The Foerder Institute for Economic Research
at Tel Aviv University

מכון למחקר כלכלי על שם ד"ר ישעיהו פורדר
על יד אוניברסיטת תל אביב

The Eitan Berglas School of Economics

בית-הספר לכלכלה ע"ש איתן ברגלס

עמותה רשומה

Non-Sequential-Non-Cascading Clicking in Online Consumer Search

Daniel Fershtman, Alessandro Pavan

Working paper 3-2026

The Foerder Institute for Economic Research
and
The Sackler Institute of Economic Studies

מכון פורדר, אוניברסיטת תל-אביב, תל אביב יפו 6997801 | foerder@tauex.tau.ac.il
Tel Aviv University, Tel Aviv Yafo 6997801 Israel | foerder@tauex.tau.ac.il

Non-Sequential-Non-Cascading Clicking in Online Consumer Search*

Daniel Fershtman† Alessandro Pavan‡

June 2026

Abstract

We consider the problem of a consumer searching online for the best alternative. We show how this problem can be formalized as a variant of Pandora's boxes problem where the decision maker retains control over the order in which boxes are opened, but where, contrary to the original version, the set of boxes grows over time and is controlled by a third party (the platform). We use the model to (a) endogenize click-through-rates, (b) explain the phenomenon of non-sequential-non-cascading clicking, (c) illustrate why the generalized second-price auction may lead to inefficient assignments even under its ascending-clock implementation, and (d) show why a firm's profits need not be increasing in the number of ads it displays on a platform even when they are free.

Keywords: Online consumer search, experimentation, clicking behavior, sponsored search, click-through-rates, value per click, generalized second-price auction.

*For comments and useful suggestions, we thank Teddy Kim, Stephan Lauermann, Michael Ostrovsky, Andrew Rhodes, Andy Skrzypacz, and Jidong Zhou. Fershtman gratefully acknowledges financial support from the Foerder Institute.

†Tel Aviv University. Email: danfershtman@gmail.com

‡Department of Economics, Northwestern University. Email: apavan31@gmail.com

1 Introduction

Consider the problem of a consumer searching online for the best possible alternative. The consumer alternates between (a) reading new ads (bringing the corresponding products to her consideration set), (b) clicking on the ads of those products already in the consideration set (revealing the products' value to the consumer, net of the purchasing price), and (c) finalizing the purchase with one of the visited vendors.

This problem can be formalized as a variant of Pandora's boxes problem in which the boxes are gradually introduced to the decision maker's consideration set in an order that is outside the decision maker's control. We show that this problem explains various phenomena observed in online sponsored search.

First, the model delivers a novel formula for the click-through-rates (CTRs) that accounts not only for the heterogeneity in the ads' positions but also for the externalities exerted by the ads that occupy them.[1] Other models of online consumer search, instead, either assume that CTRs are position-dependent but exogenous (e.g., Edelman et al., 2007) or that they depend only on the ads displayed upstream (e.g., Athey and Ellison, 2011). This last property is justified if consumers click on the ads in the same order in which they are displayed, but not when consumers choose the order the ads are clicked and possibly click on upstream ones only after clicking on a few downstream ones, a phenomenon referred to in the empirical literature as *non-cascading clicking*. For example, Jeziorski and Segal (2015) show that (1) approximately half of the searchers do not click on ads sequentially in the order they are displayed (*non-sequential clicking),* (2) over half of the users who click more than once, click on a higher position's ad after clicking on a lower position's ad (*non-cascading clicking*), and (3) the rate at which an ad at a given position is clicked upon depends not only on the ads displayed at higher positions but also on the ads displayed at lower positions (*externalities from lower positions*). We show that these phenomena are consistent with the dynamics of consumer search in our model. Perhaps surprisingly, this is the case even when the consumer expects lower positions, on average, to be occupied by ads that are less attractive than those

[1] A position's CTR is the probability that an ad displayed at such position is clicked.

displayed at higher position (in a sense made precise below), and the cost of reading increases with the number of ads brought to the consideration set.

The above results hinge on the assumption the products are brought to the consumer's attention sequentially in an order that is outside the consumer's control (as when the order is controlled by a platform). A sequential search model in which all products are in the consumer's consideration set from the outset (such as Weitzman (1979)'s original version of Pandora's boxes problem) fails to deliver any structural relationship between the ads' positions and the corresponding CTRs. In contrast, the model described above permits one to endogenize the CTRs and explains why the latter need not be monotone in the positions, even when the purchasing probabilities are, and why the "values per click" (VPCs) are also position-specific. These results are obtained by establishing a structural relationship relating the probability each product is purchased to the primitives of the problem (consumers' realized values for the products advertised, ads' positions, search costs, and consumer's beliefs over the attractiveness of the products displayed at the various positions, including those displayed downstream—externalities from below).

Next, we use the model to derive a few implications for bidding in sponsored-search auctions and for the efficiency of the allocations sustained in equilibrium. The analysis uncovers that the generalized second-price auction may lead to inefficient assignments even under the ascending-clock implementation considered in the literature (see, e.g., Edelman et al., 2007 and Gomes and Sweeney, 2014).

Finally, the model explains why a firm advertising on a platform may experience a decline in its profits when it is given additional ad space in the form of multiple slots for additional products advertised on the same platform. Interestingly, we show that such a decline can occur even if the extra space is provided for free and even if profits unambiguously increase with ad space when the firms' products are introduced exogenously to the consumer's consideration set.

1.1 Related literature

The paper is related to the literature on sponsored search. In addition to the papers by Edelman et al., 2007, Athey and Ellison, 2011, and Gomes and Sweeney, 2014, cited above, see also Edelman and Schwarz, 2010 and the references therein. Our contribution is in showing how to endogenize the click-through-rates and the value-per-click (with the latter naturally varying with the position occupied by the ads), and then using the characterization to explain dynamics of clicking and purchases that are consistent with the empirical literature, as reported in Jeziorski and Segal (2015), but not with standard search models.

In terms of the methodology, the consumer's problem in the present paper is a special version of the one studied in Fershtman and Pavan (2025). That paper considers a general multi-armed bandit problem with an endogenous set of arms (in which learning each arm's value may require multiple—possibly infinitely many—explorations and where batches of alternatives are brought stochastically to the consideration set as a function of past expansions) and derives conditions for the optimal policy to be "indexable," with a novel index for the expansion of the consideration set. In the present paper, we specialize the structure of our earlier work to the online consumer search problem under examination and use it to derive the various predictions described above.

Despite its many applications, relatively few extensions of Weitzman (1979)'s Pandora boxes problem have been studied in the literature. Notable exceptions include Olszewski and Weber (2015), Choi and Smith (2016), and Doval (2018). In these papers, though, the set of boxes is fixed. In the marketing literature, Greminger (2022) studies a consumer search problem with an endogenous set of products similar to ours, focusing on the comparison between directed and undirected search. Instead, we use the model to explain the phenomenon of non-sequential-non-cascading clicking observed in the empirical literature, and to draw novel implications for sponsored search auctions.

