behavior: they stop searching after the first inspection unless they encountered p_2 and had no match, as along the equilibrium path. It follows that such deviation cannot be profitable, as it simply replaces the monopoly profit π^m with a lower profit $\pi(p_1) < \pi^m$ in case of a match with the product 1.

A single deviation from $p_1 = p^m$ to $p_1 = p_2^*$ induces instead consumers to keep searching in the absence of a match (and stop searching otherwise); hence, it yields:

$$\begin{aligned}\hat{\Pi} &\equiv \frac{1}{2}\left[\mu_1\pi_2^* + (1-\mu_1)\,\mu_2\pi_2^*\right] + \frac{1}{2}\left[\mu_2\pi_2^* + (1-\mu_2)\,\mu_1\pi_2^*\right] \\ &= \left[\mu_1 + \mu_2 - \mu_1\mu_2\right]\pi_2^*.\end{aligned}$$

We have:

$$\begin{aligned}\Pi^* - \hat{\Pi} &= \frac{1}{2}\mu_1\pi^m + \frac{1}{2}\left[\mu_2\pi_2^* + (1-\mu_2)\,\mu_1\pi^m\right] \\ &\quad -\frac{1}{2}\left[\mu_1\pi_2^* + (1-\mu_1)\,\mu_2\pi_2^*\right] - \frac{1}{2}\left[\mu_2\pi_2^* + (1-\mu_2)\,\mu_1\pi_2^*\right] \\ &= \frac{(2-\mu_2)\,\mu_1\left(\pi^m - \pi_2^*\right) - (1-\mu_1)\,\mu_2\pi_2^*}{2}.\end{aligned}$$

This deviation is therefore unprofitable as long as:

$$(2-\mu_2)\,\mu_1\left(\pi^m - \pi_2^*\right) > (1-\mu_1)\,\mu_2\pi_2^*. \tag{24}$$

As μ_2 goes to zero, the left-hand tends to $2\mu_1\left(\pi^m - \pi_2^*\right) > 0$ whereas the right-hand side tends to 0. Hence, there exists $\hat{\mu}_2$ such that this deviation is unprofitable as long as $\mu_2 < \hat{\mu}_2$.

B.2 Other deviations

Deviating on the price of the second product—in isolation or combined with deviating on the price of the first product—induces all consumers to stop searching after the first inspection, regardless of whether there is a match; hence, a deviation to $(\tilde{p}_1, \tilde{p}_2)$ (where $\tilde{p}_1 = p^m$ in case of an isolated deviation on p_2, and $\tilde{p}_1 \neq p^m$ in case of a simultaneous deviation on both prices) yields:

$$\tilde{\Pi} \equiv \frac{1}{2}\mu_1 \pi(\tilde{p}_1) + \frac{1}{2}\mu_2 \pi(\tilde{p}_2).$$

Using $\pi(\tilde{p}_i) \leq \pi^m$ and $\pi_2^* \geq 0$, we have:

$$\begin{aligned} \Pi^* - \tilde{\Pi} &\geq \frac{1}{2}\mu_1 \pi^m + \frac{1}{2}(1-\mu_2)\mu_1 \pi^m - \frac{1}{2}\mu_1 \pi^m - \frac{1}{2}\mu_2 \pi^m \\ &= (\mu_1 - (1+\mu_1)\mu_2)\frac{\pi^m}{2}. \end{aligned}$$

It follows that this deviation is unprofitable whenever

$$\mu_2 < \tilde{\mu}_2 \equiv \frac{\mu_1}{1+\mu_1}.$$

C Existence

For any $\mu_1 \in (0, 1/2)$ and any $p_2^* < p^m$, the above strategies constitute an equilibrium for any $\mu_2 \in (0, \min\{\bar{\mu}_2, \hat{\mu}_2, \tilde{\mu}_2\})$ and any $c \in (\underline{c}(\mu_2), \bar{c}(\mu_2))$.

