# What Platforms Learn from Consumer Choices?

Maarten C.W. Janssen* Eeva Mauring†

March 15, 2024

**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 if a monopoly platform wants to maximize consumer's utility in the steady state, it will first experiment with product rankings, so that it can provide 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. Relative to this benchmark, learning is much more restricted if the platform maximizes revenues from a sponsored position or if there are competing platforms.

**JEL Classification**: D40, D83, L10

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

*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, like Google, Amazon, Booking.com and Tripadvisor, possess large amounts of data on what items people have searched for and what they eventually bought. Platforms potentially can use this information to help future consumers economize on their search activities, but there are increasing concerns that platforms may use this information to their own benefit without benefiting consumers. Regulators around the world are putting rules in place to restrict the ways platforms can use this type of 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 requires "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 depends on how they use the information. In particular, if they present consumers always with the same type of product ranking consumers may make similar choices, restricting the type of inferences platforms make as it is not clear whether the choices consumers make reflect "true preferences" or mostly reflect the ranking the platforms themselves chose. By experimenting with different rankings, and observing the search behaviour and consumer choices given different rankings, platforms potentially may learn more about consumers' 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 focus on how the extent of optimal experimentation by platforms depends on their objectives and on the environment they face. Platforms may care about consumer well-being or about selling prominent slots in their ranking to firms who may be willing to pay for being ranked first. In terms of environment, platforms may enjoy quite some monopoly power or they may face competition by other platforms. We also address the welfare consequences of different objectives and environments.

We focus on online markets where consumers' valuations for a product have a common and an idiosyncratic component. Consumers have to pay an inspection (or search) cost to learn their value for the product of each firm. They inspect products sequentially and their behaviour is characterized by an optimal stopping rule. Platforms observe the items the consumer has inspected and which item the consumer eventually has bought. They choose an algorithm that maps the observed history into a ranking of items for the next consumer. Consumers

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

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

can follow the ranking if they like, but are free to search on the platform in any order they prefer. Consumers enter the market sequentially, but they do not know their position in the queue: only the platform knows how many previous consumers there are 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 the platform's objectives and environments, i.e., we look at the steady state of the process described.

Before presenting our results, it is important to mention that the environment we study has some similarity 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, consumers do not know their position in the queue and therefore cannot condition their behaviour on it.Second, the platform does not observe the value of the consumer, but only whether the consumer inspected a product or bought it. As consumers have to pay an inspection cost, it may well be that they simply bought a product as it was "good enough" in the sense that they did not want to pay another inspection cost to inspect another unknown item. Third, there is an idiosyncratic component to consumers' product valuation, so what is good enough for some consumers may not be good enough for others.[3]

We have three sets of main results. First, if a monopoly platform aims to eventually provide the best possible recommendation for consumers, then the larger the idiosyncratic component, the better it is able to learn which products have the highest common values. 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 get 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 randomising 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 values 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.

[3]In some papers in the social learning literature, e.g., Smith and Sørensen (2000) and Goeree et al. (2006), agents' preferences are heterogeneous.

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 their common value component, the more the platform can learn about it.

Second, we consider competition between platforms. We take it that platforms choose their ranking algorithms and that consumers choose to use the platform that provides the best expected ranking.[4] We show that competition significantly limits the possibility for platforms to experiment. The reason is best understood by reconsidering the result on monopoly platforms. Here, by experimenting, the platform consciously chooses at any point in time not to provide the best possible ranking 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. Even though the monopoly analysis also requires to carefully address the issue of learning, learning is even more critical under competition. In both settings consumers may learn from inspecting some products they do not value whether they are early users of the platform, but under competition there is an additional effect in that learning may give rise to consumers switching platforms.

Our final set of results pertains to the platform having another objective, namely to maximize the revenue from selling a sponsored position. We show that in this case the platform has no incentive to experiment. As consumers are free to search on the platform in any order they like, the platform actually has an incentive to lower the continuation value of search because it wants to maximize the revenue from the sponsored position. Not providing information in the organic (i.e., non-sponsored) slots by uniformly randomizing them, is the best way to do so. Maximizing revenue means maximizing the probability the consumer buys from the sponsored slot. To maximise this probability, the common value component of the firm in the sponsored position is not so important as long as it exceeds the expected continuation value of searching other prod-

[4]That is, consumers know the algorithms the platforms use. This assumption is in line with the Digital Service Act (paragraph 70) that stipulates "online platforms should consistently ensure that recipients of their service are appropriately informed about how recommender systems impact the way information is displayed, and can influence how information is presented to them. They should clearly present the parameters for such recommender systems in an easily comprehensible manner to ensure that the recipients of the service understand how information is prioritised for them" (European Parliament, 2022b).

ucts in the ranking. Thus, the platform has no incentive to experiment under this alternative objective function: it recommends a product until the moment a consumer buys another product. Sooner or later, and certainly in the steady state, it recommends 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).

*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 a private 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.[5]

More generally, our papers 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-maximising platform sometimes makes consumers search more than they would like (in order to influence sellers' behaviour), but there is no learning.