Our results linking consumer's eventual purchasing outcomes to the primitives of the problem are related to Choi, Dai and Kim (2018)'s "eventual purchase theorem." That paper shows that, in a consumer search setting with an exogenous set of alternatives, eventual purchasing decisions can be derived through a comparison of the products' "effective values" (see also Armstrong (2017)). We show that such a characterization extends to search problems in which the products are introduced to the consideration set gradually over time, with an order that is outside the consumer's control. This is established by introducing "discovery values" for the products that differ from the "effective values" by accounting for the fact that products need to be first introduced to the consideration set before they can be explored.

Related is also Au and Whitmeyer (2023) who endogenize the collection of prize distributions from which a consumer samples in a model of oligopolistic competition with search frictions, where competing firms with products of unknown quality advertise how much information a consumer's visit will reveal.

2 Model

Consider the problem of a consumer (she) using an online search engine to identify the best product to purchase. The consumer submits a query on the search engine, without knowing which products are advertised at the various positions. She reads the ads in the order they are displayed by the platform but clicks on the ads in the order of her choice. Reading a product's ad is what brings the product to the consumer's consideration set (hereafter CS).

Time is discrete and indexed by $t = 0, 1, 2, \ldots$. At each period before the purchase of one of the products, the consumer must choose among one of the following three actions: (a) reading the next ad displayed by the platform; (b) clicking on one of the products' ads the consumer has read already and that have not been clicked yet; (c) purchasing a product among those whose ads have been clicked already. Clicking on a product's ad (after reading it) directs the consumer to the vendor's website where the consumer learns her value for the vendor's product, net of the product's price. The purchase of a product brings to an end the consumer's search. Importantly, while the consumer naturally reads the ads in the order in which they are displayed, she clicks on the links of the products whose ads have been read in the order of her choice.[2]

As anticipated in the Introduction, this problem can be formalized as a multi-armed bandit problem with a set of arms (equivalently, of products) constructed gradually over time with an order that is beyond the decision maker's control (the sequencing of the products' ads on the platform is outside the consumer's control).

The product that occupies each position $m \in \mathbb{N}$ in the list provided by the platform in response to the consumer's query has a type $\xi \in \Xi$ that is unknown to the consumer before reading the corresponding ad. To each product's type $\xi \in \Xi$ corresponds (a) an inspection cost $\lambda_\xi \in \mathbb{R}$ (this is the cost the consumer incurs, when clicking on the product's link, to learn about her net value $v \in \mathbb{R}$ for the product), and (b) an absolutely continuous cumulative distribution function F_ξ from which the consumer's net value v is drawn.[3] For simplicity, we assume that Ξ is a finite totally ordered set, and then, without loss of generality, let $\Xi = \{1, ..., K\}$.[4] We assume that for any $\xi', \xi'' \in \Xi$, with $\xi' < \xi''$, $\lambda_{\xi'} \leq \lambda_{\xi''}$ and $F_{\xi'}$ (weakly) first-order-stochastically dominates $F_{\xi''}$. In other words, a product of type ξ' is perceived as weakly superior to one of type ξ'', prior to being inspected.

Each position $m \in \mathbb{N}$ is occupied by the ad of one and only one firm, with the same firm possibly advertising at multiple positions. Let $\xi(m)$ denote the type of the product advertised at the m-th position. The consumer expects $\xi(m)$ to be drawn from Ξ according to a probability measure $\rho(m) \in \Delta(\Xi)$, with $\rho_\xi(m) \in [0, 1]$ denoting the probability that the product is of type ξ. We assume that $\rho(m)$ is (weakly) first-order-stochastically dominated by $\rho(m')$, for any $m' > m$. Because higher ξ index lower-quality ads, this assumption means that the consumer expects lower positions to be occupied, on average, by lower-quality products (that is, products that are more costly to learn about, and that deliver, on average, lower

[2]That ads are read in the order they are displayed seems uncontroversial, especially when downstream positions are expected to be occupied by less attractive products delivering, on average, lower net values and whose inspection costs is higher. Furthermore, the assumption is without loss when reading an ad is equivalent to scrolling over it, and the platform displays each downstream ad only after the consumer scrolled on the upstream ones.

[3]The assumption that each F_ξ is absolutely continuous is to avoid the need to keep track of possible indifferences in the consumer's optimal behavior which affect the formulas but not the qualitative results.

[4]The results below do not hinge on Ξ be finite but some of the expressions are less transparent when Ξ is infinite dimensional.

net values).

The assumption that $\rho(m)$ is invariant in the type of ads encountered at upstream positions and that lower positions are expected to be occupied by weakly less attractive products is not crucial to the results but simplifies the analysis. As we show in Section 6, these beliefs are consistent with the firms' equilibrium bidding behavior under an auction format often used in online sponsored search, under appropriate conditions.

Reading the m-th ad reveals to the consumer the type $\xi(m) \in \Xi$ of product advertised at the m-th position and comes with a cost $c(m) \in \mathbb{R}_+$, with $c(\cdot)$ weakly increasing reflecting limiting attention or fatigue.

The consumer cannot click on an ad's link without first reading the ad and cannot purchase a product without first clicking on the corresponding ad's link.

Let $\delta \in \mathbb{R}_+$ denote the discount factor. The consumer maximizes the expected discounted value of the product purchased, net of the discounted sum of the reading and inspection costs. For example, when the consumer clicks on each ad's link immediately after reading the ad (what in the literature is called sequential-cascading behavior) and stops after reading $m \in \mathbb{N}$ ads by purchasing a product with net value v, her ex-post payoff is

$$U = \delta^{2m} v - \sum_{0 \leq s \leq m-1} \delta^{2s} c(s+1) - \sum_{0 \leq s \leq m-1} \delta^{2s+1} \lambda_{\xi(s+1)}.$$

When, instead, the consumer reads $m \in \mathbb{N}$ ads in a row (without clicking on any of them), with $m \geq 2$, then clicks on the m-th ad link, then clicks on the first ad's link, and then purchases the product advertised on the m-th position with a net value of v (an example of what in the literature is called non-sequential-non-cascading behavior), her ex-post payoff is

$$U = \delta^{m+2} v - \sum_{0 \leq s \leq m-1} \delta^{s} c(s+1) - \delta^{m} \lambda_{\xi(m)} - \delta^{m+1} \lambda_{\xi(1)}.$$

More generally, the consumer can follow any strategy of her choice, subject to the constraints described above, namely that ads are read in the order they are presented, that each ad's links can be clicked upon only after reading the corresponding ad, that a product can be purchased only after clicking on the link directing the consumer to the vendor's webpage, and that no more than a single product can be purchased. Given any sequence of decisions satisfying the above restrictions, the ex-post payoff is the discounted sum of the value v of the selected product and of all reading and inspection costs, with the value discounted by the total time $(m+n)$ it took the consumer to finalize the purchase (with m denoting the number of ads read and n the number of products inspected), and with each reading and inspection cost accounting for the precise time $t \leq m+n$ at which the reading or clicking decision occurred.

Note that we are assuming that clicking and reading each take a single period. This assumption is only for convenience. The results extend to a setting in which the two activities take a different amount of time.[5]

Finally, we assume the consumer's outside option is zero.

