How (Much) Can Platforms Learn from Consumer Choices?*

Maarten C.W. Janssen† Eeva Mauring‡

November 15, 2024
(VERY PRELIMINARY VERSION)

Abstract

Consumers search on a platform to learn how a product fits their preferences. Consumers' value for the product has a common and an idiosyncratic component. The platform observes which products consumers inspect and what they eventually buy. Based on these observations on past consumer choices the platform ranks products. We find that monopoly platforms maximizing long-term consumer's utility will first experiment with product rankings, providing future consumers with rankings that only list products that many consumers have bought before. This guarantees that consumers are more picky and only buy products they really like. The more important the idiosyncratic component, the better the platform is able to assist consumers in their search. In contrast, learning is much more restricted if platforms compete.

JEL Classification: D40, D83, L10

Key-words: Online platforms; learning; consumer search; product rankings; steering.

*We have benefitted considerably from comments by Mark Armstrong, Heski Bar-Issac and Cole Williams and by participants to the virtual CEPR IO seminar and the 2024 consumer search workshop (Istanbul).

†University of Vienna and CEPR. E-mail: maarten.janssen@univie.ac.at.

‡University of Bergen and CEPR. E-mail: Eeva.Mauring@uib.no.

1 Introduction

Online platforms possess large amounts of data on items people have searched for. This can be data about products, services, news or other items. Platforms also observe whether or not a consumer buys a product or service or how long a consumer spends on a news item (as an indication of whether the consumer appreciates the item). Platforms may employ this information to help users find items they like, but there are increasing concerns that platforms may use this information in other ways without benefiting consumers. Regulators around the world are putting rules in place to restrict the ways platforms can use information (see, e.g. the EU's Digital Market Act (DMA) and the Digital Service Act (DSA)). For example, the Digital Market Act forbids that gatekeepers "cross-use personal data from the relevant core platform service in other services provided separately by the gatekeeper, including other core platform services, and vice versa",[1] and states "Gatekeepers should therefore be required to provide access, on fair, reasonable and non-discriminatory terms, to those ranking, query, click and view data in relation to free and paid search generated by consumers on online search engines to other undertakings providing such services, so that those third-party undertakings can optimise their services and contest the relevant core platform services."[2]

What platforms can infer from the information they collect does not only depend on the search behaviour of their users, but also on how they rank the items. In particular, if they present consumers always with the same ranking of items ussers may make similar choices, restricting the type of inferences platforms make as it is not clear whether the user choices reflect "true preferences" or the ranking that is chosen by the platforms themselves. By experimenting with different rankings, and observing behaviour given different rankings, platforms potentially may learn more about users' preferences. How much platforms learn depends, however, also on how much consumers are willing to search for alternative items, which depends, among other things, on their search cost.

In this paper we analyze what a platform can learn by optimally choosing its ranking. We consider that a platform wants to maximize the well-being of its users in the long run, for example because its main revenue source derives from advertising or selling information to interested parties. It may also be that it is interested in user well being to increase the market size. We compare how much can be learned both in case a platform is a monopolist

[1] See, the Digital Market Act, Article 5 2(c) (European Parliament, 2022b).

[2] See, the Digital Markets Act, paragraph 61 (European Parliament, 2022a).

or faces competition.

We focus on online markets where consumers' valuations for a product have a common and an idiosyncratic component, but the ideas of the paper can equally be applied to news aggregators or recommendation platforms. Platforms observe the items consumers inspect and which item they buy. They choose an algorithm that maps the observed history into a ranking of items for the next consumer. Consumers inspect products sequentially and their behaviour is characterized by an optimal stopping rule. They can follow the ranking if they like, but are free to search on the platform in any order they prefer. Consumers use the platform sequentially (one by one), but they themselves may not know their position in the queue; while the platform(s) always know(s) how many consumers have visted their platform and how they behaved. Platforms do not observe, however, the actual value a consumer attaches to the product. They only know whether or not consumers have inspected a good and whether it was bought. We are interested in how much the platform eventually may learn, and how this depends on whether platforms face competition, i.e., we look at the steady state of the process described. As individual consumers only use the platform once, learning is about the common value component of consumers' preferences.

Before presenting our results, it is important to mention that the environment we study has some similarities to the social learning literature with seminal contributions by Banerjee (1992), Smith and Sørensen (2000), Bikhchandani et al. (1992), Kremer et al. (2014) and Glazer et al. (2021), among others, with some important modifications. First, in online markets, and in our paper, consumers can choose one out of many products and a platform *communicates* their information through a *ranking* of these products. This type of communication (and choice) is different from advising a consumer whether or not to adopt one product, which is the core of the social learning literature. Second, consumers have to pay an inspection cost if they want to learn their value for a product and they are not restricted to search only a limited set of products. Consumers may decide to stop searching and buy a product simply because they believed it was "good enough" in the sense that they did not want to inspect another unknown item at an inspection cost. In contrast, if there is information acquisition in the social learning literature, as in Glazer et al. (2021), it is limited to one signal about one option. These two factors impose natural limits on what a platform may learn. Third, our focus is on how much a platform may learn, and unlike most of the social learning literature there is no right or wrong choice.

Our main result for a monopoly platform is a characterization of a unique steady state that shows that by experimenting with rankings the platform can learn much more about the common value of the best products than by

following other ranking algorithms. To see why, consider that the idiosyncratic component is absent so that all consumers have the same value for the product. *Without experimentation*, the platform would suggest to all future consumers the item that was bought first, which is a product that has a (common) value that is at least equal to the consumers' ex ante reservation value, as in Wolinsky (1986). Subsequent consumers would always buy the same product as they expect to inspect a random product on their next search. *With experimentation*, the platform can hide for a first set of consumers the products that consumers before them have bought by randomizing recommendations. It can gradually shift to showing consumers a ranking where products that have been bought before are placed first. Doing so increases consumers' continuation value of search as they believe that if they continue to search they are more likely seeing products with a high value compared to the situation where the platform does not experiment. Thus, in expectation consumers are offered products with a much higher common value than without experimentation. However, for any positive inspection cost, there is a limit to what the platform can learn and it is not only the very best products that are ranked first. The larger the inspection cost, the less the platform learns.

The steady state is such that the larger the weight on the idiosyncratic component, the better the platform is able to learn which products have the highest common values. If consumers' value incorporates an idiosyncratic component, they themselves have a larger incentive to search and experiment. In particular, some consumers may be unhappy with products the platform has ranked high, simply because they draw a low idiosyncratic component of their match value. The larger the weight of the idiosyncratic component, the more consumers experiment themselves and if this weight is above a critical value, the platform eventually will learn which products have a common value component that is arbitrarily close to the highest possible value. Thus, the less the consumers value the common value component, the more the platform can learn about it.

We show that these results are robust to a variety of different settings. We consider that consumers know or do not know their position in the queue, that consumers have different search cost and that only a fraction of consumers are informed about the algorithm that is chosen by the platform. We also consider that a fraction or all consumers have the options to buy blindly.

Our second set of results pertains to competition between platforms. We show that competition significantly limits the possibility for platforms to experiment. The reason is best understood by reconsidering the result on monopoly platforms for the case where consumers know their position in

the queue.[3] Here, by experimenting, the platform consciously chooses at any point in time not to provide the best possible ranking for the current consumer using the current information. For example, without idiosyncratic component, once a consumer has bought a certain product, the next consumer would be best served by a ranking where this product is offered first. Experimentation therefore trades off the benefit of present consumers in favour of future consumers. Competition, however, implies that when they have a choice consumers prefer to use a platform that is not experimenting, leaving platforms that do experiment without consumers to experiment on. Proposition 7 formalises this intuition and shows that under competition platforms recommend a product until the moment a consumer buys another product. Sooner or later, and certainly in the steady state, platforms recommend a product that all (future) consumers buy as the common component of the valuation is high enough so that even consumers with the lowest possible idiosyncratic component will buy (given a random continuation value of search).

Competition may also come in the form of a platform entering the market after an incumbent already has served consumers. We show that a more efficient entrant that has created a more convenient website design where consumers have lower search costs may not be able to gain market share, because the incumbent platform already has been able to experiment with consumers until entry took place. Consumers may therefore prefer to use the incumbent platform as they are more likely to get better recommendations even if it means to incur higher search costs.

Even though the monopoly analysis also requires to carefully address the issue of learning, learning is even more critical under competition, especially when consumers do not know their position in the queue. Receiving a bad draw from a product that the platform recommended makes consumers understand it is more likely they are located earlier in the queue which may give rise to consumers switching platforms. In addition, coordination incentives may mean that a consumer prefers to join a platform that offers an inferior ranking (e.g., that experiments a lot) just because he believes many other consumers to have joined the platform. We show, however, that if consumers have some information about their position in the queue, but do not know their positions precisely (specifically, we consider consumers knowing their position up to two positions), consumers still prefer a platform that experiments less to one that experiments more which leads to limited learning under competing platforms.

[3]Proposition **??** in the Online Appendix shows that the same steady state is reached under a monopoly whether consumers do or do not know their position in the queue.

Related Literature. As described above, our paper relates to the vast literature on social learning, with important differences. In most of the classic social learning papers, agents know their position in the queue and learn by directly observing the actions of all past agents (e.g., Banerjee (1992), Bikhchandani et al. (1992), Smith and Sørensen (2000)). In Kremer et al. (2014) and Glazer et al. (2021), agents do not observe the choices of preceding agents, but a principal does and recommends to future agents choices based on this information. In contrast to these papers, we allow for the consumers' (agents') payoffs to include an idiosyncratic component. We also allow consumers to search products in any order they like and to continue searching beyond the recommended product. Maglaras et al. (2021) studies an environment where consumers learn from reviews of previous consumers and pay a larger "search cost" for an item that is ranked lower by the platform. There is no real search in the model as consumers simply buy the product that maximizes their expected utility before inspecting. Goeree et al. (2006) study an (otherwise classic social learning) environment where the agents' valuations comprise both ann idiosyncratic and a common part. They find that in contrast to the herding result of the social learning literature, with an idiosyncratic component to their utility function, agents will eventually always learn the true state of the world. In contrast, our paper focuses on platforms trying to influence the outcome of the social learning process and on consumers searching different alternatives. Our paper is also closely related to other papers on platforms that recommend products to consumers and learn about product quality.[4]

The interaction between a platform's rankings and consumer search behaviour is also studied from an empirical angle (see, e.g., De Los Santos and Koulayev (2017) and Ursu (2018)). The main issue in that literature is how to estimate the impact of a platform's ranking on what consumers click on or what they buy as the ranking is also influenced by what consumers' past behaviour. The question in this paper, namely how a platform may choose rankings to learn consumer preferences and how that depends on the market structure, is not addressed in the empirical literature, however.

There is also an important literature on directed search and the effects of prominence (see, e.g., Arbatskaya (2007), Armstrong et al. (2009), Wilson (2010), Haan and Moraga-Gonzalez (2011), Rhodes (2011) and Zhou (2011) for early papers on the topic). Most of that literature has focussed on the effect of prominence (or search order more generally) on prices, how-

[4]A literature also exists, such as Teh and Wright (2022), Janssen and Williams (2023), Janssen et al. (2023) and Bar Isaac and Shelegia (2022), where platforms or a social influencer steer consumers to products, but there is no learning.

ever, whereas our focus is on how platforms can use rankings to learn about consumer preferences.

There is also a rapidly growing literature on how platforms can use their knowledge about consumers searching for keywords to create rankings that maximize their profits (see, e.g., Anderson and Renault (2021), Ke et al. (2023) and Janssen et al. (2024)). To the best of our knowledge, our paper is the first to develop a framework analyzing how a platform should adjust its rankings to best learn from consumer choices.

More generally, our paper relates to the literature on dynamic information design (see, for example, Ely (2017), Renault et al. (2017), Ely and Szydlowski (2020), Orlov et al. (2020), Smolin (2021)). In most of this literature, the principal's learning is exogenous to the actions of the agent(s), while in our model the consumers' actions determine what and how much the platform learns.

We find that a monopoly platform has an incentive to experiment so that one consumer can (indirectly) learn from the actions of another, as in strategic experimentation literature (see, e.g., and Bolton and Harris (1999) or Keller et al. (2005)). In our model, experimentation by the platform means that consumers on average search more than they would if the platform does not experiment. In Hagiu and Jullien (2011), a profit-maximizing platform sometimes makes consumers search more than they would like (in order to influence sellers' behaviour), but there is no learning.

The rest of the paper is organized as follows. The next Section describes the monopoly model, while Section 3 states the main monopoly results and its extensions. Section 4 discusses the results pertaining to competition. Section 5 concludes with a discussion, while proofs can be found in the Appendix.

2 The Baseline Monopoly Model

We consider a population of infinitely many consumers using a platform to look for a product. Products are horizontally and vertically differentiated, i.e., they have features that all consumers value in the same way (a common component) and other features that consumers value differently (an idiosyncratic component). Thus, a consumer i's utility for product j is denoted by

$$u_{ij} = (1-\delta)v_j + \delta\varepsilon_{ij},$$

where v_j is the common component, ε_{ij} is the idiosyncratic component and $\delta \in [0,1]$ is the relative weight on the idiosyncratic component. Before inspecting a product, the consumer is uncertain about both components of the product. The common component v_j is a random draw from the CDF

$G(\cdot)$ with support $[\underline{v}, \bar{v}]$, while the idiosyncratic component ε_{ij} is a random draw from the CDF $F(\cdot)$ with support $[\underline{\varepsilon}, \bar{\varepsilon}]$. We assume that $f(\cdot)$ and $g(\cdot)$ are strictly positive over the whole support and that $1-F$ and $1-G$ are logconcave (as is usual in the consumer search literature; see, e.g., Anderson and Renault (1999)). To know a product's value the consumer has to pay an inspection (or search) cost s and then learns the utility u_{ij} that product j gives. In the baseline model all consumers have he same search cost s and have to inspect the prroduct before buying. In extensions, we consider that consumers have different search cost and have the option to "buy blindly", i.e., to buy without knowing the product's value. The consumer chooses an optimal sequential search strategy and stops searching when the utility exceeds a certain threshold reservation utility. To avoid an uninteresting analysis where consumers never inspect products if the products are randomly ordered, we take s to be small enough, i.e., $s < \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} (1-F(\varepsilon))d\varepsilon = E\varepsilon - \underline{\varepsilon}$ and $s < \int_{\underline{v}}^{\bar{v}} (1-G(v))dv = Ev - \underline{v}$. We sometimes illustrate our results for the ε's and the v's being uniformly distributed over the interval $[0,1]$, and for this case this implies that $s < \frac{1}{2}$.

The platform observes which products a consumer inspects on their platform and which product they eventually buy, but it does not know the utility scores of each product with each consumer. The only thing the platform can infer is that the consumer's utility for products that are inspected, but not purchased is below a reservation utility, while the product that is bought has a utility level above the threshold. For each consumer visiting the platform, it can create a (different) ranking of firms and this ranking may influence the order in which consumers search. For example, if a product is bought by a consumer before, then the platform may (or may not) rank that product first, and a first ranked product may be interpreted by consumers as a product that the platform thinks they should inspect first.

The platform observes which consumers already have inspected and bought products. Consumers know that others are also looking for products on the platform and that therefore the ranking of products may be informative of what others have inspected and bought in the past. In the baseline model, consumers know their place in the queue and they also know the ranking algorithm the platform uses. In extensions, we consider that consumers do not know their position in the queue and that a positive fraction of consumers do not know the algorithm the platform uses.[5] The platform maximizes the

[5]It is clear that if none of the platform's users are informed about the ranking algorithm, the outcome fully depends on consumer expectations. For example, if all consumers expect the platform to rank the products randomly, then the platform cannot learn much. We will consider this case in subsection 2.2 where we provide a preliminary analysis. Note

long-run utility of consumers. That is, it may experiment with rankings for some finite number of consumers without impacting the value of its objective function. It can do so to better learn consumer preferences to be able to make better rankings in the future. Firms are passive.