[5] 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.

The rest of the paper is organized as follows. The next Section describes the baseline model of a platform aiming to maximize long-run consumer welfare, while Section 3 states the main results related to this baseline model. Section 4 then discusses the results pertaining to competition, while Section 5 provides the results in case the platform maximizes revenues it gets from the recommended position. Section 6 concludes with a discussion, while proofs can be found in the Appendix.

# 2 The Baseline Model

We consider a population of infinitely many consumers using a platform to look for a product. Products are sold by (infinitely) many firms, and 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. 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, $\varepsilon_{ij}$ 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 $[\underline{v}, \overline{v}]$, while the idiosyncratic component $\varepsilon_{ij}$ is a random draw from the CDF $F(\cdot)$ with support $[\underline{\varepsilon}, \overline{\varepsilon}]$. 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 inspect a product the consumer has to pay a (search) cost $s$ and then learns the utility $u_{ij}$ that product $j$ gives. The consumer chooses an optimal sequential search strategy and stops searching when the utility exceeds a certain threshold reservation utility. We assume that $s$ is small enough. In particular, if consumers' utility only consists of either the idiosyncratic or the common component, i.e., $\delta = 1$ (as in much of the search literature) or $\delta = 0$, then they will still search a next firm if they so far have only inspected firms whose products generate low match values, i.e., $s < \int_{\underline{\varepsilon}}^{\overline{\varepsilon}}(1-F(\varepsilon))d\varepsilon = E\varepsilon - \underline{\varepsilon}$ and $s < \int_{\underline{v}}^{\overline{v}}(1-G(v))dv = Ev - \underline{v}$. We sometimes illustrates our results for the $\varepsilon$'s and the $v$'s being uniformly distributed over the interval $[0,1]$, and for this case the assumption implies that $s < \frac{1}{2}$. This condition also guarantees that consumers want to start inspecting products if they are randomly ordered.

The platform observes that consumers search for products on their platform and which product they eventually buy. Thus, for each individual consumer, the platform knows the identity of the firms whose product the consumer has inspected and the eventual purchase, 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 the 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 also observes which consumers already have inspected and bought products, but consumers do not know their place in the queue. The only thing they know is that other consumers 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, the platform wants to maximize the 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 firms 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 $r_t$ : $\{S(\tau), b(\tau)\}_{\tau=1}^{\tau=t-1} \mathrm{x} Z \to X$, 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 of time $t$, the platform creates a ranking for consumer $t$ and the consumer engages in optimal sequential search given this ranking. The consumer understands the platform's objective and therefore may know that the platform's ranking generates valuable information to the consumer. Formally, we define a steady state as follows.

**Definition 1** *A steady state is a ranking $r^*$ and a stationary stopping rule for consumers characterized by a utility $u^*$ such that*

*(i) the ranking $r^*$ places an item in the first position that has a common utility component $v \geq v^*$ for some constant $v^*$,*

*(ii) the consumers' stopping rule tells consumers to buy any item with total utility $u \geq u^*$, with $u^* = (1-\delta)v^* + \delta\underline{\varepsilon}$ whenever $v^* < \overline{v}$.*

It is clear that if $v^* < \overline{v}$, 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. If $v^* = \overline{v}$, 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 $(1-\delta)v^* + \delta\varepsilon_{ij} < u^*$.

# 3 Results of the Baseline Model

As the platform maximizes long-term average consumer surplus, its pay-off 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. As consumers do not know their position in the ranking, they may use their observations to learn and update their beliefs about their position in the ranking. Initially, they rationally believe that the platform ranking is sufficiently informative 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 provided).

The platform can, for example, experiment more with the first consumers using the platform and provide them with a different (random) ranking (of first products). 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, consumers are rightfully more picky on accepting to buy 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 2** *The steady state is unique. If $\delta(E\varepsilon - \underline{\varepsilon}) < s$, the steady state is characterized by a ranking $r^*$ and a reservation utility $u^*$ such that $v^* < \overline{v}$ solves*

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

*$v^*$ is strictly increasing in $\delta$ and decreasing in $s$. If, on the other hand, $s < \delta(E\varepsilon - \underline{\varepsilon})$, then the steady state is characterized by $v^* = \overline{v}$.*

*The following algorithm leads to this steady state: at $t$, rank on any position a random unsampled item with probability $\pi_t$ 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 $\lim_{t\to\infty} \pi_t = 0$. If at $t$ a random item is ranked first, the nonrandom candidate item for rank 2 is the item with the highest (not second-highest) posterior $v$.*

The idea about the proposition can be easily illustrated by considering for the uniform distribution that the common component of the utility function in a steady state is given by $v^* < \overline{v} = 1$. As in such a steady state, the consumer will always buy from the first firm in the platform's ranking, the platform can