Illustration: Suppose that match valuations are uniformly distributed over $[0, 1]$; we have:

$$\pi(p) = p(1-p) \text{ and } s(p) = \frac{(1-p)^2}{2},$$

and thus:

$$p^m = \frac{1}{2}, s^m = \frac{1}{8} \text{ and } \pi^m = \frac{1}{4}.$$

Fix $p_2^* = p^m/2 = 1/4$; we have

$$\begin{aligned} s_2^* &= \left[\frac{(1-p)^2}{2}\right]_{p=\frac{1}{4}} = \frac{9}{32}, \\ \pi_2^* &= [p(1-p)]_{p=\frac{1}{4}} = \frac{3}{16}, \end{aligned}$$

and:

$$\bar{\mu}_2 = \mu_1 \frac{\frac{1}{8}}{\frac{9}{32}} = \frac{4\mu_1}{9} < \tilde{\mu}_2 \equiv \frac{\mu_1}{1+\mu_1}.$$

Condition (24) amounts to:

$$\begin{aligned}
0 &< \left[(2-\mu_2)\,\mu_1\,(\pi^m - \pi_2^*) - (1-\mu_1)\,\mu_2\pi_2^*\right]_{\pi^m=\frac{1}{4},\pi_2^*=\frac{3}{16}} \\
&= \frac{1}{8}\mu_1 - \frac{3}{16}\mu_2 + \frac{1}{8}\mu_1\mu_2 \\
&\Leftrightarrow \mu_2 < \hat{\mu}_2 = \frac{2\mu_1}{3-2\mu_1},
\end{aligned}$$

where $\hat{\mu}_2$ also exceeds $\bar{\mu}_2$:

$$\hat{\mu}_2 - \bar{\mu}_2 = \frac{2\mu_1}{3-2\mu_1} - \frac{4\mu_1}{9} = \frac{2\mu_1\,(3+4\mu_1)}{9\,(3-2\mu_1)} > 0.$$

Finally, we have:

$$\begin{aligned}
\hat{c}\,(\mu_2) &= \left[\frac{(1-\mu_2)\,\mu_1 s^m + (1-\mu_1)\,\mu_2 s_2^*}{2-\mu_1-\mu_2}\right]_{s^m=\frac{1}{8},s_2^*=\frac{9}{32}} = \frac{4\mu_1+9\mu_2-13\mu_1\mu_2}{32\,(2-\mu_1-\mu_2)}, \\
\tilde{c}\,(\mu_2) &= \left[\mu_2\,(p^m - p_2^*)\right]_{p^m=\frac{1}{2},p_2^*=\frac{1}{4}} = \frac{\mu_2}{4}, \\
\bar{c}\,(\mu_2) &= \left[\frac{(2-\mu_2)\,\mu_1 s^m + \mu_2 s_2^*}{3-\mu_2}\right]_{s^m=\frac{1}{8},s_2^*=\frac{9}{32}} = \frac{8\mu_1+9\mu_2-4\mu_1\mu_2}{32\,(3-\mu_2)}
\end{aligned}$$

As expected, $\bar{c}\,(\mu_2) > \hat{c}\,(\mu_2)$:

$$\frac{8\mu_1+9\mu_2-4\mu_1\mu_2}{32\,(3-\mu_2)} - \frac{4\mu_1+9\mu_2-13\mu_1\mu_2}{32\,(2-\mu_1-\mu_2)} = \frac{(4\mu_1-9\mu_2)\,(1-2\mu_1+\mu_1\mu_2)}{32\,(2-\mu_1-\mu_2)\,(3-\mu_2)},$$

where the right-hand side is positive for $\mu_1 \in (0, 1/2)$ and $\mu_2 \in (0, 4\mu_1/9)$. In addition, $\bar{c}\,(\mu_2) > \tilde{c}\,(\mu_2)$ as long as:

$$\begin{aligned}
0 &> \frac{8\mu_1+9\mu_2-4\mu_1\mu_2}{32\,(3-\mu_2)} - \frac{\mu_2}{4} = \frac{8\mu_1-15\mu_2+8\mu_2^2-4\mu_1\mu_2}{32\,(3-\mu_2)} \\
&\Leftrightarrow \mu_2 < \check{\mu}_2 \equiv \frac{15+4\mu_1-\sqrt{225-136\mu_1+16\mu_1^2}}{16}
\end{aligned}$$

where the threshold $\breve{\mu}_2$ exceeds $\bar{\mu}_2$:

$$\begin{aligned}
\breve{\mu}_2 - \bar{\mu}_2 &= \frac{15 + 4\mu_1 - \sqrt{225 - 136\mu_1 + 16\mu_1^2}}{16} - \frac{4\mu_1}{9} \\
&= \frac{15 - \frac{28}{9}\mu_1 - \sqrt{225 - 136\mu_1 + 16\mu_1^2}}{16} \\
&= \frac{\left(15 - \frac{28}{9}\mu_1\right)^2 - \left(225 - 136\mu_1 + 16\mu_1^2\right)}{16\left(15 - \frac{28}{9}\mu_1 + \sqrt{225 - 136\mu_1 + 16\mu_1^2}\right)} \\
&= \frac{128}{1296}\frac{\mu_1\left(27 - 4\mu_1\right)}{15 - \frac{28}{9}\mu_1 + \sqrt{225 - 136\mu_1 + 16\mu_1^2}} \\
&> 0,
\end{aligned}$$

where the inequality stems from the fact that the numerator and the denominator of the last expression are both positive for any $\mu_1 \in (0, 1)$.

It follows that, for $p_2^* = 1/4$, there exists an equilibrium such as described above for any $\mu_1 \in (0, 1/2]$, any $\mu_2 \in (0, 4\mu_1/9)$ and any $c \in (\underline{c}(\mu_2), \bar{c}(\mu_2))$, where $\underline{c}(\mu_2) = \max\{\tilde{c}(\mu_2), \hat{c}(\mu_2)\}$.

D Profitability

The above equilibrium can be more profitable than the monopoly price equilibrium, by encouraging more consumers to search. The upper bound on the search cost, given by (22), is indeed lower than that for monopoly pricing, which is given by:

$$\underline{c}_2 = \frac{\mathrm{M}_2}{\Gamma_2} s^m,$$

where

$$\mathrm{M}_2 = \mu_1 + \mu_2 - \mu_1\mu_2 \text{ and } \Gamma_2 = 2 - \frac{\mu_1 + \mu_2}{2}.$$

Indeed, we have:

$$\begin{aligned}
\bar{c}(\mu_2) - \underline{c}_2 &= \frac{(2 - \mu_2)\mu_1 s^m + \mu_2 s_2^*}{3 - \mu_2} - \frac{\mu_1 + \mu_2 - \mu_1\mu_2}{1 + \frac{1-\mu_1}{2} + \frac{1-\mu_2}{2}} s^m \\
&> \frac{(2 - \mu_2)\mu_1 + \mu_2}{3 - \mu_2} s^m - \frac{\mu_1 + \mu_2 - \mu_1\mu_2}{1 + \frac{1-\mu_1}{2} + \frac{1-\mu_2}{2}} s^m \\
&= (1 - \mu_1)\frac{2 - \mu_2}{3 - \mu_2}\frac{\mu_1 - \mu_2}{4 - \mu_1 - \mu_2} s^m \\
&> 0.
\end{aligned}$$

It follows that, in the above illustration, for any $c \in \left(\max\left\{\underline{c}\left(\mu_2\right), \underline{c}_2\right\}, \bar{c}\left(\mu_2\right)\right)$:

- the monopoly pricing equilibrium would generate no search, and therefore yield zero profit.[3]
- by contrast, the above equilibrium induces all consumers to search, and yields a positive profit.

[3]Although the firm is there indifferent about the product portfolio and its prices, introducing an infinitesimal number of consumers with low enough search costs would not materially affect the analysis but would eliminate this indifference.