To be more precise, let products have a number $k \in \mathbb{N}$ and denote by N the set of all subsets of $\mathbb{N}$. Let $S(t) \in N$ be the set of firms that a consumer t has inspected who had $t-1$ consumers making inspection and purchase decisions before him and by $b(t) \in \mathbb{N} \cup \varnothing$ the purchase decision of that consumer. For every consumer t the platform chooses a ranking r_t, which is a function $r_t : \{S(\tau), b(\tau)\}_{\tau=1}^{\tau=t-1} \times Z \to X$, where Z is the realization of a random device and X denotes the set of firm permutations. The overall ranking algorithm r of the platform is a sequence of those rankings: $r = \{r_t\}_{t=1}^{\infty}$. In the baseline model, the platform maximizes average utility

$$\lim_{T\to\infty} \sum_{t=1}^{T} \frac{Eu(t)}{T},$$

where $Eu(t)$ is the expected utility of the consumer in position t.

We consider the steady state of the dynamic process where at each period t, the platform creates a ranking for consumer t and the consumer engages in optimal sequential search given this ranking. To this end, let the distribution over the values of the items' common components in the platform's ranking at time t be denoted by R_t, i.e., $R_t = \{R_t(v_k \leq v)\}_{k=1}^{\infty}$ for all v where $R_t(v_k \leq v) := Pr(\text{item on rank } k \text{ at time } t \text{ has a common value less than or equal to } v)$. Formally, we define a steady state as follows.

Definition 1 *A steady state is a stationary equilibrium where R_t is time invariant, i.e., $R_t = R_{t+1} = R$, and the consumer's optimal stopping rule is a stationary cutoff rule u^*.*

It is clear that if $v^* < \bar{v}$, then in a steady state the platform does not learn anymore from consumer decisions as consumers will always first inspect the first item in the platform's ranking and immediately buy from the firm in that position. In that case a stationary stopping rule is characterised by $u^* = (1-\delta)v^* + \delta\underline{\varepsilon}$ for some $v^* < \bar{v}$ and $R_t(v_k \leq v) = \frac{G(v)}{1-G(v^*)}$ for $v \geq v^*$ for some $k < \infty$ and zero otherwise. If $v^* = \bar{v}$, then consumers may continue to search after inspecting the first item, namely if $u^* > (1-\delta)v^* + \delta\underline{\varepsilon}$ and their idiosyncratic component is so low that $(1-\delta)v^* + \delta\varepsilon_{ij} < u^*$. In that

that the DSA requires that platforms are transparent about the ranking algorithm they use so that the assumption that at least some users are informed about the algorithm is not farfetched.

case a stationary stopping rule is characterised by $u^* > (1-\delta)\bar{v} + \delta\underline{\varepsilon}$, and $R_t(v_k \leq v) = 0$ for all $v < \bar{v}$ and $R_t(v_k \leq \bar{v}) = 1$ for some $k < \infty$. We show later that these distributions R_t constitute steady states.

2.1 Comments on the Baseline Model

This paper focuses on the platform learning the common component in consumers' preferences for one product. The platform may, of course, also be interested in other forrms of learning, such as trying to understand an individual consumer's preferences by observing many searches for different products by the same consumer. such learning requires that there are consumer features that show up in all the choices the consumer makes across different product categories. We see the current paper as making a beginning in asking the question what platforms may learn by creating rankings, but it is certainly not the last paper on the topic.

The modeling does not explicitly include prices, but one may think that price is one of the elements that makes up the common component in consumers' utility function. Thus, we implicitly model thaat prices can differ across different products. The reason we do not explicitly introduce prices is that we want to focus on what platforms may learn from consumer choices, while if we would include price setting behaviour by firms, the platform may potentially also learn about product characteristics from firms' price choices. Such a dual source of learning would complicate matters beyond what we can address in this paper.

We focus on long-run consumer surplus as the platform's objective. This may capture the starting phase of a platform interested in increasing market size, or a platform making profit from using information about consumer preferences in other markets. A proper welfare analysis should take these other markets into account when assessing the desirability of a platform learning consumer characteristics. A platform may have other objectives, however, such as maximizing sales commissions and/or maximizing revenue from a sponsored slot. Another way to interpret the current paper is that it determines the maximum amount of learning as with any other type of objective, the platform is, at best, only indirectly interested in assisting consumers to find the best match quickly. For instance, if the platform would be interested in maximizing revenue from a sponsored position, then it may purposefully obfuscate information beyond the sponsored position to inncrease the likelihood the consumer buys from the sponsored position (see Janssen et al. (2024)).

As mentioned in the Introduction, we focus on ranking products, but the analysis could equally be applied to other platforms, such as news aggregators

or recommendation services, where rankings are important ways to communicate information. What is crucial for our analysis to hold is that (i) users search among many items, (ii) that platforms can influence how users search and (iii) that platforms get a clear indication whether consumer positively assess certain items, and not others. In our model, this positive aassessment is the act of purchasing an product, but in alternative application this could be the time a user spends at an item.

Finally, the baseline model makes many more detailed assumptions, such as that δ and s are identical across consumers, that consumers have to inspect a product before purchase and cannot engage in "blind buying" and that consumers know their position in queue. In extensions we show that our qualitative results continue to hold and we detail how the results need to be modified, if at all.

2.2 Preliminary analysis

We first consider as an important reference point for future analysis a myopic platform that simply recommends the item the last consumer has bought (and puts this item prominently on the first position) and places all other products randomly. Note that given the common component in the consumers' utility function such a ranking contains relevant information to consumers and consumers find it optimal to start searching the recommended product. Define u_R^* as the reservation utility under random search, i.e.,

$$u_R^* = E[max\{(1-\delta)v + \delta\varepsilon, u_R^*\}] - s.$$

Given the platform's myopic algorithm, it is clear that any consumer buys an item if the utility it generates is larger than or equal to u_R^*. For any $s > 0, u_R^* < (1-\delta)\bar{v} + \delta\bar{\varepsilon}$.

Using the definition of u_R^* we can define v_R^* by $v_R^* \equiv (u_R^* - \delta\underline{\varepsilon})/(1-\delta)$ if $u_R^* < (1-\delta)\bar{v}+\delta\underline{\varepsilon}$ and $v_R^* \equiv \bar{v}$ if $u_R^* \geq (1-\delta)\bar{v}+\delta\underline{\varepsilon}$. For an arbitrary reservation utility $\hat{u}$ we can depict the situation as in Figure 1. The reservation utility is a downward sloping line in the (v, ε) space. It crosses the horizontal axis at $\hat{v}$ as in Figure 1 if $\hat{u} < (1-\delta)\bar{v}+\delta\underline{\varepsilon}$. If consumers draw a utility level above this line, they accept. Note that this area consists of a rectangle where the utility level has a common component $v > v^*$ (where all consumers immediately accept such products) and a triangle where the utility level has a common component $v < v^*$ (where only those consumers immediately accept for whom the idiosyncratic component is large enough).

If $v_R^* < \bar{v}$, then it is clear that the myopic platform eventually always recommends the same product. The first consumer buys a product A with

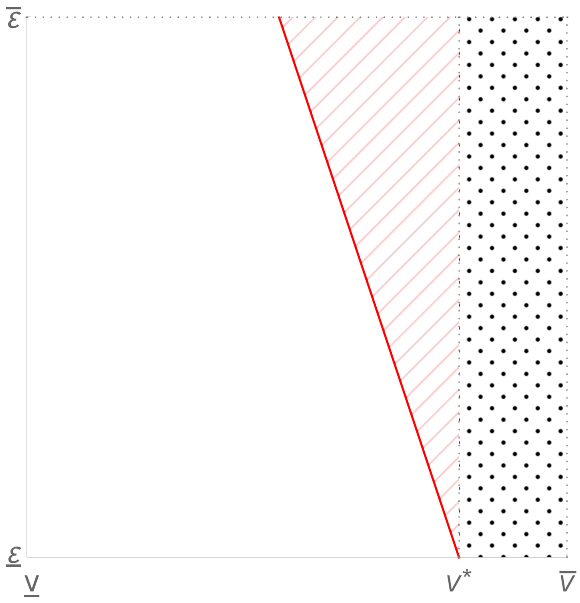


Figure 1: *Consumers accept items north-east from the thick (red) line: in the triangle (hatched) and rectangle (dotted);* $s = 0.2$, $\delta = 0.25$.

$u_A \geq u_R^*$ and this product will be recommended to the second consumer. The second (and all future consumers) will certainly buy this product after inspection if it has a common component $v_A \geq v_R^*$. If on the other hand, $v_A < v_R^*$, then eventually there will be a consumer that has a low idiosyncratic component and that consumer will buy a different product B. This product may also be eventually rejected if $v_B < v_R^*$, but if $v_R^* < \bar{v}$ at some point a product with a $v \geq v_R^*$ will be purchased and then all future consumers will be recommended that product and buy it immediately after inspection. Thus, using such a myopic ranking, the platform will eventually learn that the recommended product has a $v \geq v_R^*$ if $v_R^* < \bar{v}$ or, if $v_R^* \geq \bar{v}$, it will eventually learn which product has a $v = \bar{v}$. Thus, with a myopic ranking, the platform is only able to learn whether or not an item has a $v \geq v_R^*$, with v_R^* implicitly defined by

$$\begin{aligned}(1-\delta)v_R^* + \delta\underline{\varepsilon} \;\geq\; & (1 - \Pr(u < u_R^*))E(u|u \geq u_R^*) \\ & + \Pr(u < u_R^*)\left((1-\delta)v_R^* + \delta\underline{\varepsilon}\right) - s,\end{aligned} \tag{1}$$

or

$$v_R^* \geq E(u|u \geq u_R^*) - \frac{\delta\underline{\varepsilon}}{1-\delta} - \frac{s}{(1-\delta)\Pr(u \geq u_R^*)}.$$

Figure 2 illustrates the cutoffs v_R^* for different values of s in the case of the uniform distribution. For all δ below some critical value $\hat{\delta}_R$ (which depends

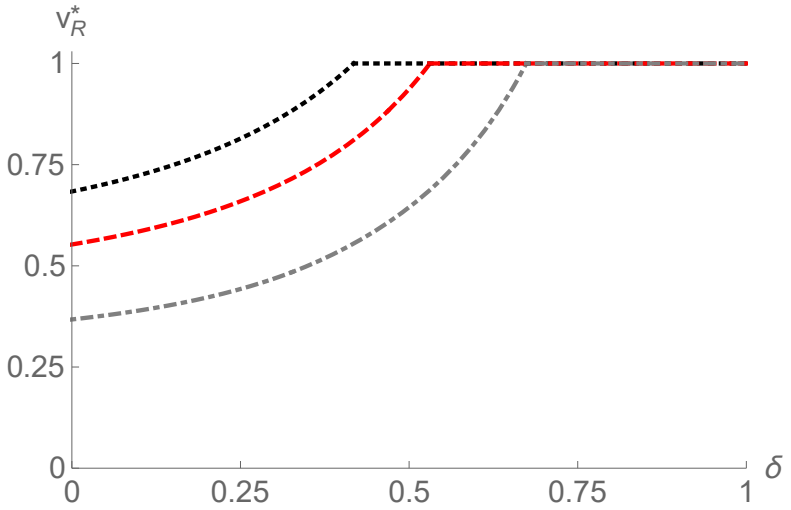


Figure 2: *The myopic cutoff v_R^* as a function of δ for $s = 0.05$ (black dotted), $s = 0.1$ (red dashed), and $s = 0.2$ (grey dot-dashed).*

on s), a cutoff is below the maximum common value of 1. The main question we ask in this paper is how much more the platform may learn by using a different ranking algorithm, when it is a monopolist and/or when it faces competition. From the above analysis it is clear that if no consumer knows the algorithm the platform uses and their behaviour is based on expectations regarding this algorithm that in equilibrium turn out to be correct, then there exists an equilibrium where the platform chooses a myopic algorithm. The reason is that given these expectations, all consumers use the reservation utility u_R^*. The platform can then not do better than using the myopic algorihm as consumers will anyway accept products generaing a utility $u \geq u_R^*$. In the next section we show that the monopoly platform can do (much) better if at least a positive fraction of consumers knows the algorithm.

Using Figure 1 another important preliminary result is that if δ is small enough a consumer never inspects a position k in the ranking of which he knows it contains a product that has been *inspected but not bought* by some previous consumer, i.e., he does not inspect items in the (red) hatched "triangle" region in Figure 1. A platform therefore will not use rankings where it displays previously rejected items in its experiments to learn the consumers' common preferences.

To gain intuition, consider a hypothetical situation where the consumer's reservation utility of picking a random item from a certain subset of products equals $\tilde{u} \geq u_R^*$, with $\tilde{v}$ being defined by $\tilde{v} = (\tilde{u} - \delta\underline{\varepsilon})/(1-\delta)$. Suppose also a consumer knows the common value of an item equals $\tilde{v}$ and the only

uncertainty he faces in picking this item is the value of the idiosyncratic component. A consumer will not want to pick this item above picking a random item with reservation utility $\tilde{u}$ if $(1-\delta)\tilde{v}+\delta E(\varepsilon)-s\leq\tilde{u}$, which is equivalent to $\delta\left(E(\varepsilon)-\underline{\varepsilon}\right)\leq s$, or defining

$$\hat{\delta}=\frac{s}{E(\varepsilon)-\underline{\varepsilon}}$$

if $\delta\leq\hat{\delta}$.[6]

A specific case in point arises when considering $\tilde{u}=u_R^*$. In that case the consumer has always the option of inspecting a random product giving a reservation utility of u_R^* and if $\delta\leq\hat{\delta}$ he would prefer that to buying a product with a common component v_R^*, i.e., a product that under the myopic algorithm considered above would eventually be rejected by a consumer.

The following Lemma now states that if a consumer has the option of inspecting a random product from a certain subset or a product from that same subset that earlier consumers have considered and at least one of them has rejected, he prefers inspecting the random product.

Lemma 2 *For any subset of products if $\delta<\hat{\delta}$, a consumer prefers to inspect a random product from that subsset to a product that has been inspected but not bought by some previous consumer.*

An implication of the lemma is that a platform will never be able to induce a consumer to inspect a product that is rejected by a previous consumer who inspected the product if δ, the weight on the idiosyncratic component, is relatively small.[7] As there are infinitely many products, a consumer can always inspect a random product on which the platform does not have information and would prefer to do so because of the upward potential such a product has. The upper limit $\hat{\delta}$ is such that a consumer does not want to search for a better idiosyncratic component of the utility function even if it gets the lowest possible value $\underline{\varepsilon}$.

Another implication of the Lemma is that, if $\delta<\hat{\delta}$, a platform will only experiment (if at all) with items that have not yet been inspected by placing them prominently in its ranking instead of prominently placing items that have been bought but not yet rejected.

[6]Note that if the consumer has the option of blind buying an item, then if $\delta\leq\hat{\delta}$ he would prefer that option to inspecting the same product.

[7]As consumers are free to inspect products in the order they like, this does not necessarily mean that a platform will not prominently rank such products as consumers may skip prominently ranked products.