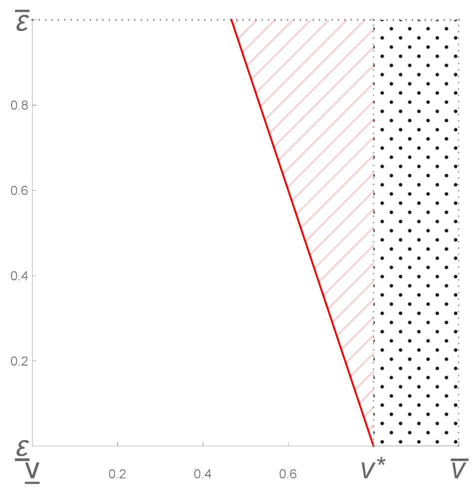


Figure 1: *Consumers accept items north-east from the thick (red) line: in the triangle (hatched) and rectangle (dotted); $s = 0.2$, $\delta = 0.25$.*

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 $(1-\delta)v^* + \delta\underline{\varepsilon}$. A consumer that draws this utility level and expects that on their next searches they only encounter products that the platform possibly also could have recommended have a common component that is at least as large as $v^*$. Thus, they will buy if, and only if,

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

where the RHS gives the expected utility level of continuing to search. This is illustrated in Figure 1. The reservation utility $u^*$ is depicted by the downward sloping line, crossing the horizontal axis at $v^*$. 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 if the idiosyncratic component is large enough).

In a steady state with $v^* < \overline{v}$ this inequality should hold with equality. For $v$ and $\varepsilon$ being uniformly distributed this gives

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

or

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

This expression 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 recommend the one with the largest $v$. In the steady state, this largest value is $\overline{v} = 1$. The proposition states the same result for general value distributions.

The analysis regarding the part of the proposition on the dynamic process towards the steady state is more intricate. The reason is that during the process towards the steady state, the consumer may encounter items that he did not expect to encounter if he was already in the steady state. As consumers do not know their place in the queue, these observations will therefore make him update his beliefs about his position in the queue. To see that this algorithm gets the system to the steady state, consider first $\delta < \hat{\delta}$ so that 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 $u_1 < u^*$, rejects this item. He also updates his beliefs, however, 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 $v_1 > v^*$, thus, $u_1 > u^*$. The worst that the buyer can believe is that he is the first one in the queue and always has to search over randomly ranked items. Thus, the cutoff $\hat{v}_t$ that the buyer uses for his purchasing decision is definitely in $[v_R^*, v^*]$. Eventually, the buyer buys some item $A$ with $v_A \geq \hat{v}_t$. 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, the platform will end up with an item $B$ on rank one that has $v_B \geq v^*$ because, first, all consumers 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 $t \to \infty$.

Now consider $\delta > \hat{\delta}$. Even an item that has $v = \bar{v}$ is 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 first an item with $v$ arbitrarily close to $\bar{v}$.

A few other interesting observations related to the proposition are in place. First, if we define

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

as the weight on the idiosyncratic component for which $v^* = \bar{v}$, then for any $\delta < \hat{\delta}$ the stead state $v^*$ is interior and increasing in $\delta$ and decreasing in $s$. 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 $\delta$ 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

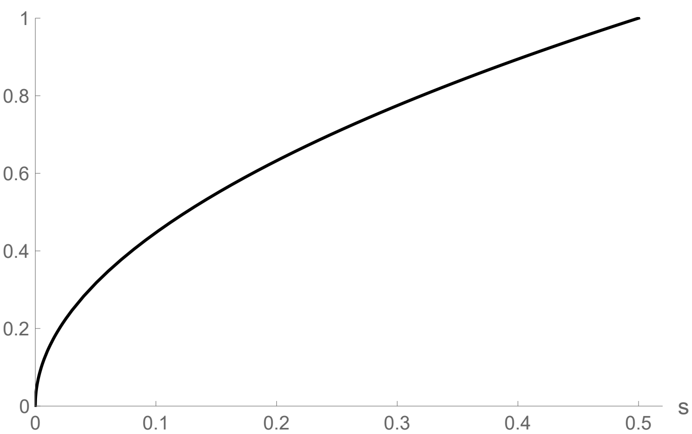


Figure 2: *The value to consumers of having (relative to not having) a platform as a function of s;* $\delta = 0$.

more inclined to experiment, making it easier for the platform to discern firms with really high values of $v$. If the weight $\delta$ 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 value $v^*$ with the standard consumer search problem in the absence of a 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 $s = \int_{\hat{v}}^{\overline{v}} (1 - G(v))dv$, which for the uniform distribution results in a value $\hat{v} = 1 - \sqrt{2s}$. Figure 2 below depicts $\frac{(1-2s)-(1-\sqrt{2s})}{1-\sqrt{2s}}$, which is the expected value 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%.

# 4 Competition

In this section we consider markets where platforms are competing with each other for consumers and ask to what extent the results of the previous Section continue to hold. In particular, 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 keep the platforms' objective as in the baseline model.

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 would 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 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. Roughly speaking, consumers will avoid using the platform that experiments most as they know this experimentation is not to their own advantage. Without consumers making inspection and purchase choices, however, a platform does not learn anything. 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. The analysis has to take into account, however, that consumers may update their beliefs about how much a platform is experimenting (and what their position in the queue is), especially after observing products that upon inspection have a low utility.