3 Consumer Optimal Search Policy

As anticipated in the Introduction, the consumer's optimal policy for the search problem described in the previous section is an index rule with the following structure. After discovering that a product is of type ξ, the consumer assigns a "clicking index" $\mathcal{I}^c(\xi)$ to the product, which solves the following equation

$$\mathcal{I}^C(\xi) = \frac{-\lambda_\xi + \delta \int_{\frac{\mathcal{I}^C(\xi)}{1-\delta}}^{\infty} v \mathrm{d}F_\xi(v)}{1 + \frac{\delta}{1-\delta}\left(1 - F_\xi\left(\frac{\mathcal{I}^C(\xi)}{1-\delta}\right)\right)}. \tag{1}$$

This "clicking index" has the same structure as Weitzman (1979)'s "reservation prize."[6] In our problem where the set of products is expanded over time (through the reading of new

[5]See Fershtman and Pavan (2025) for a discussion of how to accommodate for different time lengths for searching for new alternatives and exploring existing ones in a general multi-armed bandit problem with an endogenous set of arms.

[6]Weitzman defines the reservation prize $\hat{\mathcal{I}}^C(\xi)$ for an arm of type ξ as the solution to $\lambda_\xi = \delta \int_{\hat{\mathcal{I}}^C(\xi)}^{\infty}(v - \hat{\mathcal{I}}^C(\xi))\mathrm{d}F_\xi(v) - (1-\delta)\hat{\mathcal{I}}^C(\xi)$, which yields
$\hat{\mathcal{I}}^C(\xi) = [-\lambda_\xi + \delta \int_{\hat{\mathcal{I}}^C(\xi)}^{\infty} v\mathrm{d}F_\xi(v)]/[1 - F_\xi(\hat{\mathcal{I}}^C(\xi))]$. The reservation prizes in (1) are thus equal to those in Weitzman (1979) multiplied by $(1-\delta)$; that is, $\mathcal{I}^C(\xi) = (1-\delta)\hat{\mathcal{I}}^C(\xi)$.

ads), the consumer also assigns a "reading index" $\mathcal{I}^R(m)$ to each ad's position which is given by the solution to the following equation

$$\mathcal{I}^R(m) = \frac{-c(m) + \delta \sum_{\xi \in \Xi(\mathcal{I}^R(m))} \rho_\xi(m) \left(-\lambda_\xi + \delta \int_{\frac{\mathcal{I}^R(m)}{1-\delta}}^{\infty} v \mathrm{d}F_\xi(v) \right)}{1 + \sum_{\xi \in \Xi(\mathcal{I}^R(m))} \rho_\xi(m) \left[\delta + \frac{\delta^2}{1-\delta} \left(1 - F_\xi \left(\frac{\mathcal{I}^R(m)}{1-\delta} \right) \right) \right]}, \tag{2}$$

where, for any $l \in \mathbb{R}$, $\Xi(l) \equiv \left\{ \xi \in \Xi : \ \mathcal{I}^C(\xi) > l \right\}$ denotes the set of ads' types whose clicking index is strictly greater than l.[7]

The following example (whose proof is in the Appendix) illustrates the structure of these indexes:

Example 1. Suppose $\Xi = \{1, 2\}$, with F_1 denoting the cdf of a uniform distribution with support $[0, 1 + \alpha]$, $\alpha > 0$, and F_2 the cdf of a uniform distribution with support $[0, 1]$. The only relevant cost is the discounting time cost: $\lambda_\xi = c(m) = 0$, for $\xi = 1, 2$ and all m. The clicking indexes are then equal to

$$\mathcal{I}^C(1) = (1 + \alpha) \left(\frac{1 - \delta}{\delta} \right) \left(1 - \sqrt{1 - \delta^2} \right)$$

and

$$\mathcal{I}^C(2) = \frac{1 - \delta}{\delta} \left(1 - \sqrt{1 - \delta^2} \right),$$

whereas the reading indexes are equal to

$$\mathcal{I}^R(m) = \frac{(1 - \delta)(1 + \alpha) \left[1 - \sqrt{1 - \frac{\delta^4}{1+\alpha} (1 + \alpha\rho_1(m)) (1 + \alpha\rho_2(m))} \right]}{\delta^2 (1 + \alpha\rho_2(m))}$$

when

$$\frac{\delta \left(1 - \sqrt{1 - \delta^2} \right) (1 + \alpha\rho_2(m))}{1 + \alpha} > 1 - \sqrt{1 - \frac{\delta^4}{1 + \alpha} (1 + \alpha\rho_1(m)) (1 + \alpha\rho_2(2))} \tag{3}$$

[7]See Theorem 1 and Proposition 1 in Fershtman and Pavan (2025), as well as Proposition S.1 in the Supplement of Fershtman and Pavan (2025).

and

$$\mathcal{I}^R(m) = \frac{(1-\delta)(1+\alpha)\left[1-\delta+\delta\rho_1(2)-\sqrt{(1-\delta+\delta\rho_1(2))^2-\delta^4\left(\rho_1(2)\right)^2}\right]}{\delta^2\rho_1(2)}$$

when the inequality in (3) is reversed. □

The consumer's optimal strategy is summarized in the following proposition:

Proposition 1 (***Index policy***)**.** *The consumer's optimal search strategy admits the following index structure. Let $m-1 \in \mathbb{N} \cup \{0\}$ denote the number of ads that have already been read.*

1. ***Clicking***. *If the highest clicking index among the products whose ads have been read but not yet clicked exceeds both (i) the reading index $\mathcal{I}^R(m)$ and (ii) the flow value $(1-\delta)v$ of each product whose ad has already been clicked (as well as the outside option), then the consumer clicks on one of the products with the highest clicking index.*

2. ***Reading***. *If the reading index $\mathcal{I}^R(m)$ exceeds both (i) the clicking indices of all products whose ads have been read but not clicked, and (ii) the flow value $(1-\delta)v$ of each product whose ad has already been clicked, then the consumer reads the m-th ad.*

3. ***Stopping***. *If neither of the above conditions holds, the consumer stops searching. She then selects the product with the highest flow value $(1-\delta)v$ among those whose ads have been clicked, provided it exceeds the outside option; otherwise, she chooses the outside option.*

Using Proposition 1, one can express the consumer's eventual purchase decisions in terms of comparisons of the products' "discovery values," with the latter accounting for the need to read the ads and clicking on them before learning the net payoff the consumer obtains by purchasing the products.

Proposition 2 below extends Armstrong (2017)'s and Choi, Dai and Kim (2018)'s "eventual purchase theorem" to the search problem under consideration, where products are brought gradually to the CS.[8] Let v_m denote the net value to the consumer for the product sold by the firm advertising at the m-th position. For all $m \geq 1$, let

$$w_m \equiv \min\{\mathcal{I}^C(\xi(m)), v_m(1-\delta)\}$$

be the "effective value" of the product advertised at the m-th position (for brevity, product m) when the product's ad has been read already, and

$$d_m \equiv \min\{w_m, \mathcal{I}^R(m)\}$$

the product's "discovery value," before the consumer learns the product's type by reading its ad. Let product $m = 0$ correspond to the consumer's outside option, with $w_0 = d_0 = 0$. Note that w_m and d_m are learned by the consumer only after reading the m-th ad (which reveals its type $\xi(m)$) and clicking on it which reveal its value v_m.