3 Results under Monopoly

As the platform maximizes long-term average consumer surplus, its payoff is unaffected by the utility level of the first, finite number of consumers. It can use observations about their search and purchase behaviour to learn about the common component of the consumers' utility. The platform can, for example, experiment with the first consumers and provide them with a (random) ranking (of first products) that is different from the myopic ranking considered above. It can then record the products these consumers inspected and bought. In a later stage, the platform can gradually move to order the products such that only those products that have been bought before are ranked in top positions. These are products that tend to have higher values of the common component v. Under such a ranking algorithm, later consumers are rightfully more picky about buying products as their continuation value of search is higher than when the products are randomly ranked.

The next proposition states what the steady state is of the dynamic process that unfolds.

Proposition 3 *The steady state is unique. If $\delta \leq \hat{\delta}$, the steady state is characterized by a ranking r^* and a reservation utility u^* such that $v^* < \bar{v}$ solves*

$$\int_{v^*}^{\bar{v}} \frac{1-G(v)}{1-G(v^*)} dv = \frac{s-\delta(E\varepsilon-\underline{\varepsilon})}{1-\delta}; \tag{2}$$

v^* is strictly increasing in δ and decreasing in s. If, on the other hand, $\delta > \hat{\delta}$, then the steady state is characterized by $v^* = \bar{v}$.

The following algorithm leads to this steady state: at t, rank on any position a random unsampled item with probability π_t and rank (potentially sampled) items in a decreasing order according to their posterior v with probability $1-\pi_t$, where $\pi_1 = 1$, $\pi_2 < 1$, $\pi_t > 0$ is strictly decreasing in t and $\lim_{t\to\infty} \pi_t = 0$.

As in a steady state with $v^* < \bar{v}$, the consumer will always buy from the first firm in the platform's ranking, the platform can only infer whether firms have a product with a common component that is larger than v^*. The minimum possible utility level a consumer gets in a steady state is then given by $(1-\delta)v^* + \delta\underline{\varepsilon}$. As a consumer drawing this utility level in a steady state expects that on his next searches he only encounters products that the platform possibly also could have recommended, i.e., that have a common component that is at least as large as v^*. Thus, the consumer will buy if, and only if,

$$(1-\delta)v^* + \delta\underline{\varepsilon} \geq (1-\delta)E(v|v \geq v^*) + \delta E\varepsilon - s, \tag{3}$$

where the RHS gives the expected utility level of continuing to search. In a steady state with $v^* < \bar{v}$ this inequality should hold with equality. Note the important difference with the steady state under the myopic ranking rule as defined in (1) in the previous section: instead of knowing that in a steady state only items with a large v are positioned in top spots, under the myopic rule a next searched item may also have a low v.

For v and ε being uniformly distributed with $\bar{v} = 1$, $E\varepsilon = 0.5$ and $\underline{\varepsilon} = 0$ the (3) gives

$$v^* = \frac{1 + v^*}{2} + \frac{\delta - 2s}{2\left(1 - \delta\right)},$$

which can be explicitly solved for

$$v^* = 1 + \frac{\delta - 2s}{1 - \delta}.$$

The expression for v^* is smaller than 1 for $\frac{\delta}{2} < s < \frac{1}{2}$, while for $s < \frac{\delta}{2}$ there is no interior steady state: the platform will eventually learn the common component of many products and recommends the one with the largest v. In the steady state, this largest value is $\bar{v} = 1$. The proposition states the same result for general value distributions.

In contrast, using (1) the myopic rule for v and ε being uniformly distributed gives

$$(1-\delta)v_R^* = \frac{\frac{(1-\delta)\left(1-v_R^{*2}\right)+\delta\left(1-v_R^*\right)}{2} + \frac{1}{2}\delta v_R^* + \frac{\delta^2}{6(1-\delta)}}{1 - v_R^* + \frac{\delta}{2(1-\delta)}} - \frac{s}{1 - v_R^* + \frac{\delta}{2(1-\delta)}},$$

or

$$\begin{aligned} v_R^* &= \frac{-(2-\delta) + \sqrt{(2-\delta)^2 - 4(1-\delta)\left(1 + \frac{\delta^2}{3(1-\delta)} - 2s\right)}}{-2(1-\delta)} \\ &= \frac{(2-\delta) - \sqrt{8(1-\delta)s - \frac{1}{3}\delta^2}}{2(1-\delta)} \end{aligned}$$

Figure 3 demonstrates that the monopolist can learn considerably more than a myopic monopolist. First, if neither learns everything, the platform learns more: the cutoff $v^* < \bar{v}$ is always above v_R^*. And second, the platform learns everything for more values of the parameters: v^* is equal to the maximum possible value $\bar{v}$ for more values of δ than is v_R^*.

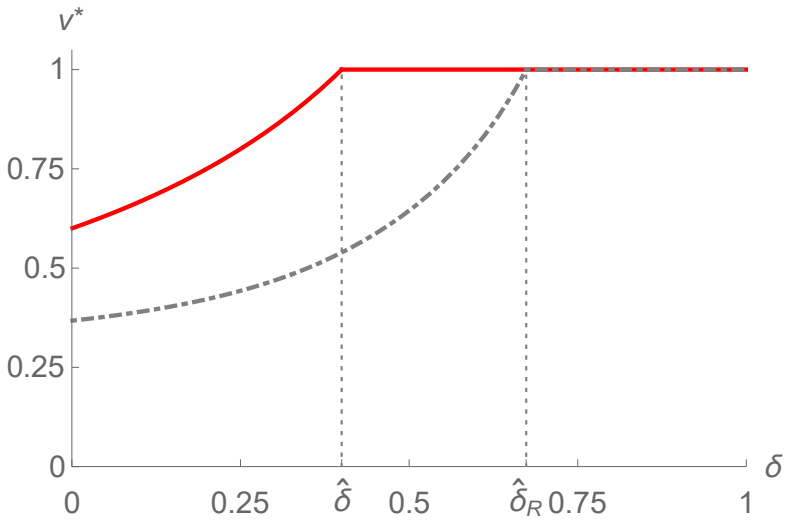


Figure 3: *The monopoly cutoff v^* (red solid) and the myopic cutoff v_R^* (grey dot-dashed) as a function of δ for $s = 0.2$.*

As in the baseline model consumers know their position in the queue, the dynamics is relatively straightforward. If the items are ranked according to the process outlined in the proposition, then consumers have an incentive to search the items in line with the way they are presented. What is important is that the platform experiments with many different items, but that the probability of experimentation eventually vanishes. In the long run, only items with a $v \geq v^*$ are placed in top positions.

A few observations are in place. First, for any $\delta < \hat{\delta}$ the steady state v^* is interior and increasing in δ, the utility function's weight on the idiosyncratic component, and decreasing in s, the search cost. When s is larger, consumers are less inclined to inspect products beyond what is recommended to them, limiting the learning possibilities of the platform. On the other hand, when δ is larger, consumers tend to have more dissimilar tastes and therefore the probability that consumers are satisfied with the product that is recommended to them is smaller. This implies that consumers themselves are more inclined to experiment, making it easier for the platform to discern firms with really high values of v. If the weight δ on the idiosyncratic component of the utility function becomes large enough, or if s is small enough, the platform will always be able to find out the highest value of the common component.

Second, it is interesting to see how consumers can benefit from a platform that serves their interests by comparing for $\delta = 0$ the above steady state value v^* with the standard consumer search problem in the absence of a

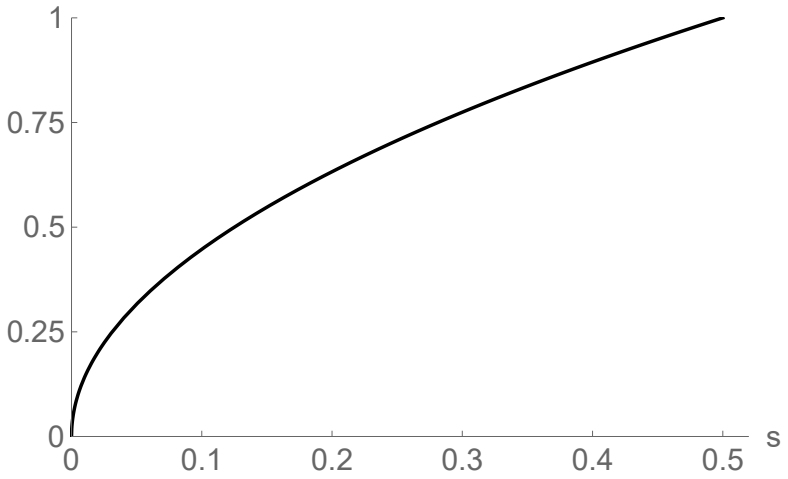


Figure 4: *The value to consumers of having (relative to not having) a platform as a function of s; $\delta = 0$.*

platform where the products to be inspected next are randomly drawn. In the latter case the threshold value $\hat{v}$ that is acceptable for consumers is given by $s = \int_{\hat{v}}^{\bar{v}} (1 - G(v))dv$, which for the uniform distribution results in a value $\hat{v} = 1 - \sqrt{2s}$. Figure 4 below depicts $\frac{(1-2s)-(1-\sqrt{2s})}{1-\sqrt{2s}}$, which is the expected utility of the platform to the consumer relative to the expected utility the consumer gets without the platform. From the Figure it is easy to see that the larger the search cost, the larger the benefit to the consumer, but that even at relatively low inspection costs of say 0.05 consumers benefit around 31.6%.

We will now discuss some generalizations of our monopoly result.

3.1 Extension I: Consumers do not know their position in the queue and not all consumers know the platform's algorithm

First consider that consumers do *not* know their position in the ranking and that not all consumers know the platform's algorithm. We will argue that the steady state of the unfolding dynamic process is unaffacted, but that the dynamic process towards the steady state is more intricate. Not knowing their position in the ranking and possibly not knowing the algorithm, consumers will use their observations to learn and update their beliefs about their position, and possibly the algorithm. Initially, they rationally believe that they are in a steady state and the platform ranking is sufficiently in-

formative to follow the ranking that is provided (until it becomes clear after their inspection that they have been recommended products that in a steady state should not be recommended).

To see that the dynamic path mentioned in Proposition 3 gets the system to the steady state even if consumers do not know their position, consider first that $\delta < \hat{\delta}$ and that all consumers know the algorithm. As at the start of their search, it is infinitely more likely that they have a position later on in the queue, consumers that start searching use the cutoff u^* that satisfies $u^* = (1-\delta)v^* + \delta\underline{\varepsilon}$ for some $v^* < \bar{v}$. Then a consumer (active in period t) that draws a first product and uncovers $u_1 > u^*$ simply buys this product. But if he uncovers $u_1 < u^*$, he will update his beliefs about his position in the queue and puts more weight on the chance that he is early on in the queue because the later he is in the queue, the more likely that the first item has $v_1 > v^*$. The worst that the consumer can believe is that he is the first one in the queue and always has to search over randomly ranked items. Thus, the cutoff $\widehat{u}_t$ that the consumer uses for his purchasing decision is definitely in $[u_R^*, u^*]$, and that $\widehat{v}_t$ defined by $\widehat{u}_t = (1-\delta)\widehat{v}_t + \delta\underline{\varepsilon}$ is in $[v_R^*, v^*]$. Eventually, the consumer buys some item A with $u_A \geq \widehat{u}_t$. With probability $1-\pi_{t+1}$, at $t+1$ the platform ranks item A first if it becomes the item with the highest posterior v and the previously first-ranked item if it doesn't.

If later some consumer rejects A, the platform knows that $v_A < v^*$. In that case another item is ranked first as the non-random item and A is ranked lower. Otherwise, A remains on the first rank for ever (because $v_A > v^*$). In any case, eventually, as shown in the next proposition, the platform will end up with an item B on rank one that has $v_B \geq v^*$ because, first, all consumers initially use cutoff v^* when deciding whether to buy the first-ranked item. Second, $P(v \geq v^*) > 0$ for all $v^* < \bar{v}$. Thus, the probability that none of t consumers buys an item with $v > v^*$ is less than $G(v^*)^t$ which goes to zero as $t \to \infty$.

If $\delta > \hat{\delta}$ even items with $v = \bar{v}$ are rejected with positive probability. Thus, the platform always keeps changing its ranking, learning about the common components of items (including previously uninspected items). Eventually, the platform will be able to rank items with v arbitrarily close to $\bar{v}$ first.

The above discussion assumes that all consumers know the platform's algorithm. However, if only a fraction of consumers know the algorithm, then we can repeat the same argument by only considering those consumers who know the algorithm and who do not update their beliefs about the algorithm the platform uses.

Thus, we have next proposition.

Proposition 4 *Even if consumers do not know their position in the queue and not all consumers know the platform's algorithm, Proposition 3 continues to hold.*

3.2 Extension II: Consumer heterogeneity in s and δ

As a next generalization of the baseline model, consider that consumers have different search cost and that they differ in the weight they attach to the idiosyncratic component. In particular, suppose that the search costs are distributed over an interval $[s_L, s_H]$, with $s_L > 0$, that the relative weight δ is distributed over an interval $[\delta_L, \delta_H]$, with $\delta_L > 0$, that s and δ are independently distributed from each other and that the platform does not know the individual values of these parameters.

This generalization implies that different consumers have different cut-off values v_i^* and a consumer with (s_i, δ_i) buys if

$$(1-\delta_i)v_i^* + \delta_i \underline{\varepsilon} \geq (1-\delta_i)E(v|v \geq v_i^*) + \delta_i E\varepsilon - s_i.$$

we have shown above that if $\delta < \hat{\delta}$ the cut-off value v^* in the baseline model is decreasing in s and increasing in δ. translating that result to the context of hetereogeneous s and δ, it is clear that consumers with lower s and higher δ have larger cut-off values v_i^*, with the highest cut-off value being $v_i^*(s_L, \delta_H)$. As the platform mainly learns when items are rejected, and items with a $v < v_i^*(s_L, \delta_H)$ eventually will be rejected, it follows that we will have an interior steady state if, and only if, $\delta_H \leq \hat{\delta}_L \equiv \frac{s_L}{E(\varepsilon)-\underline{\varepsilon}}$:

Proposition 5 *If consumers differ in their search cost s and their relative weight on the idiosyncratic component δ, then if $\delta_H \leq \hat{\delta}_L$, the steady state is characterized by a ranking r_H^* and a reservation utility u_H^* such that $v_H^* < \bar{v}$ solves*

$$\int_{v_H^*}^{\bar{v}} \frac{1-G(v)}{1-G(v^*)} dv = \frac{s_L - \delta_H(E\varepsilon - \underline{\varepsilon})}{1-\delta_H}.$$

If, on the other hand, $\delta_H > \hat{\delta}_L$, then the steady state is characterized by $v_H^* = \bar{v}$.

3.3 Extension III: Blind buying

Finally, consider that consumers can buy a product without inspecting it (as they may trust the platform's "recommendation") and that the platform cannot observe whether or not the consumer inspected the product, i.e., the platform can only observe whether the the consumer buys the product and

make inferences about the consumer's optimal inspection behaviour. We refer to buying without inspecting as "blind buying". If the platform only observes that the consumer bought a product, but does not know whether he has inspected the product's fit (or believes that he has not), then it does not learn from the consumer choice. Clearly, if consumers have the option of blind buying the platform cannot learn more than what we have considered so far. However, if only a fraction of consumers consider "blind buying", while another fraction always inspects before buying (as we previously assumed all consumers would do), then the previous analysis remains valid. What matters for the steady state analysis is that eventually some consumers inspect some products and these consumers would make the same choices.

In the remaining of this subsection we show that if all consumers consider buying blindly (but they will do so only when this is rational to do), then the platform may still learn, but only for small values of the search cost s. We will show that for small search costs some consumers prefer to inspect, while for large search costs consumers would want to deviate from inspecting in the steady state of Proposition 3. Thus, if all consumers are allowed to buy blindly, the platform still learns more than a myopic platform if the search cost is small. But if the search cost is large it does not learn as much as when consumers cannot buy blindly.