To study the implications of competition, in this section we model competition between platforms in the following stark way. First, platforms choose how to rank products in every period, i.e., they choose $r = \{r_t\}_{t=1}^{\infty}$. Second, not knowing their place in the ranking, a consumer chooses which platform to use for their search and which products to inspect and when to buy. This way of modelling assumes that all consumers know exactly what ranking algorithm the platforms use. 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 that (some) consumers will have some information about the platform's ranking (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 the previous section 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.

The main result in this section is that competition limits the extent to which platforms can learn about the common value of consumers' utility. Before we prove this result, we state the following Lemma:

**Lemma 3** *if $\delta < \hat{\delta}$, i.e., when there is an interior solution in the monopoly model, with competing platforms a consumer never wants to inspect a product* *that has been inspected but not bought by some previous consumer.*

The lemma critically uses the notion of reservation values (utilities). To see this, consider first the case where $\delta = 0$. In that case $v^* = u^*$ and it is equal to the consumers' reservation value, determined by $\int_{v^*}^{\overline{v}}(1 - G(v))dv = s$. If a product has been inspected before and not purchased, it is known to have a common component, and thus a reservation value, which is smaller than $v^*(s)$. The application of Weitzman's (1979) optimal search rule then says that any random product, which by definition has a reservation value of $v^*(s)$, should be inspected first. Note that this is true independent of $s$ as the impact of $s$ is incorporated in the reservation value: even if $v^*(s)$ itself may be relatively high (for small values of $s$) and even if consumers know a product has a value close to it, they want to inspect a random product as it may have a value higher than $v^*(s)$ and in case it has a lower value they will continue to search other products. An item with a known value $v < v^*(s)$ does not have the upward potential.

Now consider a positive, but small value of $\delta$. Let $r_C^*$ and $u_C^*$ represent a platform's ranking with $v_C^*$ the common component of the product in the first position, respectively, the stationary stopping rule in a steady state under competing platforms (where both platforms use the same ranking). Thus, we have that $u_C^* = (1-\delta)v_C^* + \delta\underline{\varepsilon}$. If a product has been inspected before, but not purchased, then it is known to have a common value smaller than $v_C^*$. In line with Figure 1, consumers may still want to inspect such items as now there is the possibility that they draw a higher idiosyncratic value than the previous consumer who did not purchase the product. However, if $\delta$ is small the value of this upward potential is small and if $\delta < \hat{\delta}$ it is not large enough to warrant spending the search cost on it.

The implication of the lemma is that because of competition, a platform will never attempt to have a consumer inspect a product that is rejected by a previous consumer who inspected the product when $\delta$, the weight on the idiosyncratic component, is relatively small.[6] If a platform would do so it gives the competitor the possibility to provide a ranking that is preferred by consumers and they will therefore not visit the platform. 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 Lemma 3 is that as the platforms cannot experiment and $\delta$ is small, they cannot rank the items in such a way that the consumers inspect items that were previously inspected and not bought. We will use this lemma to prove our main result of this section.

**Proposition 4** *If $\delta < \hat{\delta}$, the steady state of the process under competition is lower than under monopoly, i.e., $v_C^* < v^*$ and $u_C^* = u_R^* < u^*$, where $u_R^* \equiv E\max\{(1-\delta)v + \delta\varepsilon, u_R^*\} - s$ is the reservation utility under random search.*

[6] 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.

The idea behind the proposition is that a platform will always create rankings such that after the first item, consumers will inspect random products that have not yet been inspected before. To help consumers, platforms will always put a product in the top spot that the previous consumer has inspected and purchased. This product 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). As the continuation value of search is, however, based on inspecting random products after the first inspection, the consumers' reservation utility is as if there is no platform making recommendations, and is therefore lower than under a monopoly platform. Still, platforms do play an important role. First, they increase social surplus in the steady state (relative to a world without platform) as by recommending a product with a high common component they lower expected search cost as eventually all consumers buy from the firm they inspect first. Second, in the presence of platforms it becomes increasingly important for firms to be in the top spot as only firms in that spot have positive sales.

When $\delta \geq \hat{\delta}$ consumers may be willing to inspect products that previously have been rejected. However, the extent to which depends on the precise value of $\delta$. If $\delta$ 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 $v_C^*$. 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 $v_C^*$ have *not* been inspected before. Thus, for $\delta$ values marginally larger than $\hat{\delta}$ competing platforms continue to learn less about the common value relative to a monopoly platform. When $\delta$ gets larger, consumers are more inclined to inspect previously once rejected products (after having been bought several times) and platforms are more inclined to put such products high up in their rankings.