Proposition 2 (**Eventual purchase with endogenous CS**). *The consumer purchases product m if, for all $l \in \mathbb{N} \cup \{0\}$, $l \neq m$, $d_l < d_m$ (and only if $d_l \leq d_m$, for all $l \neq m$).*

As in Choi, Dai and Kim (2018), purchasing decisions are determined by a static condition, as in canonical discrete-choice models. Contrary to Choi, Dai and Kim (2018), however, the relevant values account for the uncertainty the consumer faces over the products occupying the various positions (equivalently, over each product's type $\xi(m)$ which is revealed to the consumer only after reading the product's ad). Allowing for such uncertainty is important. When all products are in the consumer's CS from the outset (i.e., when the consumer is aware of them and only needs to optimize over the sequence of the clicks before finalizing the purchase), positions play no role. In contrast, when the consumer faces uncertainty about the ads displayed at the various positions and chooses how to alternate between reading new ads and clicking on the links of those ads she has already read, the model delivers useful

[8] Greminger (2022) also finds that a certain version of the eventual purchase theorem extends to problems with an endogenous CS. However, none of the results in the next sections have a counterpart in that paper which instead focuses on the comparison between directed and undirected search.

structural relationships linking the positions to the primitives of the problem, as we show in the next section.

The result in the proposition implies that the further down a product is on the list, the lower the ex-ante probability the product is purchased (and hence its ex-ante demand), a property that is often assumed, but not micro-founded, in the models considered in the pertinent literature. Heuristically, if a consumer reads the m-th ad, it must be that the clicking index $\mathcal{I}^C(\xi(l))$ of all products $l < m$ whose ad has been read already, as well as the flow values $v_l(1-\delta)$ of those products $l < m$ whose ad has been clicked upon already, are no greater than $\mathcal{I}^R(m)$. Because the reading index $\mathcal{I}^R(\cdot)$ is non-increasing, $\mathcal{I}^R(m+1) \leq \mathcal{I}^R(m)$. Hence, if after reading the m-th ad, the m-th product clicking index exceeds the same product's reading index (i.e., if $\mathcal{I}^C(\xi(m)) \geq \mathcal{I}^S(m)$), the consumer clicks on the m-th ad right after reading it, thus learning the product's value v_m. Once v_m is learned, if $(1-\delta)v_m \geq \mathcal{I}^C(\xi(m))$, the consumer stops the search and purchases product m. These properties also help explain why purchasing decisions are indeed determined by the comparisons of the discovery values defined above.

4 Click-Through-Rates (CTRs)

Proposition 2 can be used to arrive at a structural formula for the positions' click-through rates (hereafter, CTRs). The latter measure the fraction of ads at each position that, once read, are clicked upon, integrating over a large number of products and consumers. Formally, suppose that the fraction of products of type ξ displayed on the m-th position coincides with the probability $\rho_\xi(m)$ that each such product is of type ξ, and that the fraction of consumers who click on position m-th product coincides with the probability that each such a consumer clicks. For each position m, the corresponding CTR is then equal to

$$CTR(m) \equiv \Pr\left(m\text{'s ad is clicked}\right).$$

This formulation assumes that the probability is computed using only the knowledge of the rules used by the search engine to assign the ads to the various positions (as when such a probability is computed by a platform that does not know the attractiveness of the firms' ads, or by a firm that also lacks such information). In some other cases of interest, these probabilities may be computed using specific information about the type of firms occupying the different positions (as when the above probability is computed by firms knowing each other's type and the assignment induced by a specific bidding outcome in a given auction).

The next proposition relates the CTRs to the effective and discovery values introduced above.

Proposition 3 (Click-through rates). *For each position $m \geq 1$, the click-through-rate is given by*[9]

$$CTR(m) = \Pr\left(\mathcal{I}^R(m) \geq max_{l<m}\{w_l\} \cap \mathcal{I}^C(\xi(m)) \geq \{max_{l<m}\{w_l\}, max_{l>m}\{d_l\}\}\right).$$

To understand the formula, note that, in order for the ad in position m to be read, it must be that $\mathcal{I}^R(m) \geq \max_{l<m}\{w_l\}$, for otherwise the consumer purchases one of the products advertised at one of the upstream positions before reading the ad displayed at the m-th position. Once product m is read, in order for it to be clicked upon, it must be that its clicking index $\mathcal{I}^C(\xi(m))$ exceeds the effective value w_l, $l < m$, of each product advertised at upstream positions, but also the discovery value d_l, $l > m$, of all products advertised further down the list, for otherwise the consumer selects one of the other products before clicking on m. Note that the property that $\mathcal{I}^R(\cdot)$ is weakly decreasing is important. It implies that, if for some position $l > m$, $d_l > \mathcal{I}^C(\xi(m))$, then for all $j = m+1, ..., l$, $\mathcal{I}^R(j) > \mathcal{I}^C(\xi(m))$, meaning that the consumer will necessarily read the ad of any product displayed between position m and position l before clicking on m. If for any of such products the discovery

[9] For simplicity, the formula in the proposition assumes that, in case of indifference, the consumer favors position m (both when it comes to reading and clicking it). This is what justifies the weak inequalities in the formula. The proof discusses how alternative ways of breaking the indifferences must be accounted for if one were to compute bounds for such probabilities.

value d_l exceeds $\mathcal{I}^C(\xi(m))$, the consumer purchases one of these products before clicking on m, and never clicks on the m-th product.

Take the perspective of an observer (e.g., a platform, a firm, or a savvy consumer) knowing the rules used to assign the ads to the positions but not the type of ad occupying each position.

Corollary 1 (**Non-monotonicity of CTRs in ads' positions**). *The probability each ad is read is decreasing in m. However CTRs need not be monotone in positions, consistently with what has been noticed in the empirical literature.*

The result follows from the fact that, while the probability the m-th ad is read – which is given by $\Pr\left(\mathcal{I}^R(m) \geq \max_{l<m}\{w_l\}\right)$ – is decreasing in m, the probability that the ad's clicking index exceeds all downstream ads' discovery values (which is given by $\Pr\left(\mathcal{I}^C(\xi(m)) \geq \max_{l>m}\{d_l\}\right)$) need not be decreasing in m.

5 Non-sequential-non-cascading clicking and externalities from lower positions

We now show how the consumer behavior of Proposition 1 can rationalize the phenomenon of non-sequential-non-cascading clicking documented in empirical work. In their analysis of consumer demand for search advertising, Jeziorski and Segal (2015) document three properties of consumers' behavior that, while ubiquitous in online search, are inconsistent with existing models of search advertising:

1. *Non-sequential clicking*: Nearly half of users do not click on ads sequentially in the order of the positions in which ads are presented.

2. *Non-cascading clicking*: Over half of users who click more than once click on a higher position after having clicked on a lower position.

3. *Externalities from lower positions*: The rate at which an ad at a given position is clicked depends on which ads are displayed below it.