To this end, for every $\delta > 0$ define $\hat{s}_R(\delta)$ as the largest search cost for which $v_R^* = \bar{v}$. For $s < \hat{s}_R(\delta)$, in Figure 1 only the "triangle" region exists. We have that $\lim_{\delta \to 0} \hat{s}_R(\delta) = 0$. We show that if the search cost is below some critical level (which we denote by $\bar{s}(\delta)$), but above $\hat{s}_R(\delta)$,then the platform can learn strictly more than the myopic cutoff v_R^*, but never as much as v^* – the cutoff that the platform reaches when blind buying is not allowed. If that critical value $\bar{s}(\delta) < \hat{s}_R(\delta)$, then

Proposition 6 *For any $\delta > 0$ if $s < \bar{s}(\delta) \leq \hat{s}_R(\delta)$, then the platform learns the steady-state cutoff $\hat{v}$ with $\hat{v} = v_R^* = v^* = \bar{v}$. If $s \in (\hat{s}_R(\delta), \bar{s}(\delta))$ and δ is small enough, then the platform learns a +steady-state cutoff $\hat{v}$ with $v_R^* < \hat{v} \leq v^*$. The platform never learns the cutoff v^* if $v^* < \bar{v}$.*

If the search cost is sufficiently small, inspecting products means that any consumer can guarantee themselves a utility close to $(1-\delta)\bar{v} + \delta\bar{\varepsilon} - s$. Blind buying carries the risk of getting a low idiosyncratic draw of the utility so the expected utility that a consumer gets is at most $(1-\delta)\bar{v} + \delta E\varepsilon$. Thus, for all $\delta > 0$ inspecting is preferred to blind buying if the search cost is small.[8]

[8]We do not explicitly have to derive how the platform ranks items for this argument because the platform is always able to make better recommendations if consumers inspect rather than buy blindly: consumers can expect a higher common component in a steady state if more previous consumers inspect rather than buy blindly.

As a result, the platform learns as much when consumers are able to blind buy as when they are not: it achieves the cutoff $\bar{v}$.

If, on the other hand, $s > \hat{s}_R(\delta)$, then the myopic cutoff value v_R^* is below $\bar{v}$. Consider first that $\delta = 0$ and $s > 0$, but small. The first consumer strictly prefers inspecting to blind buying for small enough search costs (as we have assumed) and uses the cutoff value v_R^*. Later consumers would prefer to buy blindly *if they expected the platform to prominently rank the product that the first consumer bought.* But the platform can achieve a higher steady-state utility by ranking products randomly for the first large group of consumers and then ranking the products bought buy this group highly for a next group of consumers, etc. Since consumers in the second have a higher continuation value than u_R^*, they would prefer to inspect the highly-ranked items rather than settle for $E(u|u \geq u_R^*)$. As a result, even if they can buy blindly some consumers strictly prefer to inspect rather than buy blindly for sufficiently small positive search costs and the platform learns a higher steady-state value than v_R^* if $v_R^* < \bar{v}$. By continuity a similar argument continues to hold if δ is a small positive number.

A simple argument shows that $\hat{v} < v^*$ if $v^* < \bar{v}$, i.e., if the monopoly cutoff is interior when blind buying is not allowed. Recall that v^* is such that $u^* = (1-\delta)v^* + \delta\underline{\varepsilon}$ and all items are accepted in the steady state. The cutoff v^* is the highest possible cutoff that can be reached when blind buying is allowed because the platform can learn weakly less when blind buying is allowed versus prohibited. Suppose that in a steady state the cutoff v^* is supposed to be reached and consider a T that is so large that the steady state is almost reached. Then consumers $t > T$ expect to get a utility (almost) equal to $E(u|u > u^*) - s$ when inspecting, whereas they expect to get (almost) $E(u|u > u^*)$ by buying blindly. As a result, all consumers $t > T$, i.e., sufficiently far in the queue, prefer to buy blindly and the steady state with $v^* < \bar{v}$ cannot be reached and the platform learns less than when consumers cannot buy blindly.

4 Competition

Next, we consider markets where platforms are competing with each other for consumers and ask to what extent the results of Section 3 continue to hold. can platforms learn from observed consumer search and purchase behavior about the underlying common component to the same extent as under monopoly? To answer that question, we start -as in the monopoly section- by letting consumers know their position in the queue.

An important part of the analysis under monopoly platforms is that the

platform can experiment with different rankings to learn consumer choices in a variety of settings. This form of experimentation implies that the utility of early consumers is sacrificed to the benefit of later consumers. This is easiest seen by considering the first and second consumer using the platform. As the platform has no information about products before the first consumer inspects products and makes a choice whether or not to buy, it randomly ranks the products. The situation is, however, different for the second consumer. If the first consumer bought a particular product (whether or not it is the first inspected product), then the platform gets positive information about the common value of that product and *if* it wants to serve the second consumer best, it will prominently place that product in the ranking that is shown to the second consumer so that this consumer inspects this product first. However, doing so limits the scope of learning from the choices the second consumer makes. Thus, a monopoly interested in the long-run benefits of consumers may want to obfuscate the product the first consumer bought and put other products prominently in the ranking for the second consumer.

Under competition, platforms are limited in the extent to which they can experiment with their rankings as consumers will avoid using the platform that experiments most as they know this experimentation is not to their own advantage. Without consumers inspecting different products a platform cannot learn, however. In other words, as consumers only care about their own utility and not about the utility of future consumers they force platforms to create rankings that optimally use the information they have to cater to current consumers.

To study the implications of competition, we start this section by modeling competition between platforms in the following stark way. In some extensions, we show how the results continue to hold in a generalized setting. First, platforms choose how to rank products in every period, i.e., they choose $r = \{r_t\}_{t=1}^{\infty}$. Second, consumers choose which platform to use for their search, which products to inspect, and when to buy. Because of switching costs, the full individual consumer's search process takes place on the platform of the consumer's initial choice. This way of modelling assumes that all consumers know exactly what ranking algorithm the platforms use.[9] For platforms to compete for consumers, each platform maximizes the consumers'

[9]This is, of course, a stark assumption, but it captures in the most simple way one of the implications of the E.U.'s DSA (cf., footnote 4). In the unlikely event that none of the consumers knows anything about the platforms' choice of r, then there would be no competition between platforms and the results of Section 3 apply. Thus, to the extent the DSA makes consumers more aware of the ranking algorithms platforms use, this section may be viewed as an evaluation of the possible effects of the DSA in markets with competing platforms.

average long-run utility that is generated on the platform, i.e., platform k maximizes

$$\lim_{T\to\infty} \sum_{t=1}^{T} \frac{Eu(t)\mathbb{1}_k(t)}{T},$$

where $\mathbb{1}_k(t) = 1$ if consumer t uses platform k, and otherwise it is zero. This objective function captures the idea that platforms can generate more advertising revenue if consumers are more satisfied with it.

The main result in this section is that competition severely limits the extent to which platforms can learn about the common value of consumers' utility and that it cannot do better than under the myopic ranking algorithm introduceed in subsection 2.2.

Proposition 7 *If $\delta < \hat{\delta}$ and consumers know their position in the queue, the steady state of the process under competition is u_R^*, which is strictly smaller than under monopoly, i.e., $u_R^* < u^*$. The steady state is unique.*

The idea behind the proposition is that if $\delta < \hat{\delta}$, a platform will always want to optimally assist *current* consumers. If they do not do so, a competing platform has an incentive and the ability to provide a better ranking for early consumers. Thus, they will use the platform that provides the better ranking for them, and (importantly) they will be expected to do so by later consumers. Optimally assisting *current* consumers implies putting the product in the top spot that previous consumers have inspected and purchased. Other items are either not inspected or inspected and at least once rejected, but from lemma 2 we know that consumers would not want to inspect items in this latter category. The product that is put first can either generate a utility "in the triangle" in Figure 1 (in which case it will ultimately be replaced as some consumers will draw a low idiosyncratic component of the utility function), or a utility "in the rectangle" (in which case it will be purchased by all subsequent consumers after being inspected). In this way, at any point in time there is only one item that has been inspected before and not been rejected and this item is placed as the first item. Thus, under competition the platforms will use the myopic ranking rule introduced in Section 2.

The situation is different and more complicated when $\delta \geq \hat{\delta}$ as consumers may be willing to inspect products that previously have been rejected ahead of random products. However, the extent to which they want to do so depends on the precise value of δ. If δ is marginally larger than $\hat{\delta}$, then consumers only prefer inspecting a previously rejected product to random products if that product has been bought by many consumers before it was eventually rejected by another consumer as this would be an indication that

the common component is very close to v_R^*. The chance that such a product will be discovered is, however, very small as it relies on the fact that products with a common component larger than v_R^* have *not* been inspected before and on the chance that a sufficient number of consumers with a relatively high idiosyncratic component have been presented a ranking where this item has been ranked highly.

We will now discuss the robustness of our competition result to similar generalizations as in the previous section.

4.1 Extension I: Consumers do not know their position in the queue

If consumers do not know their position in the queue, then the analysis is more complicated as consumers will update their beliefs and learn whether they are early searchers or later searchers based on the observed utility of a product. Learning arises when consumers inspect an item that turns out to generate low utility. In that case, consumers should attach a higher probability to them having a position early in the queue. After they learn they may switch platforms, at least if the cost of doing so is not prohibitively high. Due to learning, consumers may face a coordination problem in that if many consumers are expected to choose to use a platform that experiments more, then it may be optimal for an individual consumer to also use the more experimenting platform as it will create more surplus for later consumers if the chance of being a late consumer is sufficiently high.

In many applications, however, consumers may have some idea where they are in the queue, i.e., they may know whether they are early or late adopters, without knowing precisely their position in the queue. One way to model this is to have consumers receive an informative signal of their position in the queue. It is clear that if this signal is informative and precise enough, then the the result in Proposition 7 should continue to hold. To analyze the robustness further, we here consider that different platforms have some initial users (formally, they all have one first consumer visiting them) and that consumers 2 and 3 know that they are early adopters, but do not know whether they are second or third in the queue. Moreover, other consumers know that they are either fourth or fifth in the queue, sixth and seventh, etc, and they know the "market share" of the platform, i.e., they know the fraction of consumers before them choosing a particular platform. The next proposition shows that it continues to be true that competition implies that platforms do not want to experiment.

Proposition 8 *If $\delta = 0$[10] and consumers 2 and 3 in the queue do not know whether they are 2nd or 3rd in the queue, the steady state under competition remains as in Proposition 7.*

The ideas that are exploited in the proof are as follows. First, to attract later consumers it is important for platforms to attract earlier consumers. So, if a platform did not attract consumers 2 and 3, then it will not attract later consumers either. Second, to maximize the chance of attracting consumers 2 and 3, the platform has no incentive to experiment with consumer 3 and therefore if it experiments, it will experiment with consumer 2. Denote by π the probability the platform experiments with consumer 2 (and suppose that in that case it fully randomizes its ranking).[11] Clearly, consumer 2 will be worse off on the experimenting platform, but consumer 3 may be better off. The main question is whether the positive effect on consumer 3 outweighs the negative effect on consumer 2. Third, as the consumer does not know whether he is 2nd or 3rd, he has to choose one reservation value that characterizes optimal search behaviour. Given the observed value in the top spot, he will also update his beliefs about whether he is second or third. For example, if he observes a $v < v_R^*$ then he knows that the platform experimented with him and that implies that he is second in the queue. Fourth, a change in the probability π has a direct effect on the expected utility of the two consumers, but also an indirect effect through the choice of the reservation value. This indirect effect equals zero, however, as – using the envelope theorem – the reservation value is chosen optimally. The proof then shows that the direct effect of π on the utility of the consumer who doesn't know if he is second or third is negative as the negative effect on the second consumer dominates the positive effect on the third consumer.

4.2 Extension II: Consumer heterogeneity in s and δ

In this subsection we show that the result that platform learning is hampered if platforms compete continues to hold when consumers differ in s and δ. As when consumers are homogeneous, competing platforms have an incentive to use algorithms that best serve current consumers in order to attract them. However, what it means to serve consumers best may now differ across consumers. For example, if consumers only differ in their search cost and this difference is large enough, then a high-search-cost consumer may prefer that an item that wa bought by many previous consumers, but eventually rejected

[10] We still plan to consider whether a similar result holds true for $\delta > 0$.

[11] The proof shows that other forms of experimentation are strictly worse for the consumer who does not know whether he is in position 2 or 3.

is ranked highly, while a low-search-cost consumer may prefer not to inspect such an item. Whether or not a platform that faces competition prominently ranks previously rejected items then depends on the distribution of s and δ. To avoid these complications below we present a result that holds if this conflict of interest is not present as the consumer heterogeneity is not that large.

In the argument preceding Lemma 2, we used a hypothetical item with common value $\tilde{v}$ that a consumer was assumed to know and asked when the consumer prefers inspecting a random item with some reservation utility over picking the item with a known $\tilde{v}$. Similarly, suppose here that consumer i of type (δ_i, s_i) "knows" that an item has the common component equal to $\tilde{v}$. Recall that consumer i's value of searching over randomly ordered uninspected items is $u_R(\delta_i, s_i) = (1-\delta_i)v_R(\delta_i, s_i) + \delta_i\underline{\varepsilon}$, where the cutoff $v_R(\delta_i, s_i)$ solves

$$v_R(\delta_i, s_i) = E(u|u \geq u_R(\delta_i, s_i)) - \frac{\delta_i\underline{\varepsilon}}{1-\delta_i} - \frac{s_i}{(1-\delta_i)\Pr(u \geq u_R(\delta_i, s_i))}.$$

Thus, she prefers to search through such items rather than pick the item with known $v = \tilde{v}$ iff

$$(1-\delta_i)v_R(\delta_i, s_i) + \delta_i\underline{\varepsilon} > (1-\delta_i)\tilde{v} + \delta_i E\varepsilon - s_i,$$

which is harder to satisfy the higher is $\tilde{v}$. As $v_R(\delta_i, s_i)$ increases in δ_i and decreases in s_i, the maximum common value $\tilde{v}$ an item can have that is previously rejected is $v_R(\delta_H, s_L)$. Thus, if

$$s_i > (1-\delta_i)(v_R(\delta_H, s_L) - v_R(\delta_i, s_i)) + \delta_i(E\varepsilon - \underline{\varepsilon}) \tag{4}$$

holds true for all consumers i, then they all prefer inspecting randomly drawn items to items that have been previously rejected.

Proposition 9 *If consumers differ in s and δ, then the steady state with competing platforms is characterized by a reservation utility $u_R(\delta_H, s_L)$ with $v_R(\delta_H, s_L) < \bar{v}$ if (4) holds for all consumer types (δ_i, s_i) and $v_R(\delta_H, s_L) < v_H^*$. If (4) does not hold for all (δ_i, s_i), then the reservation utility in the steady state is smaller than or equal to $u_R(\delta_H, s_L)$ if $v_R(\delta_H, s_L) < \bar{v}$.*

Proposition 9 shows that also under consumer heterogeneity competing platforms learn less than a monopoly platform as long as $v_R(\delta_H, s_L) < \bar{v}$. Without consumer heterogeneity, i.e., when all consumers have the same δ_H and s_L, condition (4) boils down to $\delta_H(E\varepsilon - \underline{\varepsilon}) < s_L$, or $\delta_H < \hat{\delta}_L$. Thus, it

is clear that under consumer heterogeneity there is an open set of parameter values such that Proposition 9 applies.[12]

If (4) does not hold for all (δ_i, s_i), then depending on the distributions of δ and s it could be that the platforms feel forced to continue prominently ranking previously rejected items if their common component is reasonably good. They could do so if they calculate that a large majority of consumers will buy the product and not continue to search for better fits. If this is the case the platforms may learn even less than $v_R(\delta_H, s_L)$ as previously rejected items surely have a common value $v < v_R(\delta_H, s_L)$. Together with $v_R(\delta_H, s_L) < \bar{v}$ this would imply that the platforms continue to learn less than under monopoly.