The different effects of $\delta$ imply that its comparative statics can be non-monotonic. Figure 3 below indicates that for values being uniformly distributed, the common component $v_C^*$ of the steady state value is decreasing in $\delta$ when $\delta$ and $s$ are small, while for other parameter values the effect is as in the monopoly case. This effect may seem small, but in Figure 4 we show that in a neighborhood of $\delta = 0$ the effect may actually be quite substantial if $s$ is very small.[7] The reason why competing platforms may be able to learn less when $\delta$ increases starting from small values is as follows. First, like for monopoly platforms there

[7] If $v$ and $\varepsilon$ are uniformly distributed, we can explicitly calculate the reservation value to be $v_C^* = \frac{2-\delta-\frac{\delta^2}{6(1-\delta)}-\sqrt{\left[2-\delta-\frac{\delta^2}{6(1-\delta)}\right]^2-4(1-2\delta)\left[\frac{\delta^2(2-\delta)}{6(1-\delta)^2}+1-2s\right]}}{2(1-2\delta)}$. (Details are available upon request.) It is straightforward to show that $\lim_{\delta\to 0}\frac{\partial v_C^*}{\partial \delta} = 5 - \frac{3(1+4s)}{2\sqrt{2s}}$ so that $\lim_{s\to 0}\lim_{\delta\to 0}\frac{\partial v_C^*}{\partial \delta} = -\infty$. Thus, this counterintuitive comparative statics effect of $\delta$ can bee quite large when $s$ is also small.

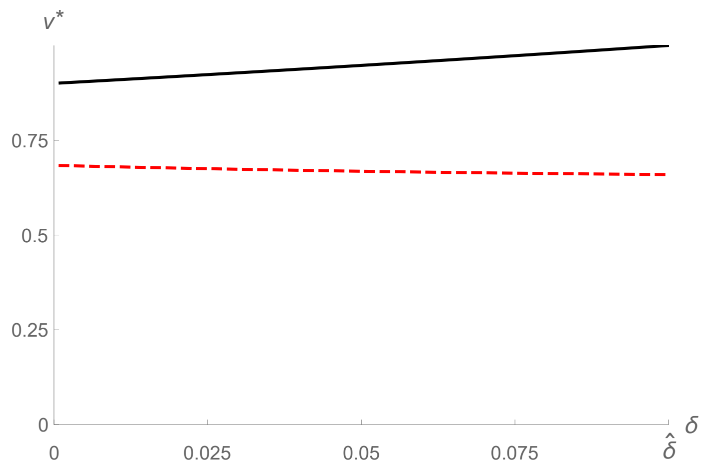


Figure 3: *Monopoly cutoff* $v^*$ *(black solid) and competitive cutoff* $v_C^*$ *(red dashed) in* $\delta$*;* $s = 0.05$*.*

is the effect that when $\delta$ gets larger consumers tend to inspect more as they know their preferences may not be fully aligned with those of their predecessors. Second, also similar to monopoly platforms consumers are willing to purchase products whose common component value is smaller than $v_C^*$, namely when their idiosyncratic value is large enough, i.e., products that are in the "triangle region" of Figure 1. However, and this is the difference with monopoly platforms, for $\delta < \hat{\delta}$ competing platforms' ranking are such that after rejecting the product in the top spot, consumers search for random products and purchase products above their reservation utility (which includes products that lie in the "triangle region" of Figure 1). Relative to the monopoly platforms, this lowers the continuation value of search for consumers. The strength of this effect depends on the size of the "triangle region" relative to the size of the "rectangle region" and this relative size is large when $s$ is small. Thus, for small $s$ this effect dominates the first effect of an increase in $\delta$ creating the nonmonotonic effect of an increase in $\delta$.

# 5 Maximizing revenue from the Recommended Position

Many platforms have different objectives than maximizing long-run consumer welfare. It is important to know how these objectives affect the platform's learning. In this section we therefore consider a platform that aims to maximize revenues from selling the top position in the ranking and ask what are the platform's incentives to learn from consumers' behaviour given this different objective function.[8] We take it that the firms, like the platform, do not have information about consumer preferences, and that, unlike the platform, do not observe consumer behavior on the platform. A firm's willingness to pay for the top spot is expressed as an amount the firm is willing to pay conditional on making a sale.[9] As this paper's focus is on what the platform can learn from observed consumer behaviour, we consider that firms are ex ante identical and that their behaviour does not form a second source of information the platform can draw upon. The analysis starts off by considering a monopoly platform, but at the end of the Section, we argue that the analysis extends beyond this case.