The above properties are consistent with what is predicted by the model of consumer search described above. Perhaps surprisingly, this is so even when the consumer expects lower positions to be occupied by less attractive ads and the cost of reading increases with the number of ads read. Suppose $\mathcal{I}^R(1) \leq \mathcal{I}^C(1)$, which, together with the fact that $\mathcal{I}^R(m)$ is decreasing in m, implies that the reading index is always smaller than the clicking index of any ad whose type is the most attractive one. We then have the following result:

Proposition 4 (**Non-sequential-non-cascading clicking and externalities from lower positions**). *The consumer's search for the optimal product is consistent with non-sequential-non-sequential clicking, and generates externalities from lower positions.*

Importantly, these dynamics do not obtain under traditional search models (see, e.g., Athey and Ellison, 2011). They can emerge in models of consumer search with a fixed set of alternatives (e.g., Weitzman, 1979) if one assumes a specific relationship between the ads' positions and the ads' types. However, in such a model, the consumer's beliefs over the ads displayed at the various positions are degenerate, as the consumer knows (exogenously) which ad occupies each position. Hence, there is no reason why firms would want to bid more for higher positions, in contrast with what is documented in the literature on bidding for sponsored search.

The model of consumer search described above can also generate dynamics under which the probability of non-cascading and non-sequential clicking is non-monotone in the positions. To see this, suppose that $\Xi = \{1, 2\}$. Because $\mathcal{I}^R(\cdot)$ is non-increasing and $\mathcal{I}^R(m) \leq \mathcal{I}^C(1)$ for all m, when $\mathcal{I}^R(1) > \mathcal{I}^C(2)$, there exists a position m^* such that: (a) for any $m < m^*$, the consumer clicks on the m-th ad immediately after reading it if and only if it is of type $\xi = 1$, whereas (b) for any $m > m^*$, the consumer clicks on the m-th ad immediately after reading it, regardless of the ad's type. When, for any $m < m^*$, the probability that the m-th ad is of type $\xi = 1$ is strictly decreasing in m, we then have that the probability of non-sequential and non-cascading clicking is single-peaked (increasing in m for $m < m^*$, and equal to zero for $m > m^*$). In turn, it can be shown that the probability that the m-th ad is occupied by a firm of type 1 is indeed decreasing in m when firms' profits from selling to the consumer are drawn from a distribution that is related to the ads' types by the MLRP and the assignment of the ads is governed by an auction that induces monotone bidding.

6 Implications for equilibrium bidding in sponsored-search auctions

We now show how the results above can be put to work to study equilibrium bidding in sponsored-search auctions. In particular, we show that the auction format often considered in the literature can lead to inefficient assignments even under its ascending-clock implementation.

For simplicity, suppose there are only two firms, with each firm advertising a single product (the results below extend to auctions with more than two positions and more than two bidders). The two firms compete to place their ads on a platform offering two different positions. The platform uses the ascending-clock version of the generalized second price (GSP) auction of Edelman et al. (2007) to allocate the two positions. The firm dropping out first is allocated the second position and pays nothing, whereas the other firm is allocated the first position and pays the price at which the other firm drops out, per click. Let z denote the profit each firm derives from selling its product and assume that z is drawn from $[\underline{z}, \bar{z}]$ according to a distribution G_z, with the draws independent across firms. Finally, assume that the payoff the consumer expects from learning her value for each firm's product (net of the exploration cost λ_ξ and of the cost $c(m)$ of bringing the product to her consideration set) exceeds her outside option, no matter the product's type ξ and the position m at which the product is advertised. Because $c(\cdot)$ is non-decreasing, this amounts to assuming that $\delta^2 \mathbb{E}[v|F_\xi] \geq c(2) + \delta\lambda_\xi$.

As shown in Section 4, the search model introduced above delivers a structural characterization of the CTRs. It can also be used to characterize the purchasing probability for each firm's product, and each firm's value per click (VPC), for each possible profile of firms' types (ξ_1, ξ_2) and each possible assignment of the ads' positions to the firms. Formally, and consistently with the notation introduced above, a type assignment is a vector $(\xi(1), \xi(2))$ where the first entry denotes the type of firm occupying the top position, and the second entry the type of firm occupying the second position. For each position $m = 1, 2$ and each assignment $(\xi(1), \xi(2))$, let

$$P(m; \xi(1), \xi(2)) = \Pr(\text{product } m \text{ purchased under } (\xi(1), \xi(2)))$$

denote the probability that the consumer purchases the product advertised at the m-th position, under the assignment $(\xi(1), \xi(2))$. Clearly, the consumer does not know the assignment $(\xi(1), \xi(2))$ at the beginning of the search process and learns it by reading the various ads.

Let $CTR(m; \xi(1), \xi(2))$ denote position m-s CTR under the assignment $(\xi(1), \xi(2))$; that is, the probability the consumer clicks on the ad displayed at the m-th position when the assignment is $(\xi(1), \xi(2))$. Finally, for any position m, assignment $(\xi(1), \xi(2))$, and unit profit z, let

$$VPC(m; \xi(1), \xi(2), z) \equiv z \frac{P(m; \xi(1), \xi(2))}{CTR(m; \xi(1), \xi(2))} \tag{4}$$

denote the **value-per-click** (VPC) that a firm with unit profit z assigns to placing its ad at position m, under the assignment $(\xi(1), \xi(2))$. Note that, contrary to Edelman et al. (2007) and Athey and Ellison (2011), the values per click here are not only heterogeneous across firms but also position-specific and depend on the quality of other firms' ads, not only those displayed at higher positions but also those at lower positions (a property consistent with what found in the empirical literature; see the discussion above).

Suppose that the two firms commonly know the attractiveness of their ads/products, possibly as a result of past experiences with consumers who searched similar products on the same platform. We then have the following result:

Proposition 5 (Firms' bidding behavior). *Assume that $\delta^2 \mathbb{E}[v|F_\xi] \geq c(2) + \delta\lambda_\xi$ for all $\xi \in \Xi$. Suppose firms know each other's type and do not follow weakly dominated strategies. For any profile of ad types (ξ_1, ξ_2), and each unit profit z, there exists a threshold $b(z; \xi_1, \xi_2)$ such that each firm with unit profit z drops out at price $b(z; \xi_1, \xi_2)$ irrespective of whether its ad is more or less attractive than the rival's.*

To gather intuition, take a type profile (ξ_1, ξ_2) in which firm 1's product is more attractive than firm 2's (in the sense that $\xi_1 < \xi_2$). Consider the interesting case in which the consumer finds it optimal to click on the ad encountered at the first position before reading the ad in the second position, irrespective of whether the ad in the first position is firm 1's or firm 2's. That is, the clicking index $\mathcal{I}^C(\xi(1))$ for the ad displayed in the first position is higher than the reading index $\mathcal{I}^R(2)$ for the ad in the second position, irrespective of the type $\xi(1)$ of ad encountered in the top position (that is, irrespective of whether $\xi(1) = \xi_1$ or $\xi(1) = \xi_2$). When this property does not hold, firms are indifferent between advertising in the first and second position given that the consumer always clicks first the ad of the most attractive firm irrespectively of the position at which the ad is displayed. Firm 1 then finds it optimal to drop out at a price b defined by

$$[VPC(1; \xi_1, \xi_2, z) - b]\, CTR(1; \xi_1, \xi_2) = VPC(2; \xi_2, \xi_1, z) CTR(2; \xi_2, \xi_1). \tag{5}$$

The above indifference condition reflects the fact that both the CTRs and the VPCs are not only position-specific but also ad-specific and depend on the type of ads displayed at all positions (externalities from lower and higher positions).