4.3 Extension III: Blind buying

Next, we consider how the results under competition are affected if consumers can buy blindly. It is not difficult to see that the issues here are similar to those we encountered in the monopoly section on blind buying. First, if s is so small that u_R^* is close to $(1-\delta)\overline{v} + \delta\bar{\varepsilon}$, then consumers will always want to inspect even if they can buy blindly and the main analysis in this competition section is not affected. Second, if $v_R^* < \overline{v}$, then the platforms can learn even less than without blind buying if δ and s are small enough. To see this note that the first consumer wants to inspect products if the search cost is relatively small. It is easy to see that if $\delta = 0$ all subsequent consumers are strictly better off to buy blindly as the first ranked product will always be bought: the search cost is simply a waste. By continuity, consumers also would prefer to buy blindly if δ is small enough. This means that platforms will only learn that the product has a $v \geq \frac{u_R^* - \delta\bar{\varepsilon}}{1-\delta}$, which is smaller than $v_R^* = \frac{u_R^* - \delta\underline{\varepsilon}}{1-\delta}$.

4.4 Extension IV: Incumbent and entrant

We finally explore a different form of competition, namely between an incumbent and a more efficient entrant and ask the question how efficient the entrant should be to be able to attract consumers to its network. We measure efficiency by how easy it ease to navigate a platform's website and for consumers to find out information about products. Thus, the incumbent's efficiency is measured by the consumers' search cost s_I on its platform and the

[12]For example, when v and ε are uniform on $[0,1]$, $\delta_H = 0.1$ and $s_L = 0.06$, then the condition is satisfied if $\delta_L \geq 0$ and $s_H \leq 0.066$.

entrant's efficiency is s_E. As the incumbent has already learned about consumer preferences, while the entrant has not, we consider the only interesting case where $s_E < s_I$.

Consider that the incumbent has been in the market long enough so that a steady state is reached and the incumbent ranks items with $v \geq v^*(s_I)$ first. Then a consumer's expected utility of searching on the incumbent platform is $(1-\delta)E(v|v \geq v^*(s_I)) + \delta E\varepsilon - s_I = (1-\delta)v^*(s_I) + \delta\underline{\varepsilon}$. As at the moment of entry, the entrant has no consumers whose choices to learn from, the first consumer on that platform has an expected utility of $u_R^*(s_E) = (1-\delta)v_R^*(s_E) + \delta\underline{\varepsilon}$. The consumer, thus, prefers to use the entrant's platform iff $v_R^*(s_E) > v^*(s_I)$. Note that this also implies that if $v^*(s_I) = \bar{v}$, the entrant cannot attract any consumers. The formal result below then follows directly from Proposition 3 and the facts that an interior monopoly cutoff v^* satisfies $v^*(s) > v_R^*(s)$ and both cutoffs decrease in s.

Corollary 10 *For any $s_I > s_E > 0$ there exists $\Delta(s_I, s_E) > 0$ such that the entrant can enter the market only if $\frac{s_E}{s_I} < \Delta(s_I, s_E)$.*

For general distribution functions F and G, it is difficult to determine exactly how large the efficiency gains of the entrant should be for it to be able to enter the market. However, for specific distributions, it is easy to see that the efficiency gains may need to be very large for entry to take place. To see this, define $\frac{s_E}{s_I}$ as the measure of inefficiency of the entrant relative to the incumbent: if the entrant is maximally efficient (i.e., $s_E = 0$), the measure is zero and if the entrant has no efficiency advantage, the measure is one. Then, if v and ε are uniformly distributed between zero and one, and considering $u_R^*(s_E)$ and $v^*(s_I)$ to be interior, the expected utility is higher on the entrant platform if

$$v_R^*(s_E) = \frac{(2-\delta) - \sqrt{8(1-\delta)s_E - \frac{1}{3}\delta^2}}{2(1-\delta)} > 1 + \frac{\delta - 2s_I}{1-\delta} = v^*(s_I),$$

or

$$\frac{s_E}{s_I} < \frac{2s_I + \frac{\delta^2}{6s_I} - \delta}{1-\delta}.$$

For $\delta = 0$, this inequality reduces to $\frac{s_E}{s_I} < 2s_I$. For a relatively inefficient incumbent with an s_I being equal to 0.1 this requires the entrant to be 5 times more efficient. For a more efficient incumbent, the entrant needs to be increasingly more efficient. Figure 5 illustrates how much more efficient the entrant has to be than the incumbent for different values of δ. The larger δ the more efficient the etrant has to be.

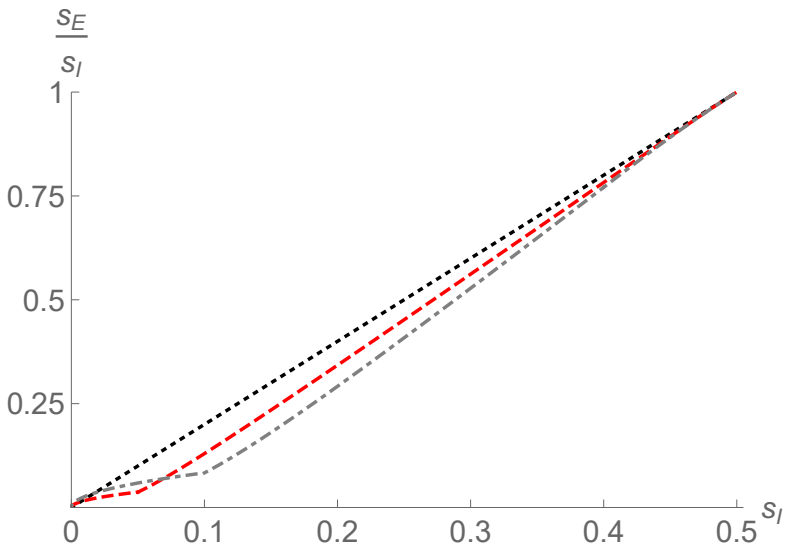


Figure 5: *The efficiency bound on the entrant's relative search cost* $\frac{s_E}{s_I}$ *for* $\delta = 0$ *(black dotted),* 0.1 *(red dashed), and* 0.2 *(grey dot-dashed).*

5 Discussion and Conclusion

In this paper we have asked to what extent online platforms may infer information about a common component in how consumers value products on the basis of the observations regarding consumer search behavior, i.e., the products they have inspected and (not) bought. If a monopoly platform cares for long-run consumer surplus, they will experiment with product rankings and eventually only rank products that consumers have bought in the past in top spots. This will benefit consumers as they will only accept to buy products that are much better than under random search. Competition, and selling top spots to firms to boost their sales, seriously limits the incentives of platforms to learn.

In order to focus on what platforms may learn from consumer search behaviour, we have kept firm behaviour as exogenous throughout the paper. An interesting next step in the analysis is to endogenize firm behaviour, for example, by endogenizing their pricing decisions or by allowing them to also have some incomplete information about consumers' common value that they could signal in their bids to acquire the top position. Another interesting avenue for future research is to allow products to be vertically differentiated and to have consumers being different in their value for quality. Will this gives rise to echo chambers where some platforms cater to consumers with a high willingness to pay for quality and other platforms that do not?

6 Appendix: proofs

Proof of Lemma 2.

According to the optimal search rule from Weitzman (1979), a consumer should search through items in descending order according to their reservation values. We show that, if $\delta < \hat{\delta}$, the reservation value of any item that has been inspected but not bought is below the reservation value of a random item from a certain subset of products. That is, if $\delta < \hat{\delta}$, a consumer optimally searches random items from this subset before searching any item that has been inspected but not bought before.

Consider the consumer has the option of inspecting a product with a reservation value $\tilde{u} \geq u_R^*$ and define $\hat{v} = \frac{\tilde{u}-\delta\underline{\varepsilon}}{1-\delta}$. Denote by $\hat{\rho}(v)$ the reservation value of inspecting a product with a common component of the product equal to $v \leq \hat{v}$. We show that if $\delta < \hat{\delta}$ and $v \leq \hat{v}$, then $\hat{\rho}(v) \leq \tilde{u}$.

If a solution exists, $\hat{\rho}(v)$ satisfies

$$E[\max\{u, \hat{\rho}(v)\}] - \hat{\rho}(v) - s = 0, \tag{5}$$

for $u = (1-\delta)v + \delta\varepsilon$.

Consider first that $\hat{\rho}(v) \leq (1-\delta)v + \delta\underline{\varepsilon}$. In this case $E[\max\{u, \hat{\rho}(v)\}] = (1-\delta)v + \delta E\varepsilon$ so that

$$\hat{\rho}(v) = (1-\delta)v + \delta E\varepsilon - s.$$

It is clear that this is a possible solution that satisfies the constraint $\hat{\rho} \leq (1-\delta)v + \delta\underline{\varepsilon}$ if $\delta E\varepsilon - s \leq \delta\underline{\varepsilon}$ or $\delta \leq \hat{\delta}$.

Consider then that $\hat{\rho}(v) > (1-\delta)v + \delta\underline{\varepsilon}$. In this case, we can define $\varepsilon(v) := \frac{\hat{\rho}(v)-(1-\delta)v}{\delta} > \underline{\varepsilon}$ such that

$$\delta \int_{\varepsilon(v)}^{\bar{\varepsilon}} (\varepsilon - \varepsilon(v)) f(\varepsilon) d\varepsilon = s.$$

The reservation value $\hat{\rho}(v)$ is then implicitly defined through $\hat{\varepsilon}$. Since the LHS of this expression is decreasing in $\varepsilon(v)$, the LHS is smaller than $\delta \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} (\varepsilon - \underline{\varepsilon}) f(\varepsilon) d\varepsilon = \delta(E\varepsilon - \underline{\varepsilon})$. Therefore, if $\delta \leq \hat{\delta}$ there is no solution $\hat{\rho}(v) > (1-\delta)v + \delta\underline{\varepsilon}$.

Thus, if $\delta \leq \hat{\delta}$, then for any $v \leq \hat{v}$ the reservation value is given by $\hat{\rho}(v) = (1-\delta)v + \delta E\varepsilon - s$, which is smaller than $(1-\delta)v + \delta\underline{\varepsilon}$ which for $v \leq \hat{v}$, in turn, is smaller than $(1-\delta)\hat{v} + \delta\underline{\varepsilon} = \tilde{u}$, the reservation value of the alternative option.

Proof of Proposition 3.

We first prove the existence and uniqueness of an interior steady state (i.e., with $v^* < \bar{v}$ where v^* is the threshold value of the common component of the utility in a steady state) and then the corner solution (i.e., with $v^* = \bar{v}$).

In an interior steady state, the consumer should always buy an item with $u \geq u^* = (1-\delta)v^* + \delta\underline{\varepsilon}$. A consumer that draws the minimum possible utility level in a steady state, $(1-\delta)v^* + \delta\underline{\varepsilon}$, buys if

$$(1-\delta)v^* + \delta\underline{\varepsilon} \geq (1-\delta)E(v|v \geq v^*) + \delta E\varepsilon - s,$$

which gives

$$v^* \geq E(v|v \geq v^*) + \frac{\delta(E\varepsilon - \underline{\varepsilon}) - s}{1-\delta}. \tag{6}$$

Clearly inequality (6) can hold only if $\delta(E\varepsilon - \underline{\varepsilon}) < s$ so consider such δ and s. We can write (6) as

$$\begin{aligned} v^* &\geq \frac{\int_{v^*}^{\bar{v}} vg(v)dv}{1-G(v^*)} + \frac{\delta(E\varepsilon - \underline{\varepsilon}) - s}{1-\delta} \\ &= \frac{\bar{v} - v^*G(v^*) - \int_{v^*}^{\bar{v}} G(v)dv}{1-G(v^*)} + \frac{\delta(E\varepsilon - \underline{\varepsilon}) - s}{1-\delta} \\ &= v^* + \frac{\bar{v} - v^* - \int_{v^*}^{\bar{v}} G(v)dv}{1-G(v^*)} + \frac{\delta(E\varepsilon - \underline{\varepsilon}) - s}{1-\delta}. \end{aligned}$$

Thus, we have that an interior v^* is implicitly defined by

$$v^* = \bar{v} - \int_{v^*}^{\bar{v}} G(v)dv + (1-G(v^*))\frac{\delta(E\varepsilon - \underline{\varepsilon}) - s}{1-\delta}. \tag{7}$$

Equation (7) can be rewritten as

$$\int_{v^*}^{\bar{v}} \frac{1-G(v)}{1-G(v^*)}dv = \frac{s - \delta(E\varepsilon - \underline{\varepsilon})}{1-\delta}. \tag{8}$$

Since the RHS of equation (8) is positive, we also need the LHS to be positive for an interior v^* to exist. Now

$$\lim_{v^* \to \underline{v}} \int_{v^*}^{\bar{v}} \frac{1-G(v)}{1-G(v^*)}dv = \int_{\underline{v}}^{\bar{v}} 1 - G(v)dv = Ev - \underline{v} > 0,$$

and

$$\lim_{v^* \to \bar{v}} \int_{v^*}^{\bar{v}} \frac{1-G(v)}{1-G(v^*)}dv = \lim_{v^* \to \bar{v}} \frac{-(1-G(v^*))}{-g(v^*)} = \frac{0}{g(\bar{v})} = 0,$$

as $g(\bar{v}) > 0$. Thus, a sufficient condition for an interior solution for v^* to exist is $Ev - \underline{v} > \frac{s-\delta(E\varepsilon-\underline{\varepsilon})}{1-\delta}$ or

$$(1-\delta)(Ev-\underline{v}) + \delta(E\varepsilon - \underline{\varepsilon}) > s.$$

This inequality holds because, by assumption, both $Ev - \underline{v} > s$ and $E\varepsilon - \underline{\varepsilon} > s$.

A sufficient condition for a unique interior v^* to exist is that the derivative of the LHS of equation (8), or

$$\frac{\partial}{\partial v^*}\int_{v^*}^{\bar{v}} \frac{1-G(v)}{1-G(v^*)}dv = -1 + \int_{v^*}^{\bar{v}} \frac{(1-G(v))g(v^*)}{(1-G(v^*))^2}dv, \tag{9}$$

is negative for all $v^* \in (\underline{v}, \bar{v})$. The derivative can be rearranged so that it is negative if for all v^*

$$g(v^*)\int_{v^*}^{\bar{v}} \frac{1-G(v)}{1-G(v^*)}dv < 1 - G(v^*). \tag{10}$$

We show that inequality (10) holds because the RHS of the inequality decreases faster than the LHS for all $v^* < \bar{v}$, while the two sides equal at $v^* = \bar{v}$. At $v^* = \bar{v}$, using l'Hopital's rule, the LHS equals $g(\bar{v})\frac{-(1-G(\bar{v}))}{-g(\bar{v})} = g(\bar{v})\cdot 0 = 0$ and the RHS is also equal to zero. Taking the derivative with respect to v^* on both sides of inequality (10) gives $-g(v^*)\left[1 - \int_{v^*}^{\bar{v}} \frac{(1-G(v))g(v^*)}{(1-G(v^*))^2}dv\right] + g'(v^*)\int_{v^*}^{\bar{v}} \frac{(1-G(v))}{1-G(v^*)}dv$ for the LHS and $-g(v^*)$ for the RHS. Thus, the derivative of the LHS of (10) is larger than the derivative of the RHS if

$$\left[\frac{g^2(v^*)}{1-G(v^*)} + g'(v^*)\right]\int_{v^*}^{\bar{v}} \frac{(1-G(v))}{1-G(v^*)}dv > 0.$$

This is indeed the case if $\frac{g^2(v^*)}{1-G(v^*)} + g'(v^*)$ which follows from the fact that $1-G(v)$ is logconcave: logconcavity implies that for the first and second derivative of $\ln(1-G(v))$ we have $\frac{-g}{1-G}$ and $\frac{-g'(1-G)-g^2}{(1-G)^2} = \frac{-1}{(1-G)}\left[g' + \frac{g^2}{(1-G)}\right] \leq 0$. In sum, an interior steady state with $v^* < \bar{v}$ exists and is unique if $\delta(E\varepsilon - \underline{\varepsilon}) < s$.