A few characteristics of the steady state should be immediately clear. First, firms will be willing to pay for the top spot if (and only if) this implies more sales, which is for example the case if consumers are willing to start their search there. Second, *ex ante* identical firms will make identical bids as they have no information about consumers' willingness to pay. Third,[10] the platform has an incentive to lower the continuation value of search as much as possible as this will make it more likely that a consumer will buy from the top spot. The lowest possible continuation value if consumers continue to search is obtained when products are randomly placed in the ranking as consumers can (and will) always skip inspecting products if they know (from the platform's ranking algorithm) that products on a certain place in the ranking are likely to be worse than a randomly chosen product. Given this property of the optimal ranking, consumers have a reservation utility equal to $u_R^* \equiv E \max\{(1-\delta)v + \delta\varepsilon, u_R^*\} - s$. Fourth, the platform also has an incentive to put the product with the highest likelihood of generating a sale in the top spot. For any consumer that is not the first in the ranking that is the product which the last consumer bought. As this is a product that has a reservation value larger than a random product, consumers will start their search there and the consumer will immediately buy if her utility $u$ satisfies $u \geq u_R^*$. Fifth, the platform can announce the following algorithm. Firms are ranked according to their bid. The firm with the highest bid will get the top spot until the moment a consumer decides not to purchase the product of that firm. If a consumer inspects the product at the top spot, but does not buy it, then the top spot goes to the next highest bid, while all other products get a random place in the ranking. Finally, defining $v_R^* \equiv (u_R^* - \delta\underline{\varepsilon})/(1-\delta)$ as the value of the common component that yields the consumer a utility equal to $u_R^*$ if the idiosyncratic component is at its lowest possible level $\underline{\varepsilon}$, it is not difficult to see that as long as $v_R^* < 1$ in the steady state of this process, a product gets the top spot that has a $v \geq v_R^*$.

Thus, if we define we $\hat{\delta}_R$ as the weight such that $v_R^* = 1$, we can state the following result.

[8] Considerations similar to those in the current section apply when the platform sells off multiple positions.

[9] Only payment rules that depend on the platform's performance in helping firms to generate sales make sense as otherwise the platform would not have an incentive to change consumer behaviour. This also implies that general payment rules that do not depend on a firm acquiring a certain prominent position also do not affect the platform's behavior.

[10] This and the next point are similar to the analysis in Janssen et al. (2023).

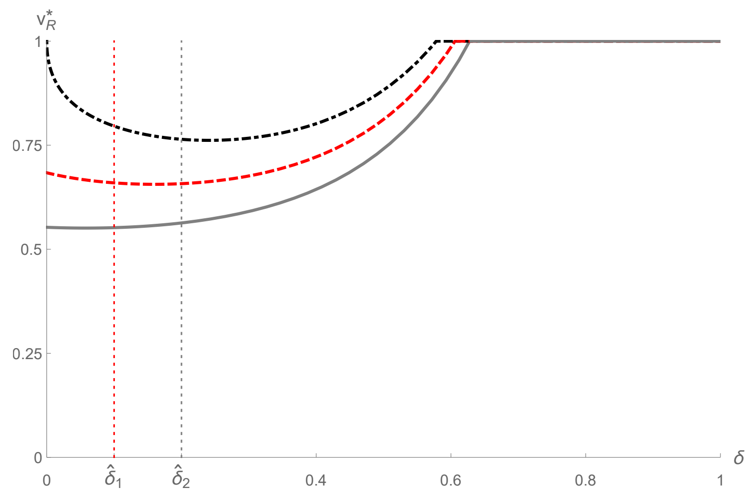


Figure 4: *The critical value $v_R^*$ as a function of $\delta$ for $s = 0$ (black dot-dashed), $s = 0.05$ (red dashed), and $s = 0.1$ (grey) where $\hat{\delta}_1 := \hat{\delta}_{s=0.05}$ and $\hat{\delta}_2 := \hat{\delta}_{s=0.1}$.*

**Proposition 5** *For any $\delta < \hat{\delta}_R$, the steady state of the process when the platform maximizes the revenue from selling the top spot has a reservation utility that is equal to $u_R^*$ and the common component is larger than or equal to $v_R^*$. We also have that $\hat{\delta}_R > \hat{\delta}$.*

The Proposition shows that if a monopoly platform sells the top spot in its ranking, the outcome mimics the outcome under competing platforms in the previous section in the sense that the role of the platform is limited to putting an item with a large enough value in the top spot. As firms are willing to pay an amount equal to their expected profits, the platform's profit is maximized. Consumers are still better off than without a platform as they economize on their search cost. The result is, in a sense, however, worse than the outcome under competition as in the present case there is no competitor constraining the platform in offering random continuation values of search. Under competition, platforms cannot offer a ranking with random products beyond the top spot if $\delta > \hat{\delta}$ as competitors could then offer a better ranking to consumers. Thus, under a monopoly platform there is a (much) wider set of $\delta$ values for which the steady state features a critical value of $v_R^* < \overline{v}$. This is illustrated in Figure 4 for the case of the uniform distribution. In the Figure the critical values of $\hat{\delta}_R$ are given by the values where (for different values of $s$) the curves hit the maximum common value of 1; whereas the critical values $\hat{\delta}$ under competition are given by $\hat{\delta}_1$ (for $s = 0.05$) and by $\hat{\delta}_2$ (for $s = 0.1$). The critical value $\hat{\delta}$ for $s = 0$ is 0 and therefore not represented.

Given our analysis so far, it is not difficult to see that the result would be similar if we consider competing platforms. The only difference is that under competition platforms should provide consumers with the best possible ranking given the information they have. Therefore, the ranking in the steady state will be exactly identical to the one in the previous section with the threshold value being $\hat{\delta}$.

# 6 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 keep our 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 product 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?