Equivalently, using the relation between VPCs, CTRs and the selling probabilities in Condition (4) above, we have that the price at which firm 1 drops out is equal to $z[P(1; \xi_1, \xi_2) - P(2; \xi_2, \xi_1)]$. Because the first position is clicked with probability one, the firm optimally drops out when the price reaches a value equal to the extra profit the firm expects from placing its ad in the first position instead of the second one. Note that the term in square brackets is the difference between the probability that the firm sells its product when listed in the top position (that is, under the assignment (ξ_1, ξ_2)) and when listed in the second position (that is, under the assignment (ξ_2, ξ_1)).

Likewise, firm 2, when its unit profit is z, drops out when the price reaches the value b implicitly defined by

$$[VPC(1; \xi_2, \xi_1, z) - b]\, CTR(1; \xi_2, \xi_1) = VPC(2; \xi_1, \xi_2, z) \cdot CTR(2; \xi_1, \xi_2). \tag{6}$$

Note that Condition (6) differs from Condition (5) because the two firms expect different CTRs and have different VPCs for the two positions. Equivalently, using again the relationship in (4), we have that the price at which firm 2 drops out is equal to $z[P(1;\xi_2,\xi_1) - P(2;\xi_1,\xi_2)]$. Clearly, the probability $P(1;\xi_2,\xi_1)$ that firm 2 assigns to selling its good when advertising at the top position is different from the probability $P(1;\xi_1,\xi_2)$ that firm 1 assigns to selling its product when advertising at the same position, reflecting the difference in the distributions from which the consumer's net values are drawn, the exploration costs, and the probability the consumer clicks on the second ad when encountering an ad of type ξ_1 or ξ_2 in the top position. However, the differential in the probability of selling when occupying the first and second position is the same for the two firms; that is,

$$P(1;\xi_2,\xi_1) - P(2;\xi_1,\xi_2) = P(1;\xi_1,\xi_2) - P(2;\xi_2,\xi_1). \tag{7}$$

This follows from the fact that the total probability the consumer purchases from one of the firms depends on the firms' types but is invariant in the assignment of the products to the positions. Under the assumption that the value of exploring each product exceeds the sum of the product's exploration cost, λ_ξ, and of the reading cost $c(m)$, for all ξ all m, this probability is equal to the probability that $\max\{v_1, v_2\} \geq 0$. Such a probability is equal to $1 - F_{\xi_1}(0)F_{\xi_2}(0)$, which naturally depends on the products' types but not on the position at which they are advertised. When this is the case, in equilibrium, the price $b(z;\xi_1,\xi_2)$ at which each firm drops out is the same, despite the fact that one firm is more attractive than the other, and that the firms' CTRs and VPCs depend on the quality of the ads displayed both at higher and lower positions.

Because, for any type profile (ξ_1,ξ_2), the two firms follow identical bidding strategies (they drop out when the price reaches $b(z;\xi_1,\xi_2)$, with $b(\cdot;\xi_1,\xi_2)$ increasing in z), a consumer who understands the rules of the auction should hold beliefs about the type of ad displayed in the second position that are invariant in the type of ad encountered in the first position, as assumed throughout the entire analysis.

The result in Proposition 5 has important implications for the efficiency of the equilibrium allocations under the ascending-clock implementation of the GSP auction considered in the literature. To see this, consider any welfare objective that assigns strictly positive weight to consumer surplus. We then have the following result:

Corollary 2 (Inefficiency of equilibrium ad allocations). *Assume that* $\delta^2\mathbb{E}[v|F_\xi] \geq c(2) + \delta\lambda_\xi$ *for all* $\xi \in \Xi$. *Suppose that firms know each other's type and do not follow weakly dominated strategies. For any type profile* (ξ_1, ξ_2) *for which the attractiveness of the firms' products is not the same across firms (i.e.,* $\xi_1 \neq \xi_2$*), the positions are assigned inefficiently with strictly positive probability.*

Efficiency requires that, in each state in which the attractiveness of the two products is different (that is, $\xi_1 \neq \xi_2$), whenever the difference $|z_1 - z_2|$ between the two firms' unit profits is small, the firm whose product is the most attractive to the consumers be assigned the top position. This is because the top position is clicked more often. Hence, when the first position is occupied by the most attractive firm, the chances the consumer purchases the product she values the most (formally, for which her ex-post net value v is the highest) are higher than when the top position is occupied by the least attractive firm. Hence, no matter the weight the planner assigns to consumer surplus in the welfare objective function, as long as the latter is strictly positive, the auction allocates the positions inefficiently with positive probability. The result is a direct consequence of the fact that, in each state, the two firms follow symmetric bidding strategies, which implies that the assignment of the two positions is based entirely on the firms' unit profits z and not the attractiveness of their products to consumers.

The result in Corollary 2 applies also to settings with more than two firms and more than two positions. To see this, it suffices to note that the situation described above continues to represent the problem each firm faces when there are only two firms left in the auction.[10] The two key properties responsible for the result are that (1) firms possess information about the

[10]The selling probabilities $P(m; \xi(1), \xi(2))$, click-through rates $CTR(m; \xi(1), \xi(2)))$, and values-per-click $VPC(1; (\xi_1, \xi_2, z)$, for positions $m = 1, 2$ should then be interpreted as given the assignment $(\xi(m))_{m>2}$ of the lower positions observed by the two remaining firms.

attractiveness of their products at the bidding stage, and (2) the differential in the selling probability across the top two positions is equalized across the two remaining firms, which is the case when the total probability these firms assign to the buyer purchasing one of their products depends on their respective types as well as the type of products occupying the lower positions, but not on which of the remaining two firms will be assigned the top position.

7 Detrimental effect of additional ad space

The model above can also be used to investigate the effects of additional ad space on firms' profits. Typically, a firm receiving additional ad space expects larger profits. This, however, is not guaranteed when consumers' search for additional products is mediated by online platforms and the type of products advertised on the platform is discovered gradually over time, possibly after visiting other vendors' websites.

To see this, consider the following situation. There are three firms in the market, each producing a product of different type. To facilitate the connection with the model above, we continue to denote a product's type (equivalently, the identify of each firm) by ξ, with $\xi \in \Xi = \{1, 2, 3\}$. However, we drop the assumption that a lower ξ indexes a superior product, which plays no role here.

The consumer's initial CS (before the exploration starts) contains three products, one from each firm. That these products are already in the CS means that the consumer is aware of these products but still needs to visit the vendors' websites to learn her net values for each of these products. The consumer can also submit a query to the platform, which will present her with a second product from one of the same three firms, with ρ_ξ denoting the probability that the product is from firm ξ.[11] The value the consumer obtains from each of firm ξ's products is drawn independently from F_ξ. The consumer must visit the same firm twice to learn her value for each of the firm's products. Each firm makes the same profit

[11] Because a single product is displayed by the platform, we drop 1 from the arguments of the $\rho(1)$ and $\mathcal{I}^R(1)$ functions to ease the notation.

from the sale of each of its products.