A final claim to prove about the interior steady state is that v^* strictly increases in δ. We have just shown that the LHS of equation (8) decreases in v^*. It is easy to show that the RHS decreases in δ. Since the LHS of (8) does not depend directly on δ and the RHS on v^*, these two facts establish that $\frac{\partial v^*}{\partial \delta} > 0$.

Now consider $\delta(E\varepsilon - \underline{\varepsilon}) > s$. For any $G(v^*) < 1$ the RHS of equation (7) is increasing in δ without bound for $\delta(E\varepsilon - \underline{\varepsilon}) > s$, so it must be that

$\lim_{\delta \to 1} v^* \to \bar{v}$. Finally, we show that both an interior solution to v^* and $v^* = \bar{v}$ cannot coexist. For an interior solution, we need $s > \delta(E\varepsilon - \underline{\varepsilon})$. For the corner solution with $v^* = \bar{v}$ we need that a consumer who draws the lowest idiosyncratic utility component continues even if he draws the highest common utility component:

$$(1-\delta)\bar{v} + \delta\underline{\varepsilon} < (1-\delta)\bar{v} + \delta E\varepsilon - s,$$

or $s < \delta(E\varepsilon - \underline{\varepsilon})$, which cannot hold together with condition $s > \delta(E\varepsilon - \underline{\varepsilon})$. In sum, a steady state with a corner solution exists and is unique if $s < \delta(E\varepsilon - \underline{\varepsilon})$.

We now turn to the dynamics leading to the steady state. In particular, we show that for each t, $\pi_{t+1} < \pi_t$ exists such that the cutoffs that the consumers' use to decide whether to accept the highest-ranked item, $\hat{u}_{t1}$, satisfy $\hat{u}_{t+1,1} > \hat{u}_{t1}$. Thus, if $\hat{v}_{t1}$ satisfies $(1-\delta)\hat{v}_{t1} + \delta\underline{\varepsilon}$, $\{\hat{v}_{t1}\}_{t=1}^{\infty}$, converges to v^* as $t \to \infty$ if $v^* < \bar{v}$. The highest cutoff that consumer t uses is the one she uses to evaluate whether to accept the item ranked highest for her because the consumer optimally inspects items in a decreasing order of their reservation values. We focus on showing that this cutoff increases in t since the platform learns the most from rejections.

First suppose that $v_R^* < \bar{v}$. We show first that clearly the platform can choose $\pi_3 < \pi_2 < 1$ so that $\hat{u}_{31} > \hat{u}_{21}$. Consider consumer $t = 1$. She uses the cutoff $\hat{u}_{11} = u_R^*$ to evaluate the highest-ranked item because she knows that she is always searching over randomly ordered previously uninspected items.

Consider consumer $t = 2$. Since consumer $t = 2$ compares her utility of accepting the highest-ranked item to her continuation value, she, too, uses the cutoff utility $\hat{u}_{21} = u_R^*$ because her second-highest-ranked item is definitely a randomly chosen previously uninspected item.

Consider consumer $t = 3$. For this consumer, it matters whether the platform experimented with consumer $t = 2$. If it didn't, then consumers $t = 1$ and $t = 2$ bought the same item so the second-highest ranked item is a previously uninspected item. But if the platform experimented with consumer $t = 2$, the second-highest-ranked item of $t = 3$ has an expected common components equal to $(1-\pi_3)E(v|v \geq v_R^*) + \pi_3 E(v)$. All items on lower ranks are randomly drawn uninspected items with probability one. Thus, to evaluate whether to accept the highest-ranked item, consumer $t = 3$ uses a cutoff utility $\hat{u}_{31}$ that satisfies

$$\hat{u}_{31} = \pi_2(1-\pi_3)[P(u \geq \hat{u}_{31}|v \geq v_R^*)E(u|v \geq v_R^*, u \geq \hat{u}_{31}) + P(u < \hat{u}_{31}|v \geq v_R^*)u_R^*]$$

$$+[1-\pi_2(1-\pi_3)][P(u \geq \hat{u}_{31})E(u|u \geq \hat{u}_{31}) + P(u < \hat{u}_{31})u_R^*] - s,$$

because the second-highest-ranked item is a nonrandom item if the platform experimented with consumer $t = 2$ and doesn't experiment with the second-spot item of consumer $t = 3$. The cutoff $\hat{u}_{31}$ clearly exceeds $u_R^* = \hat{u}_{21}$.

Thus, we have $\hat{u}_{31} > \hat{u}_{21}$. Now suppose, by way of contradiction, that buyers know the platform's algorithm, but even for $\pi_{t+1} = \pi_t$ the cutoffs are ordered $\hat{u}_{t+1,1} < \hat{u}_{t1}$ for $t \geq 3$. The crucial thing to figure out is how the information that the platform has before setting a ranking for a consumer τ changes from $\tau = t$ to $\tau = t + 1$. Clearly, the decisions of consumer $\tau = 3$ generate information for the platform because $\hat{u}_{31} > \hat{u}_{21}$.[13] Thus, the platform is able to rank the items according to their expected common values in a manner that better reflects their true common values and the consumers' preferences for consumer $t = 4$ than $t = 3$.[14] As a result, if the platform experiments with consumer $t = 4$ with the same probability as with consumer $t = 3$ (i.e., $\pi_4 = \pi_3$), $\hat{u}_{41} > \hat{u}_{31}$. Since the inequality is strict, there also exists a π_4 strictly smaller than but close enough to π_3 such that the inequality $\hat{u}_{41} > \hat{u}_{31}$ continues to hold. A similar argument holds for all subsequent buyers so that $\{\hat{u}_{t1}\}_{t=2}^{\infty}$ constitutes a monotonically increasing sequence that is bounded above and thus converges to a limit. We showed above that the limit can only be v^*.

Finally, suppose that $v_R^* = \bar{v}$. Now even consumer $t = 1$ continues searching with a positive probability even if she inspects an item with $v = \bar{v}$. However, each consumer's decisions are again informative for the platform about the common components of items: if one item is bought by the same number of people as another, but rejected by fewer, it is likely to have a higher common component. As a result, the platform can make better recommendations for consumers as t increases and the cutoff $\hat{u}_{t1}$ that consumer t uses to evaluate the highest-ranked item increases in t. Eventually, the platform is able to rank highest an item with a common value component arbitrarily close to $\bar{v}$.

[13] More specifically, if $\tau = 3$ buys the highest-ranked item, the platform updates its belief about the common value of that item (and about no other item). If the buyer buys another item, the platform learns that the highest-ranked item has $v < \hat{v}_{31}$, other rejected items have common values less than the appropriate cutoffs, and if the buyer finally accepts item on rank k, that this item has $v \geq \hat{v}_{3k}$.

[14] Specifically, the item ranked highest for $t = 4$ has a strictly higher expected common value as opposed to consumer 3 and on lower spots, either the items there were not inspected by $t = 3$ (and thus their expected common value doesn't change) or were inspected and rejected by $t = 3$ (and thus consumers $t > 3$ prefer to inspect previously uninspected items ahead of these by Lemma 2).

Proof of Proposition 4

As the main complication introduced by consumers not knowing their position in the queue is the learning regarding this position, the proof focuses on the dynamics towards the steady state. In the proof we focus on the case where all consumers know the platform's algorithm. If only a fraction of consumers know the algorithm, then we can repeat the same argument by only considering those consumers who know the algorithm and who do not update their beliefs about the algorithm the platform uses. As there is an infinite number of these consumers, the argument below continues to hold.

Since consumers do not know their positions, they must form beliefs about them. When a consumer starts searching, he believes that he is in the steady state almost surely so a priori uses the cutoff u^*. If the consumer's utility from the first item exceeds u^*, he simply buys it and does not care about what his position in the queue really was.

Conversely, a consumer that gets utility $u_1 < u^*$ for the first item that he inspects, may update his beliefs about his position, depending on the value of v^*. In particular, if $v^* = \bar{v}$, the consumer still believes that he is in the steady state with probability one because even in the steady state consumers reject all items (including those with $v = \bar{v}$) with strictly positive probability. Thus, in this case the consumer simply continues using the steady-state cutoff u^*.

But if $v^* < \bar{v}$, the consumer who gets utility $u_1 < u^*$ for the first item that he inspects, knows that he is not in the steady state and puts almost all of the probability weight to being early in the queue, i.e., at positions $t < T$ for some finite $T > 1$. His posterior belief that he is the first in the queue becomes

$$P(t=1|u_1<u^*) = \frac{P(u_1<u^*|t=1)P(t=1)}{P(u_1<u^*)}$$

$$= \frac{P(u_1<u^*|t=1)P(t=1)}{P(u_1<u^*|t=1)P(t=1)+...+P(u_1<u^*|t=T)P(t=T)}, \qquad (11)$$

where the denominator satisfies

$$P(u_1<u^*|t=1)P(t=1)+...+P(u_1<u^*|t=T)P(t=T)$$

$$\geq \pi_1 P(u<u^*)P(t=1)+\pi_2 P(u<u^*)P(t=2)+...+\pi_T P(u<u^*)P(t=T),$$

because an unsampled randomly drawn item has $u < u^*$ with strictly positive probability $P(u < u^*)$ at any t. The denominator in (11) is strictly higher than the RHS of the last inequality if some consumers use cutoffs that are strictly below u^*. But the important thing is that the denominator in (11) exceeds $P(u_1 < u^*|t = 1)P(t = 1)$. This holds because after observing

$u_1 < u^*$, the consumer puts strictly positive weight on the events that he is t'th in the queue for $t < T < \infty$ and $\pi_2, ..., \pi_T > 0$. Thus, the posterior in (11) is strictly below one: a consumer that gets utility $u_1 < u^*$ for the first item that he inspects believes that he is early on in the queue, but not that he is the first consumer with probability one. As a result, the consumer uses a cutoff strictly larger than u_R^* to decide whether to buy the first item he inspects.

Let the cutoff to decide whether to purchase the first inspected item with $u_1 < u^*$ (i.e., used by a consumer who updates his beliefs about his position) be denoted $\hat{u}_t = (1-\delta)\hat{v}_t + \delta\underline{\varepsilon}$ where t stands for the true position of the consumer (which is unknown to the consumer himself). Note that in principle, $\hat{v}_t$ could depend on u_1. We want to show that $\lim_{t\to\infty} \hat{v}_t = v^*$.

Suppose first that $\hat{v}_t = v^*$ for some finite t and consumer t buys item A. Then the item with the highest posterior v is A and A is always ranked first as the nonrandom item. Since $\pi_t \to 0$, eventually A is ranked first with probability one and, as consumers use the cutoff v^* to decide whether to buy it, the steady state is reached.

Now suppose, conversely, that $\hat{v}_t < v^*$ for all t. Let $\max_t \hat{v}_t =: \dot{v}$. Note that no consumer rejects a first item with $u_1 \geq \dot{v}$: after seeing $u_1 \in [\dot{v}, v^*)$ the consumer updates his belief about his position and then accepts the item, while after seeing $u_1 \geq v^*$ the consumer accepts it without updating his belief. Effectively, $\dot{v}$ then is the steady-state cutoff because no item with $v \geq \dot{v}$ is rejected. But then observing an item with $u \in [\dot{u}, u^*)$ cannot depress a consumer's belief about his position so he should use the original steady-state cutoff v^*, not $\hat{v}_t \leq \dot{v}$. Thus, we must have that $\dot{v} = v^*$ which means that either the argument in the previous paragraph holds or $\lim_{t\to\infty} \hat{v}_t = v^*$.

Proof of Proposition 6.

First let $\delta > 0$ and consider $s < \hat{s}_R(\delta)$. Since $v_R^* = \bar{v}$, the expected utility of searching over randomly ordered items is $u_R^* > (1-\delta)\bar{v} + \delta\underline{\varepsilon}$. For all $\delta > 0$, let $\varepsilon_R^*(s)$ satisfy $(1-\delta)\bar{v} + \delta\varepsilon_R^*(s) = u_R^*$. Note that $\varepsilon_R^*(s=0) = \bar{\varepsilon}$, $\varepsilon_R^*(s = \hat{s}_R(\delta)) = \underline{\varepsilon}$, and $\varepsilon_R^*(s)$ strictly decreases in s.

Consider any consumer $t \geq 1$. The highest expected utility that a consumer can potentially achieve by buying blindly is $(1-\delta)\bar{v} + \delta E(\varepsilon)$. Conversely, by inspecting consumers can always ensure a utility of u_R^* because they can always ignore the ranking of the platform and sample items in a random order. Inspecting is, thus, certainly better than buying blindly for all consumers $t \geq 1$ if

$$(1-\delta)\bar{v} + \delta E(\varepsilon) < (1-\delta)\bar{v} + \delta\varepsilon_R^* - s,$$

or

$$\delta E(\varepsilon) + s < \delta \varepsilon_R^*.$$

At $s = 0$, the inequality clearly holds because the LHS becomes $\delta E\varepsilon$ and the RHS becomes $\delta\bar{\varepsilon}$. At $s = \hat{s}_R(\delta)$, the inequality becomes $\delta E(\varepsilon) + \hat{s}_R(\delta) < \delta\underline{\varepsilon}$ which clearly does not hold. Since ε_R^* decreases in s, a unique $\bar{s}(\delta)$ exists such that $\delta E(\varepsilon) + \bar{s}(\delta) = \delta\varepsilon_R^*(\bar{s}(\delta))$ and the inequality holds for all $s < \bar{s}(\delta)$. Thus, for any $s < \bar{s}(\delta) < \hat{s}_R(\delta)$, all consumers will inspect and the platform is able to learn $\bar{v}$.