# 7 Appendix: proofs

**Proof of Proposition 2.**

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 buyer 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

$$v^* \geq E(v|v \geq v^*) + \frac{\delta(E\varepsilon - \underline{\varepsilon}) - s}{1-\delta}. \tag{2}$$

Clearly inequality (2) can hold only if $\delta(E\varepsilon - \underline{\varepsilon}) < s$ so consider such $\delta$ and

$s$. We can write (2) as

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

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

$$v^* = \overline{v} - \int_{v^*}^{\overline{v}} G(v)dv + (1 - G(v^*))\frac{\delta(E\varepsilon - \underline{\varepsilon}) - s}{1-\delta}. \tag{3}$$

Equation (3) can be rewritten as

$$\int_{v^*}^{\overline{v}} \frac{1-G(v)}{1-G(v^*)}dv = \frac{s - \delta(E\varepsilon - \underline{\varepsilon})}{1-\delta}. \tag{4}$$

Since the RHS of equation (4) is positive, we also need the LHS to be positive for an interior $v^*$ to exist. Now

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

and

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

as $g(\overline{v}) > 0$. 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 (4), or

$$\frac{\partial}{\partial v^*} \int_{v^*}^{\overline{v}} \frac{1-G(v)}{1-G(v^*)}dv = -1 + \int_{v^*}^{\overline{v}} \frac{(1-G(v))g(v^*)}{(1-G(v^*))^2}dv, \tag{5}$$

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

$$g(v^*) \int_{v^*}^{\overline{v}} \frac{1-G(v)}{1-G(v^*)}dv < 1 - G(v^*). \tag{6}$$

We show that inequality (6) holds because the RHS of the inequality decreases faster than the LHS for all $v^* < \overline{v}$, while the two sides equal at $v^* = \overline{v}$. At $v^* = \overline{v}$, 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 (6) 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 (6) 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 $\delta$. We have just shown that the LHS of equation (4) decreases in $v^*$. It is easy to show that the RHS decreases in $\delta$. Since the LHS of (4) does not depend directly on $\delta$ 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 (3) is increasing in $\delta$ 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 buyer 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 focus on the dynamics towards the steady state. Since consumers do not know their positions, they must form beliefs about them. When a buyer starts searching, he believes that he is 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 buyer 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

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

$$= \frac{P(u_1<u^*|t=1)P(t=1)}{P(u_1<u^*|t=1)P(t=1)+...+P(u_1<u^*|t=T)P(t=T)}, \qquad (7)$$

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 (7) 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 (7) 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 (7) 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 $v_R^*$ 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 Lemma 3.**

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. That is, if $\delta < \hat{\delta}$, a consumer optimally searches random items before searching any item about which she knows that it has been inspected but not bought before.

We derive the reservation value for an item that has been bought some (potentially many) times but eventually rejected. The consumer then "knows" that the common component of the product is some $\hat{v} \leq v_C^*$, where $v_C^*$ is the critical value of the common component in a steady state under competition between platforms. Denote by $\hat{r}$ the reservation value of inspecting a product with a common component of the product equal to $\hat{v}$. We show that if $\delta < \hat{\delta}$ and $\hat{v} \leq v_C^*$, then $\hat{r}$ is lower than $(1-\delta)\hat{v} + \delta\underline{\varepsilon}$, thus, $(1-\delta)v_C^* + \delta\underline{\varepsilon}$ which is the reservation value of inspecting a random item.

If a solution exists, $\hat{r}$ satisfies

$$E[\max\{u, \hat{r}\}] - \hat{r} - s = 0, \tag{8}$$

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

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

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

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

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

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

The reservation value $\hat{r}$ is implicitly defined through $\hat{\varepsilon}$. Since the LHS of this expression is decreasing in $\hat{\varepsilon}$, 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 $\hat{r} > (1-\delta)\hat{v} + \delta\underline{\varepsilon}$.

Thus, if $\delta \leq \hat{\delta}$, then for any $\hat{v}$ the reservation value is given by $\hat{r} = (1-\delta)\hat{v} + \delta E\varepsilon - s$, which is smaller than $(1-\delta)\hat{v} + \delta\underline{\varepsilon}$ which for $\hat{v} \leq v_C^*$, in turn, is smaller than $(1-\delta)v_C^* + \delta\underline{\varepsilon}$, the reservation value of a randomly selected product.

**Proof of Proposition 4.**

From Lemma 3 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

$$(1-\delta)v^* + \delta\underline{\varepsilon} = P(\square)\left[(1-\delta)E(v|v > v^*) + \delta E(\varepsilon)\right] \tag{9}$$
$$+P(\triangle)E(u|u \in \triangle) + (1 - P(\triangle) - P(\square))\left[(1-\delta)v^* + \delta\underline{\varepsilon}\right] - s.$$

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

$$(1-\delta)v^* + \delta\underline{\varepsilon} = (1-\delta)E(v|v \geq v^*) + \delta E(\varepsilon) - s. \tag{10}$$

Now the LHS of (10) and (9) as functions of $v^*$ are identical. The RHS of (9) is lower than in (10) 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 $E(u|v = \widetilde{v})$ increases in $\widetilde{v}$ if $\widetilde{v}$ is in the support of the triangle, i.e., if $\widetilde{v} \in \left[v^* - \frac{\delta(\bar{\varepsilon}-\underline{\varepsilon})}{1-\delta}, v^*\right]$. Let $\widetilde{\varepsilon} \equiv \underline{\varepsilon} + \frac{(1-\delta)(v_C^*-\widetilde{v})}{\delta}$. Then we can write

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

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

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

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 (5) is negative, which we did in the proof of Proposition 2. Thus, $E(u|u \in \triangle) < E(u|u \in \square)$ and the solution to (9), if it exists, is smaller than for (10).