We then have the following result:

Proposition 6. *Consider the environment described above. An increase in the probability ρ_ξ the online search query brings a product from firm ξ unambiguously increases firm ξ's ex-ante profits when the consumer's consideration set is exogenous (that is, when the consumer learns about each product's type before starting the exploration process — as in Weitzman, 1979). When, instead, the consumer's consideration set is endogenous (meaning that the consumer chooses when to go online to learn about additional products), an increase in ρ_ξ can reduce firm ξ's ex-ante profits.*

Hence, with an endogenous consideration set, additional ad space may reduce a firm's profits. This is because an increase in the probability with which a query leads to one of firm ξ's products may make the query less attractive (formally, it may reduce the reading index $\mathcal{I}^R$), inducing the consumer to visit one of firm ξ's competitor's websites before submitting the query and learning about the additional product. When strong enough, this effect may reduce the ex-ante probability that the consumer purchases one of firm ξ's products, lowering the firm's profits.

The phenomenon highlighted in the last proposition is not confined to the stylized setting considered here. More generally, when the consumer's consideration set is endogenous, the value of acquiring information about additional products by reading the ads displayed in response to online queries depends critically on the types of products consumers expect a platform to promote. If consumers anticipate that a platform allocates more ad space to lower- or intermediate-quality products, they may prefer to visit the websites of vendors they already know of rather than submitting new queries or reading additional ads. This mechanism is consistent with our search model and helps explain why granting a firm more ad space on a platform may reduce, rather than increase, its profits—even when the space is provided free of charge.

8 Conclusions

We analyze the problem of a unit-demand consumer searching for products on an online platform. The consumer learns the type of products advertised at the various positions gradually over time, after reading the corresponding ads, with the mapping from positions to ads determined by the platform (e.g., through an auction). She decides when to read additional ads and in which order to click on those already read, with clicks directing her to vendors' websites.

We show that this problem can be modeled as a multi-armed bandit one, with an endogenous consideration set, where the consumer alternates between expanding her set by reading new ads and exploring existing options by clicking on them. Within this framework, we characterize position-specific click-through rates and demonstrate that the consumer's optimal strategy rationalizes non-sequential-non-cascading clicking patterns—such as clicking on upstream ads after downstream ones. Click-through rates at a given position depend not only on upstream ads but also on those displayed at lower positions.

We then revisit generalized second-price auctions. Our analysis shows that these mechanisms need not yield efficient allocations, even under the ascending-clock implementation emphasized in prior work. Moreover, the consumer's optimal search strategy implies that expanding the number of ads on a platform may reduce a firm's profits even when the additional ad space is free.

The analysis points to several promising avenues for further research. One direction is to use the search framework to design alternative auction formats for sponsored search. Another is to study the optimal number of ad-positions when the platform seeks to maximize revenue, welfare, or a combination thereof. It would also be interesting to endogenize the information that platforms disclose about assignment rules, the share of sponsored versus organic slots, and the algorithms governing organic search. Finally, the structural formulas for click-through rates and related variables identified in the theory can be used empirically—for example, to estimate firms' willingness to pay for different positions and consumers' beliefs about product attractiveness on search engines.

References

Armstrong, M. (2017). "Ordered consumer search," *Journal of the European Economic Association, 15*(5), 989-1024.

Au, P. H., and M. Whitmeyer (2023). "Attraction versus persuasion: Information provision in search markets." *Journal of Political Economy,* 131(1), 202-245.

Athey, S. and G. Ellison (2011). "Position auctions with consumer search." *The Quarterly Journal of Economics* 126.3, 1213-1270.

Choi, M., Dai, A. Y. and K. Kim (2018). "Consumer search and price competition." *Econometrica*, 86(4), 1257-1281.

Choi, M. and L. Smith (2016). "Optimal sequential search among alternatives." mimeo University of Wisconsin.

Doval, L. (2018). "Whether or not to open Pandora's box." *Journal of Economic Theory* 175, 127–158.

Edelman, B., Ostrovsky, M., and M. Schwarz (2007). "Internet advertising and the generalized second-price auction," *American Economic Review*, 97 (1), 242–259.

Edelman, B., and M. Schwarz (2010). "Optimal auction design and equilibrium selection in sponsored search auctions." *American Economic Review* 100(2), 597–602.

Fershtman, D. and A. Pavan (2025). "Searching for Arms: Experimentation with Endogenous Considerations Sets," mimeo Northwestern University.

Gittins, J. C. (1979). "Bandit processes and dynamic allocation indexes," *Journal of the Royal Statistical Society. Series B (Methodological)* 41 (2), 148–177.

Greminger, R. (2022). "Optimal search and discovery," *Management Science* 68.5, 3904-3924.

Jeziorski, P. and I. Segal (2015). "What makes them click: Empirical analysis of consumer demand for search advertising," *American Economic Journal: Microeconomics*, 7(3), 24-53.

Olszewski, W. and R. Weber (2015). "A more general Pandora rule?," *Journal of Economic Theory* 160, 429–437.

Weitzman, M. (1979). "Optimal search for the best alternative," *Econometrica* 47 (3), 641–654.

Appendix

Proof of Proposition 1. The result follows from Proposition S.1 in the Supplement of Pavan and Fershtman (2025). ∎

Example 1. When F_ξ is the cdf of a uniform distribution with support $[0, k_\xi]$, with $k_1 = \alpha$ and $k_2 = 1$, and $\lambda_\xi = 0$, the equation in (1) reduces to

$$\mathcal{I}^C(\xi) = \frac{\frac{\delta}{2k_\xi(1-\delta)^2}\left[(1-\delta)^2(k_\xi)^2 - (\mathcal{I}^C(\xi))^2\right]}{1 + \frac{\delta}{(1-\delta)^2 k_\xi}\left[(1-\delta)k_\xi - \mathcal{I}^C(\xi)\right]}.$$

Solving for $\mathcal{I}^C(\xi)$, we have that

$$\mathcal{I}^C(\xi) = \left(\frac{1-\delta}{\delta}\right) k_\xi \left(1 \pm \sqrt{1-\delta^2}\right).$$

The relevant solution is the smallest one. This is because the largest solution exceeds the maximal flow value $(1-\delta)k_\xi$ the consumer can obtain by purchasing the product, which is inconsistent with the fact that the index is also equal to the expected discounted payoff per expected discounted amount of time, under an optimal stopping rule; see, e.g., Gittins, 1979).