Consider next that $s > \hat{s}_R(\delta)$. We first focus on $\delta = 0$ (so that $\hat{s}_R(\delta) = 0$) and consider a small $s > 0$. Now the first consumer prefers to inspect (using the cutoff $v_R^* < \bar{v}$). For later consumers the platform can rank products randomly for the first T_1 consumers and then rank first the products bought by these first T_1 consumers for the next T_2 consumers. Consumers t with $T_1 < t \leq T_1 + T_2$ have a continuation value from inspecting $\tilde{v}$ that satisfies

$$\tilde{v} \geq \frac{1 - G(\tilde{v})}{1 - G(v_R^*)} E(v|v \geq \tilde{v}) + \frac{G(\tilde{v}) - G(v_R^*)}{1 - G(v_R^*)} \tilde{v} - s, \tag{12}$$

where an interior $\tilde{v}$ solves the equation with equality. To prove that for sufficiently low search costs consumers t with $T_1 < t \leq T_1 + T_2$ would prefer to inspect the highly-ranked items to find an item with $v \geq \tilde{v}$ rather than blind buy and settle for $E(v|v \geq v_R^*)$, we need to prove that for sufficiently large T_1 and for small s

$$E(v|v \geq \tilde{v}) - \frac{(1 - G(v_R^*))\, s}{1 - G(\tilde{v})} > E(v|v \geq v_R^*),$$

or

$$(1 - G(\tilde{v}))\, E(v|v \geq \tilde{v}) - (1 - G(v_R^*))\, s > (1 - G(\tilde{v}))\, E(v|v \geq v_R^*).$$

Taking the total differential, we need to prove that in a neighborhood of $s = 0$

$$-g(\tilde{v})\tilde{v}d\tilde{v} - (1 - G(v_R^*))\, ds + g(v_R^*)sdv_R^* > (1 - G(\tilde{v}))\, \frac{g(v_R^*)s}{(1 - G(v_R^*))^2} dv_R^* - g(\tilde{v})d\tilde{v}E(v|v \geq v_R^*),$$

where we use the fact $(1 - G(\tilde{v}))\, E(v|v \geq \tilde{v}) = \int_{\tilde{v}} g(v)vdv$. As $\int_{\tilde{v}} g(v)(v - \tilde{v})dv = (1 - G(v_R^*))\, s$ gives in a first-order approximation that $-(1 - G(\tilde{v}))\, d\tilde{v} = (1 - G(v_R^*))\, ds - g(v_R^*)sdv_R^*$ we have

$$-g(\tilde{v})\tilde{v}d\tilde{v} + (1 - G(\tilde{v}))\, d\tilde{v} > (1 - G(\tilde{v}))\, \frac{g(v_R^*)s}{(1 - G(v_R^*))^2} dv_R^* - g(\tilde{v})d\tilde{v}E(v|v \geq v_R^*).$$

$$\left[1+\frac{g(\tilde{v})}{(1-g(\tilde{v}))}\left(E(v|v\geq v_R^*)-\tilde{v}\right)\right]d\tilde{v} > \frac{g(v_R^*)s}{(1-G(v_R^*))^2}dv_R^*.$$

To show that this holds, it is sufficient to show that

$$\lim_{s\to 0}\frac{E(v|v\geq v_R^*)-\tilde{v}}{(1-G(\tilde{v}))}=-\infty$$

as $d\tilde{v}/ds$ and dv_R^*/ds are negative and $\frac{g(v_R^*)s}{\left(1-G(v_R^*)\right)^2}$ is positive. So, the LHS would be positive and the RHS would be negative.

$$\begin{aligned}
&\lim_{s\to 0}\frac{E(v|v\geq v_R^*)-\tilde{v}}{(1-g(\tilde{v}))}\\
=\ &\lim_{s\to 0}\frac{-\frac{g(v_R^*)s}{(1-G(v_R^*))^3}-\frac{\partial\tilde{v}}{\partial s}}{-g(\tilde{v})\frac{\partial\tilde{v}}{\partial s}}=\frac{1}{g(\tilde{v})}\left(\lim_{s\to 0}\frac{g(v_R^*)s}{(1-G(v_R^*))^3\frac{\partial\tilde{v}}{\partial s}}+1\right)\\
=\ &\frac{1}{g(\tilde{v})}\lim_{s\to 0}\left(1+\frac{-g(v_R^*)s}{(1-G(v_R^*))^3\left(\frac{1-G(v_R^*)}{1-G(\tilde{v})}+\frac{g(v_R^*)s}{(1-G(\tilde{v}))}\frac{1}{1-G(v_R^*)}\right)}\right)\\
=\ &\frac{1}{g(\tilde{v})}\lim_{s\to 0}\left(1-\frac{1}{\frac{(1-G(v_R^*))^4}{(1-G(\tilde{v}))g(v_R^*)s}+\frac{(1-G(v_R^*))^2}{(1-G(\tilde{v}))}}\right)\\
=\ &\frac{1}{g(\tilde{v})}\lim_{s\to 0}\left(1-\frac{1}{\frac{(g(\overline{v})(\overline{v}-v_R^*))^4}{g(\overline{v})(\overline{v}-\tilde{v})g(v_R^*)s}+\frac{(g(\overline{v})(\overline{v}-v_R^*))^2}{g(\overline{v})(\overline{v}-\tilde{v})}}\right),
\end{aligned}$$

where the first step uses l'Hopital and the third step uses the fact that $-\left(1-G(\tilde{v})\right)d\tilde{v}=\left(1-G(v_R^*)\right)ds-g(v_R^*)sdv_R^*$ and $-\left(1-G(v_R^*)\right)dv_R^*=ds$ yields that in a neighborhood of $s=0$

$$\frac{\partial\tilde{v}}{\partial s}=-\frac{\left(1-G(v_R^*)\right)}{\left(1-G(\tilde{v})\right)}-\frac{g(v_R^*)s}{\left(1-G(v_R^*)\right)\left(1-G(\tilde{v})\right)}.$$

Now, as $s=\int_{v_R^*}^{\overline{v}}(1-G(v))dv\equiv h(v_R^*)$, and we know that $v_R^*(s)=\overline{v}$ at $s=0$ we can give a second-order Taylor approximation of $h(v_R^*)$ in a neighborhood of $\overline{v}$ as

$$\begin{aligned}
h(v_R^*) &= 0+(1-G(\overline{v}))(\overline{v}-v_R^*)+\frac{1}{2}g(\overline{v})(\overline{v}-v_R^*)^2+o\left((\overline{v}-v_R^*)^2\right)\\
&= \frac{1}{2}g(\overline{v})(\overline{v}-v_R^*)^2+o\left((\overline{v}-v_R^*)^2\right),
\end{aligned}$$

where $o(\cdot)$ is Laudau's little-o notation. Thus, we have that in a neighborhood of $s = 0$

$$\overline{v} - v_R^* = \sqrt{\frac{2s}{g(\overline{v})}} + o(\sqrt{s}).$$

Similarly, as $\int_{\tilde{v}}^{\overline{v}}(1 - G(v))dv \equiv h(\tilde{v})$ we have that the second-order Taylor approximation of $h(\tilde{v})$ in a neighborhood of $\overline{v}$ is

$$h(\tilde{v}) = \frac{1}{2}g(\overline{v})(\overline{v} - \tilde{v})^2 + o\left((\overline{v} - \tilde{v})^2\right),$$

so that in a second-order approximation $\int_{\tilde{v}}^{\overline{v}}(1 - G(v))dv = (1 - G(v_R^*))\, s$ yields

$$\frac{1}{2}g(\overline{v})(\overline{v} - \tilde{v})^2 + o\left((\overline{v} - \tilde{v})^2\right) = g(\overline{v})s\left(\sqrt{\frac{2s}{g(\overline{v})}} + o(\sqrt{s})\right)$$

or

$$\overline{v} - \tilde{v} = \sqrt{2s\sqrt{\frac{2s}{g(\overline{v})}}} + o(s\sqrt{s}).$$

Thus,

$$\lim_{s\to 0}\frac{E(v|v \geq v_R^*) - \tilde{v}}{(1 - G(\tilde{v}))} = \frac{1}{g(\tilde{v})}\lim_{s\to 0}\left(1 - \frac{1}{\frac{(g(\overline{v})(\overline{v}-v_R^*))^4}{g(\overline{v})(\overline{v}-\tilde{v})g(v_R^*)s} + \frac{(g(\overline{v})(\overline{v}-v_R^*))^2}{g(\overline{v})(\overline{v}-\tilde{v})}}\right)$$

$$= \frac{1}{g(\tilde{v})}\lim_{s\to 0}\left(1 - \frac{1}{\frac{o(s^2)}{o(s^{1\frac{3}{4}})} + \frac{o(s)}{o(s^{\frac{3}{4}})}}\right) = -\infty.$$

We next consider that $\delta > 0$ and small and argue that the platform can use the same algorithm. Consumer t with $T_1 < t \leq T_1 + T_2$ has a higher continuation value than u_R^*, some $\tilde{u} > u_R^*$. Again, for sufficiently low search costs, i.e., $s \in (\hat{s}(\delta), \bar{s}(\delta)]$, $\tilde{u} = \bar{u}$ while $u_R^* < \bar{u}$. Thus, these consumers would prefer to inspect the highly-ranked items to find an item with $u \geq \tilde{u}$ rather than settle for $(1 - \delta)E(v|u \geq u_R^*) + \delta E(\varepsilon)$. If δ is very small, then even for small s, $v_R^* < \bar{v}$. Thus, the platform learns more than the random cutoff, $\hat{v} > v_R^*$ also for some $\delta > 0$.

Last, we show that the platform can no longer achieve the cutoff v^* for any combination of δ and s such that $v^* < \bar{v}$, i.e., if $\delta < \hat{\delta}$. Recall that v^* is such that $u^* = (1 - \delta)v^* + \delta\underline{\varepsilon}$ and all items are accepted in the steady state. The cutoff v^* is the highest possible cutoff that can be reached when

blind buying is allowed because the platform can learn weakly less when blind buying is allowed versus prohibited. Consider a potential steady state with cutoff v^* and a $t = T$ so large that the steady state is almost reached. Then consumers $t > T$ expect to get a utility (almost) equal to $E(u|v > v^*) - s$ when inspecting, whereas they expect to get (almost) $E(u|v > v^*)$ when buying blindly. Thus, all consumers $t > T$, i.e., sufficiently far in the queue, prefer to buy blindly and the steady state with $v^* < \bar{v}$ cannot be reached.

Proof of Proposition 7.

From Lemma 2 it follows that for $\delta < \hat{\delta}$, a consumer facing competing platforms prefers to search over random items rather than items that were rejected before. Thus, platforms will not prominently display rejected items in their rankings. Accordingly, for $\delta < \hat{\delta}$ the cutoff value v_C^* solves

$$(1-\delta)v_C^* + \delta\underline{\varepsilon} = V_C(v_C^*),$$

where $V_C(v_C^*)$ is the continuation value if the next firm is a randomly drawn firm. For ease of reading, we let $P(\triangle)$ and $P(\square)$ denote the probabilities with which an item has a utility either in the triangle or in the rectangle as indicated in Figure 1:

$$P(\triangle) := \int_{v_C^* - \frac{\delta(\bar{\varepsilon}-\underline{\varepsilon})}{1-\delta}}^{v_C^*} \int_{\underline{\varepsilon}+\frac{(1-\delta)(v_C^*-v)}{\delta}}^{\bar{\varepsilon}} f(\varepsilon)g(v)d\varepsilon dv,$$

and

$$P(\square) := 1 - G(v_C^*).$$

Then we can write

$$V_C(v_C^*) = (P(\triangle) + P(\square)) \left\{ \frac{P(\square)}{P(\triangle)+P(\square)} \left[(1-\delta)E(v|v > v_C^*) + \delta E(\varepsilon)\right] \right.$$
$$\left. + \frac{P(\triangle)}{P(\triangle)+P(\square)} E(u|u \in \triangle) \right\} + (1 - P(\triangle) - P(\square)) \left[(1-\delta)v_C^* + \delta\underline{\varepsilon}\right] - s,$$

i.e., with probability $P(\triangle) + P(\square)$ the consumer gets a higher utility (the product has a value either in the rectangle or in the triangle) and in that case the expected utility is the expected utility of either being in the rectangle or in the triangle. With the remaining probability she gets a lower utility and buys at the reservation value $(1-\delta)v_C^* + \delta\underline{\varepsilon}$. For $\delta < \hat{\delta}$ the reservation value v_C^* is the v^* that solves

$$(1-\delta)v^* + \delta\underline{\varepsilon} = P(\square)\left[(1-\delta)E(v|v > v^*) + \delta E(\varepsilon)\right] \tag{13}$$

$$+P(\triangle)E(u|u\in\triangle)+(1-P(\triangle)-P(\square))\left[(1-\delta)v^*+\delta\underline{\varepsilon}\right]-s.$$

Recall that in case of a monopoly, the cutoff v^* is defined by

$$(1-\delta)v^*+\delta\underline{\varepsilon}=(1-\delta)E(v|v\geq v^*)+\delta E(\varepsilon)-s. \tag{14}$$

Now the LHS of (14) and (13) as functions of v^* are identical. The RHS of (13) is lower than in (14) as long as, for a given v^*, $E(u|u\in\triangle)<E(u|u\in\square)=(1-\delta)E(v|v>v^*)+\delta E(\varepsilon)$.

To show that $E(u|u\in\triangle)<E(u|u\in\square)$ we argue that $E(u|v=\tilde{v})$ increases in $\tilde{v}$ if $\tilde{v}$ is in the support of the triangle, i.e., if $\tilde{v}\in\left[v^*-\frac{\delta(\bar{\varepsilon}-\underline{\varepsilon})}{1-\delta},v^*\right]$. Let $\tilde{\varepsilon}\equiv\underline{\varepsilon}+\frac{(1-\delta)(v_C^*-\tilde{v})}{\delta}$. Then we can write

$$\begin{aligned}
E(u|v &= \tilde{v})=(1-\delta)\tilde{v}+\delta E\left(\varepsilon|\varepsilon>\underline{\varepsilon}+\frac{(1-\delta)(v_C^*-\tilde{v})}{\delta}\right)\\
&= (1-\delta)\tilde{v}+\frac{\delta}{1-F\left(\tilde{\varepsilon}\right)}\int_{\tilde{\varepsilon}}^{\bar{\varepsilon}}f(\varepsilon)\varepsilon d\varepsilon\\
&= (1-\delta)\tilde{v}+\delta\tilde{\varepsilon}+\delta\int_{\tilde{\varepsilon}}^{\bar{\varepsilon}}\frac{1-F(\varepsilon)}{1-F\left(\tilde{\varepsilon}\right)}d\varepsilon\\
&= (1-\delta)v_C^*+\delta\underline{\varepsilon}+\delta\int_{\tilde{\varepsilon}}^{\bar{\varepsilon}}\frac{1-F(\varepsilon)}{1-F\left(\tilde{\varepsilon}\right)}d\varepsilon,
\end{aligned}$$

where the one but last equality follows by integrating by parts. Now

$$\frac{\partial E(u|v=\tilde{v})}{\partial\tilde{v}}=(1-\delta)\left[1-\int_{\tilde{\varepsilon}}^{\bar{\varepsilon}}\frac{(1-F(\varepsilon))f\left(\tilde{\varepsilon}\right)}{\left(1-F\left(\tilde{\varepsilon}\right)\right)^2}d\varepsilon\right],$$

which is positive as long as the term in the squared brackets is positive. But showing that this term is positive is analogous to showing that the expression in (9) is negative, which we did in the proof of Proposition 3. Thus, $E(u|u\in\triangle)<E(u|u\in\square)$ and the solution to (13), if it exists, is smaller than for (14).

We now consider the dynamics. For simplicity of the arguments, we consider two competing platforms and show that a platform always has an incentive to deviate to a ranking that "experiments less" than the other's ranking. An important part of the proof is to consider what "experimenting less" means. Before that, note that since we consider $\delta<\hat{\delta}$, by Lemma 2 a platform that maximizes consumer surplus never wants to rank prominently items that have been inspected and not bought. Thus, the only question is how platforms rank uninspected items versus the items that have been bought in the past and subsequently never rejected (call these item Bs).

Note that as long as a platform does not experiment, there is only one item B on that platform. We say that platform 1 "experiments first at t" if at t platform 1 ranks the only item B on a prominent spot with a probability less than one.