THE PART ON THE DYNAMICS TO BE COMPLETED.

# 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," The 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?" The 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," The Quarterly Journal of Economics, 101, 493–511.

# 8 Online Appendix

In this online appendix we describe the dynamics that lead to the steady states in the different settings when the consumers know their position in the queue. We do this to show that for two reasons. First, for easy comparison with related literature that assumes consumers know their position in the queue, and second, to show our results are robust to this modelling assumption.

## 8.1 Monopoly platform

**Proposition 6** *If buyers know their position in the queue, then the following algorithm leads to the steady state as described in Proposition 2 as $t \to \infty$:*

- *at $t$, rank on any position a random unsampled item with probability $\pi_t$ 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 $\lim_{t\to\infty} \pi_t = 0$.*

Note that if at $t$ a random item is ranked first, the nonrandom candidate for rank 2 is the item with the highest (not second-highest) posterior $v$, etc.

**Proof.** Let the cutoffs that buyer who know his position to be $t$ be denoted $\tilde{u}_t$ and $\tilde{v}_t$ for the total utility and common utility parts respectively and note that $\tilde{v}_t \in [v_R^*, v^*]$ for all $t$.

First, let $v^* < \bar{v}$ and consider buyer $t = 1$. He knows that he samples random items, thus, optimally uses the cutoff $\tilde{v}_1 = v_R^*$. Thus, for buyer 2, the nonrandom candidate item for rank 1 is the item that buyer 1 bought because

$$E(u|v \geq v_R^*) > E(u).$$

Thus, the value of starting search for buyer 2 is

$$\begin{aligned} V(\tilde{v}_1, \pi_2) &:= \pi_2 E \max\{u, \tilde{u}_1\} + (1 - \pi_2) E(\max\{u, \tilde{u}_1\}|v > \tilde{v}_1)] - s \\ &> E \max\{u, \tilde{u}_1\} - s = u_R^*. \end{aligned}$$

Thus, the cutoff that buyer 2 uses satisfies

$$(1 - \delta)\tilde{v}_2 + \delta\underline{\varepsilon} = V(\tilde{v}_1, \pi_2),$$

with $\tilde{v}_2 > \tilde{v}_1$. By a similar argument, $\tilde{v}_3 > \tilde{v}_2$ because $E(u|v > \tilde{v}_2) > E(u|v > \tilde{v}_1)$ and $\pi_3 < \pi_2$ imply that $V(\tilde{v}_2, \pi_3) > V(\tilde{v}_1, \pi_2)$. The cutoffs of buyers that know their positions, $\{\tilde{v}_t\}_{t=1}^{\infty}$, thus, constitute a bounded strictly increasing sequence that converges to $v^*$. The steady state with $v^* < \bar{v}$ is reached in the limit as $t \to \infty$.

Now let $v^* = \bar{v}$. That is, in a steady state a buyer rejects an item with positive probability even if that item's common component has the highest possible value. Denote by $\hat{v}(\tilde{u})$ the $v$ such that $(1-\delta)\hat{v}(\tilde{u}) + \delta\bar{\varepsilon} = \tilde{u}$ and by $\check{v}(\tilde{u})$ the $v$ such that $(1-\delta)\check{v}(\tilde{u}) + \delta\underline{\varepsilon} = \tilde{u}$ if such $\check{v}(\tilde{u}) \leq \bar{v}$ exists. Both $\hat{v}(\tilde{u})$ and $\check{v}(\tilde{u})$ increase in $\tilde{u}$.

Buyer 1 uses the cutoff $\tilde{u}_1 = u_R^*$ because he knows that he faces only random items so $\check{v}(\tilde{u}_1) = v_R^*$. For buyer 2, the platform again places the item bought by buyer 1 as the the nonrandom candidate item for rank 1 because

$$E(u|v \geq \hat{v}(\tilde{u}_R)) > E(u).$$

As a result, essentially the same argument as above establishes that the buyers cutoffs $\{\tilde{u}_t\}_{t=1}^{\infty}$ constitutes a strictly increasing sequence. Thus, $\{\check{v}(\tilde{u}_t)\}_{t=1}^{\infty}$ constitutes a bounded strictly increasing sequence that converges to $v^* = \bar{v}$. The steady state with $v^* = \bar{v}$ is reached in the limit as $t \to \infty$. ■