We first show that, all else equal, a consumer prefers a ranking where an item B is on a single position with probability one rather than a ranking where item B is on position k with probability $\pi > 1/2$ and another position k' with probability $1-\pi$. To see that, let us consider the expected value that the consumer receives from searching through the latter ranking. Once the consumer draws an item with utility u, his value function is

$$V(u) = \max\{u, E_{u'}[V'(u')]\} - s,$$

where V and V' can differ because the consumer searches through the ranking in a particular order. Note that $V(u)$ is convex. Before drawing u, the expected value from searching through the ranking starting with positions k and k' is

$$E_u[V(u)] = \pi E[\max\{u, V_R(u')\}|u \in \triangle \cup \square] + (1-\pi)E[\max\{u, V''(u')\}] - s,$$

where $V''(u') = \max\{u', E[V_R(u'')|u \in \triangle \cup \square]\} - s$. The expression can be read as follows. With probability π, the consumer inspects item B first (since it is on position k) in which case he knows that it is either in the "triangle" or "rectangle" because it has been bought before. The consumer accepts B if it is better than continuing search. If he rejects B, he knows that the rest of the items in the ranking are randomly drawn uninspected items and V_R is defined as the value of searching over these. With probability $1-\pi$, a random item is on position k and the consumer knows that item B is on the next-searched position k' so his continuation value takes that into account. The fact that $V(u)$ is convex and expectation is a linear operator imply that $E_u[V(u)]$ is convex, too. But then the expected value $E_u[V(u)]$ is higher if outcomes are more spread, i.e., at $\pi = 1$ rather than an interior π. Thus, a consumer prefers a ranking that places item B on a single position with probability one rather than on two positions with some interior probabilities. By an analogous argument, a consumer prefers a ranking that places an item B on a single position with probability one rather than on more than two positions with some interior probabilities.

Now we show that if platform 1 experiments with a positive probability, then platform 2 wants to undercut it in terms of experimentation. In particular, suppose that neither platform experiments at $t < \tau$ and platform 1 experiments with a positive probability at τ: platform 1's ranking places item B on a prominent position with a probability $\pi_\tau < 1$ at $t = \tau$. We claim

that platform 2 benefits from deviating to a ranking that at τ experiments with probability $\pi'_\tau < \pi_\tau$.

Since consumer $t = 1$ expects both platforms to make equally good recommendations, he is indifferent between them. Consumer $t = 2$ then expects both platforms to have attracted the same number of consumers at $t = 1$. Thus, he, too, is indifferent between the platforms because he expects them to make equally good recommendations. This is true for all consumers $t < \tau$. Now consider consumer τ. He knows that platform 1 experiments and, as we showed above, strictly prefers searching on platform 2 if it experiments less at τ. Thus, platform 2 has an incentive to post a ranking that at τ experiments with a slightly lower probability than π_τ, for two reasons. First, it attracts consumer τ. Second, platform 2 learns (almost) the maximal amount that it would have learned if it had experimented with any other probability $\hat{\pi}_\tau$: if $\hat{\pi}_\tau > \pi_\tau$, platform 2 expected to learn nothing at τ as consumer τ would have searched on platform 1; if $\hat{\pi}_\tau < \pi_\tau$, platform 2 would have attracted consumer 2, but learned less.

As a result of this deviation, consumer $\tau + 1$ strictly prefers platform 2 (if it expects the platforms to experiment the same amount at $\tau + 1$) because he knows that platform 2 has attracted one additional consumer (in expectations) and, thus, is likely to make him a better recommendation. As a result, platform 2 strictly benefits from undercutting platform 1 in terms of experimentation in all periods where platform 1 experiments with a positive probability.

Thus, the only potential equilibrium is where no platform experiments. To see that this constitutes an equilibrium, consider a deviation by platform 1 to an algorithm that experiments first with some consumer $\tau \geq 2$ (as platforms experiment with consumer $\tau = 1$ by definition). By a similar argument as above, any consumer $\tau \geq$ strictly prefers a platform that doesn't experiment to platform 1. Thus, this deviation is not profitable and not experimenting constitutes an equilibrium for $\delta < \hat{\hat{\delta}}$. The steady state as described in Proposition 7 is reached in the limit as $t \to \infty$.

Proof of Proposition 8

Suppose that a consumer does not know if he is second or third. On a non-experimenting platform the consumer makes the same utility of

$$-s + E(v|v \geq v_R^*),$$

regardless of his true position as these consumers buy immediately a product with a value larger than v_R^*, where

$$v_R^* = -s + (1 - F(v_R^*))E(v|v \geq v_R^*) + F(v_R^*)v_R^*.$$

Consider that the experimenting platform experiments with probability π implying that it fully randomizes the ranking for consumer 2. Using the arguments of the proof of Proposition 7 it is clear that this is the optimal way for the platform to experiment. On an experimenting platform the second and third consumers use a different cut-off value $\hat{v}$. It is clear that $\hat{v} > v_R^*$ as there is a possibility that the next item they search is better than a random item. If they observe a $v < v_R^*$ then they know for sure that the platform experimented with them and that they are on the second position. They then face just a random search problem. If they observe a $v > v_R^*$ in the first spot, then they update their belief they are on the second position (and that the platform experimented with them) as $\frac{\pi\left(1-F(v_R^*)\right)+(1-\pi)}{1-\pi F(v_R^*)+1} = \frac{1-\pi F(v_R^*)}{2-\pi F(v_R^*)}$.

To determine $\hat{v}$ we have to determine the expected pay-off of a consumer who continues to search after observing a $v > v_R^*$. If they are in the second position in the queue their next search is just random. If they are in the third position, they know there is a probability $1-\pi$ that the platform did not experiment with the second consumer in which case their next search is also from a random order, while with a probability π the platform experimented with the second consumer in which case they know the next item also has a value $v > v_R^*$ and their expected pay-off of continuing to search is $\frac{1-F(\hat{v})}{(1-F(v_R^*))}E(v|v \geq \hat{v}) + \frac{F(\hat{v})-F(v_R^*)}{(1-F(v_R^*))}\hat{v}$. Thus, the overall expected value of search after observing a $v > v_R^*$ is

$$\begin{aligned} &-s + \frac{1-\pi F(v_R^*)+1-\pi}{2-\pi F(v_R^*)}\left[(1-F(\hat{v}))E(v|v \geq \hat{v}) + F(\hat{v})\hat{v}\right] \\ &+\frac{\pi}{2-\pi F(v_R^*)}\left[\frac{1-F(\hat{v})}{(1-F(v_R^*))}E(v|v \geq \hat{v}) + \frac{F(\hat{v})-F(v_R^*)}{(1-F(v_R^*))}\hat{v}\right]. \end{aligned}$$

The cut-off value $\hat{v}$ should be equal to that, or

$$\hat{v} = E(v|v \geq \hat{v}) - \frac{(1-F(v_R^*))\,(2-\pi F(v_R^*))}{((1-F(v_R^*))\,(2-\pi F(v_R^*)) + \pi F(v_R^*))\,(1-F(\hat{v}))}s.$$

The expected utility of the second consumer in the queue is now given by

$$-s+\pi F(v_R^*)v_R^*+(1-\pi F(v_R^*))\left\{\frac{1-F(\hat{v})}{1-F(v_R^*)}E(v|v \geq \hat{v}) + \frac{F(\hat{v})-F(v_R^*)}{(1-F(v_R^*))}\left[F(v_R^*)E(v|v_R^* < v \leq \hat{v})\right.\right.$$

$$\left.\left. + (F(\hat{v})-F(v_R^*))\,E(\max\{v_1,v_2\}|v_R^* \leq v_1, v_2 \leq \hat{v}) + (1-F(\hat{v}))\,E(v|v \geq \hat{v}) - s\right]\right\},$$

while for the third consumer it is

$$-s+\frac{1-F(\hat{v})}{1-F(v_R^*)}E(v|v\geq\hat{v})+\frac{F(\hat{v})-F(v_R^*)}{1-F(v_R^*)}\left\{-s+(1-\pi)F(v_R^*)E(v|v_R^*<v\leq\hat{v})\right.$$

$$\left.+[1-(1-\pi)F(v_R^*)]\left[\frac{F(\widehat{v})-F(v_R^*)}{1-F(v_R^*)}E(\max\{v_1,v_2\}|v_R^*\leq v_1,v_2\leq\widehat{v})+\frac{1-F(\widehat{v})}{1-F(v_R^*)}E(v|v\geq\widehat{v})\right]\right.$$

Let the sum of the expected utilities of the second and third consumers be denoted K. Let also $x=v_R^*$ and $y=\hat{v}$ and $w_1:=E(v|v\geq y)$, $w_2:=E(\max\{v_1,v_2\}|x\leq v_1,v_2\leq y)$, $w_3:=E(v|v\geq x)$, and $w_4:=E(v|x<v\leq y)$ so that $w_1>w_2>w_4>x$ and $w_1>w_3>w_4$. The total effect of π on K is composed of a direct effect and an indirect effect via $\hat{v}$. However, as $\hat{v}$ is chosen optimally for every π this indirect effect is zero by the envelope theorem. To evaluate the total effect of π on K we therefore only have to evaluate the direct effect. Using the notation introduced above, this direct effect is

$$\frac{\partial K}{\partial\pi}=F(x)x-F(x)\left\{\frac{1-F(y)}{1-F(x)}w_1+\frac{F(y)-F(x)}{1-F(x)}\left[-s+F(x)w_4+(F(y)-F(x))\,w_2\right.\right.$$

$$\left.\left.+(1-F(y))\,w_1\right]\right\}+\frac{F(y)-F(x)}{1-F(x)}\left\{-F(x)w_4+F(x)\left[\frac{F(y)-F(x)}{1-F(x)}w_2+\frac{1-F(y)}{1-F(x)}w_1\right]\right\},$$

or, since $s=(1-F(x))(w_3-x)$,

$$\frac{(1-F(x))}{F(x)}\frac{\partial K}{\partial\pi}=(1-F(x))x-\frac{[1-F(y)]^2}{1-F(x)}w_1$$

$$-(F(y)-F(x))\left[-(1-F(x))(w_3-x)+F(x)w_4+(F(y)-F(x))\,w_2+(1-F(y))\,w_1\right]$$

$$+[F(y)-F(x)]\left[-w_4+\frac{F(y)-F(x)}{1-F(x)}w_2\right].$$

This can be rewritten as

$$\frac{(1-F(x))}{F(x)}\frac{\partial K}{\partial\pi}=(1-F(x))x[1-(F(y)-F(x))]+[F(y)-F(x)]w_1$$

$$-\{1-F(x)-F(x)(F(y)-F(x))\}w_1-(F(y)-F(x))\left[-(1-F(x))w_3+F(x)w_4\right]$$

$$-[F(y)-F(x)]w_4+\frac{[F(y)-F(x)]^2F(x)}{1-F(x)}(w_2-w_1).$$

Let $A:=F(y)$ and $B:=F(x)$, with $A>B$. *Then*

$$\frac{(1-B)}{B}\frac{\partial K}{\partial\pi}=(1-B)x[1-(A-B)]+[-1+A+A-B]w_1-(A-B)(1+B)w_4$$

$$+(A-B)(1-B)(w_3-w_1)+\frac{(A-B)^2B}{1-B}(w_2-w_1)$$
$$=-(1-A)(w_1-x)+(A-B)(w_1-w_4)-B(A-B)(w_4-x)$$
$$-(A-B)(1-B)(w_1-w_3)-\frac{(A-B)^2B}{1-B}(w_1-w_2).$$

Note that, since $(1-A)w_1+(A-B)w_4=(1-B)w_3$, we can write

$$\frac{(1-B)}{B}\frac{\partial K}{\partial \pi}=-(1-A)(w_1-x)+(A-B)(w_3-w_4)-B(A-B)(w_4-x)$$
$$+B(A-B)(w_1-w_3)-\frac{(A-B)^2B}{1-B}(w_1-w_2)$$
$$=-(1-A)(w_3-x)-B(A-B)(w_4-x)+B(A-B)(w_1-w_3)-\frac{(A-B)^2B}{1-B}(w_1-w_2)$$

$$=-(1-A)(w_3-x)-B(A-B)(w_4-x)+\frac{(A-B)B}{1-B}[(1-B)(w_1-w_3)-(A-B)(w_1-w_2)]$$

$$=-(1-A)(w_3-x)-B(A-B)(w_4-x)+\frac{(A-B)^2B}{1-B}(w_2-w_4)$$
$$<-(1-A)(w_3-x)-\frac{(A-B)^2B}{1-B}(w_4-x)+\frac{(A-B)^2B}{1-B}(w_2-w_4)$$
$$=-(1-A)(w_3-x)-\frac{(A-B)^2B}{1-B}(2w_4-x-w_2).$$

Recall the definitions for w_2 and w_4: we can rewrite

$$2w_4-w_2-x=2(w_4-x)-(w_2-x)$$
$$=2E(v-x|x\leq v\leq y)-E(\max\{v_1-x,v_2-x\}|x\leq v_1,v_2\leq y).$$

As the maximum of some distribution $\tilde{F}$ is given by $2\int f(v)F(v)vdv$ it is clear that this expression is larger than zero so that $\frac{\partial K}{\partial \pi}<0$.

References

Anderson, S. P. and R. Renault (1999): "Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model," RAND Journal of Economics, 719–735.

Banerjee, A. V. (1992): "A simple model of herd behavior," Quarterly Journal of Economics, 107, 797–817.

Bar Isaac, H. and S. Shelegia (2022): "Monetizing Steering," CEPR Discussion Paper DP17685.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A theory of fads, fashion, custom, and cultural change as informational cascades," Journal of political Economy, 100, 992–1026.

Bolton, P. and C. Harris (1999): "Strategic experimentation," Econometrica, 67, 349–374.

Ely, J. C. (2017): "Beeps," American Economic Review, 107, 31–53.

Ely, J. C. and M. Szydlowski (2020): "Moving the goalposts," Journal of Political Economy, 128, 468–506.

European Parliament (2022a): "Digital Markets Act," `https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2022.265.01.0001.01.ENG`, [Online; accessed 01.03.24].

——— (2022b): "Digital Services Act," `https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022R2065`, [Online; accessed 01.03.24].

Glazer, J., I. Kremer, and M. Perry (2021): "The wisdom of the crowd when acquiring information is costly," Management Science, 67, 6443–6456.

Goeree, J. K., T. R. Palfrey, and B. W. Rogers (2006): "Social learning with private and common values," Economic Theory, 28, 245–264.

Hagiu, A. and B. Jullien (2011): "Why do intermediaries divert search?" RAND Journal of Economics, 42, 337–362.

Janssen, M., T. Jungbauer, M. Preuss, and C. Williams (2023): "Search Platforms: Big Data and Sponsored Positions," Working paper.

Janssen, M. and C. Williams (2023): "Influencing Search," RAND Journal of Economics, forthcoming.

Keller, G., S. Rady, and M. Cripps (2005): "Strategic experimentation with exponential bandits," Econometrica, 73, 39–68.

Kremer, I., Y. Mansour, and M. Perry (2014): "Implementing the "wisdom of the crowd"," Journal of Political Economy, 122, 988–1012.

Maglaras, C., M. Scarsini, D. Shin, and S. Vaccari (2021): "Social learning from online reviews with product choice," Columbia Business School Research Paper.

Orlov, D., A. Skrzypacz, and P. Zryumov (2020): "Persuading the principal to wait," Journal of Political Economy, 128, 2542–2578.

Renault, J., E. Solan, and N. Vieille (2017): "Optimal dynamic information provision," Games and Economic Behavior, 104, 329–349.

Smith, L. and P. Sørensen (2000): "Pathological outcomes of observational learning," Econometrica, 68, 371–398.

Smolin, A. (2021): "Dynamic evaluation design," American Economic Journal: Microeconomics, 13, 300–331.

Teh, T.-H. and J. Wright (2022): "Intermediation and steering: Competition in prices and commissions," American Economic Journal: Microeconomics, 14, 281–321.

Weitzman, M. L. (1979): "Optimal Search for the Best Alternative," Econometrica, 641–654.

Wolinsky, A. (1986): "True monopolistic competition as a result of imperfect information," Quarterly Journal of Economics, 101, 493–511.