Next, consider the reading index. Recall that $\mathcal{I}^R(m)$ is given by the solution to equation

(2), which, when $c(m) = \lambda_\xi = 0$, $\xi = 1, 2$, reduces to

$$\mathcal{I}^R(m) = \frac{-\delta^2 \sum_{\xi \in \Xi(\mathcal{I}^R(m))} \rho_\xi(m) \left(\int_{\frac{\mathcal{I}^R(2)}{1-\delta}}^{\infty} v \mathrm{d}F_\xi(v) \right)}{1 + \sum_{\xi \in \Xi(\mathcal{I}^R(m))} \rho_\xi(m) \left[\delta + \frac{\delta^2}{1-\delta} \left(1 - F_\xi \left(\frac{\mathcal{I}^R(m)}{1-\delta} \right) \right) \right]},$$

with $\Xi(\mathcal{I}^R(m)) \equiv \left\{ \xi \in \Xi : \mathcal{I}^C(\xi) > \mathcal{I}^R(m) \right\}$. Because $\mathcal{I}^C(1) > \mathcal{I}^C(2)$, it cannot be that $\Xi(\mathcal{I}^R(m)) = \{2\}$. Furthermore, it cannot be that $\Xi(\mathcal{I}^R(m)) = \emptyset$, for, in this case, the equation has no solution. Hence, there are two cases to consider: (i) $\Xi(\mathcal{I}^S(m))) = \{1, 2\}$ and (ii) $\Xi(\mathcal{I}^S(m)) = \{1\}$. Below, we characterize $\mathcal{I}^R(m)$ for each of these two cases.

Case (i): $\Xi(\mathcal{I}^S(m)) = \{1, 2\}$. In this case,

$$\mathcal{I}^R(m) = \frac{-\delta^2 \left[\rho_1(m) \left(\int_{\frac{\mathcal{I}^R(m)}{1-\delta}}^{\infty} v \mathrm{d}F_1(v) \right) + \rho_2(m) \left(\int_{\frac{\mathcal{I}^R(2)}{1-\delta}}^{\infty} v \mathrm{d}F_2(v) \right) \right]}{1 + \rho_1(2) \left[\delta + \frac{\delta^2}{1-\delta} \left(1 - F_1 \left(\frac{\mathcal{I}^R(m)}{1-\delta} \right) \right) \right] + \rho_2(m) \left[\delta + \frac{\delta^2}{1-\delta} \left(1 - F_2 \left(\frac{\mathcal{I}^R(m)}{1-\delta} \right) \right) \right]},$$

which, after some algebra, can be rewritten as

$$\mathcal{I}^R(m) = \frac{(1-\delta)^2 (1+\alpha) \delta^2 (1 + \alpha\rho_1(m)) - \delta^2 \left(\mathcal{I}^R(m) \right)^2 (1 + \alpha\rho_1(m))}{2(1-\delta)(1+\alpha) - 2\delta^2 (1 + \alpha\rho_1(m)) \mathcal{I}^R(m)}.$$

Solving for $\mathcal{I}^S(m)$, we have that

$$\mathcal{I}^R(m) = \frac{(1+\alpha)(1-\delta)}{\delta^2(1+\alpha\rho_2(m))} \left(1 \pm \sqrt{1 - \delta^4 \frac{(1+\alpha\rho_1(m))(1+\alpha\rho_2(m))}{1+\alpha}} \right).$$

Note that the largest solution exceeds $\mathcal{I}^C(2)$ in contradiction to the fact that we are in case (i) where $\mathcal{I}^S(2) < \mathcal{I}^C(2)$. We conclude that, in case (i),

$$\mathcal{I}^S(m) = \frac{(1+\alpha)(1-\delta)}{\delta^2 (1+\alpha\rho_2(m))} \left(1 - \sqrt{1 - \delta^4 \frac{(1+\alpha\rho_1(m))(1+\alpha\rho_2(m))}{1+\alpha}} \right).$$

Case (ii): $\Xi(\mathcal{I}^S(m)) = \{1\}$. In this case,

$$\mathcal{I}^S(m) = \frac{-\delta^2 \rho_1(m) \left(\int_{\frac{\mathcal{I}^R(2)}{1-\delta}}^{\infty} v \mathrm{d}F_1(v) \right)}{1 + \rho_1(m) \left[\delta + \frac{\delta^2}{1-\delta} \left(1 - F_1 \left(\frac{\mathcal{I}^R(m)}{1-\delta} \right) \right) \right]},$$

which, after some algebra, can be written as

$$\mathcal{I}^R(2) = \frac{\frac{1}{2}\delta^2(1-\delta)^2 \rho_1(2)(1+\alpha) - \frac{\delta^2 \rho_1(2)}{2(1+\alpha)} \left(\mathcal{I}^R(2)\right)^2}{(1-\delta)^2 + \rho_1(2) \left[\delta(1-\delta) - \delta^2 \left(\frac{\mathcal{I}^R(2)}{1+\alpha} \right) \right]}.$$

Solving for $\mathcal{I}^R(m)$, we have that

$$\mathcal{I}^R(m) = \frac{(1-\delta)(1+\alpha)}{\delta^2 \rho_1(m)} \left(1 - \delta + \delta\rho_1(m) \pm \sqrt{\left(1 - \delta + \delta\rho_1(m)\right)^2 - \delta^4 (\rho_1(m))^2} \right).$$

The solution is again the smallest root, since the largest exceeds $(1-\delta)(1+\alpha)$, which is inconsistent with the fact that $\mathcal{I}^R(m) \leq \mathcal{I}^c(1) < (1-\delta)(1+\alpha)$. We conclude that, in case (ii),

$$\mathcal{I}^R(m) = \frac{(1-\delta)(1+\alpha)}{\delta^2 \rho_1(m)} \left(1 - \delta + \delta\rho_1(m) - \sqrt{\left(1 - \delta + \delta\rho_1(m)\right)^2 - \delta^4 (\rho_1(m))^2} \right).$$

Now, to determine which of the two cases applies, observe that case (i) is the relevant case when, and only when, $\mathcal{I}^C(2) > \mathcal{I}^S(m)$, which is the case if and only if condition (3) holds. ■

Proof of Proposition 2. Since product 0 corresponds to the outside option, a product is always purchased. Let $l \neq m$ be such that $d_l < d_m$. We show that product l will not be purchased.

Case 1: $l > m$ (i.e., l is read after m is read). First, suppose that $d_l = \mathcal{I}^R(l)$. Because $\mathcal{I}^R(l) \leq \mathcal{I}^R(m)$ and because $\min\{\mathcal{I}^C(\xi(m)), (1-\delta)v_m\} \geq d_m > \mathcal{I}^R(l)$, under the index policy of Proposition 1, product l is read only after product m is clicked upon. Once m is clicked, however, because $(1-\delta)v_m > \mathcal{I}^R(l)$, l is never read. Hence, l will not be purchased.

Next suppose that $d_l = \mathcal{I}^C(\xi(l))$. Then, $\min\{\mathcal{I}^C(\xi(m)), (1-\delta)v_m\} \geq d_m > \mathcal{I}^C(\xi(l))$. Thus, product l is clicked only after m is clicked. But again, once m is clicked, because $(1-\delta)v_m > \mathcal{I}^C(\xi(l))$, l is never clicked, implying that l is not purchased. Finally, suppose $d_l = (1-\delta)v_l$. Then, because $\min\{\mathcal{I}^C(\xi(m)), (1-\delta)v_m\} \geq d_m > (1-\delta)v_l$, m must be clicked before l is purchased. Because $v_m > v_l$, l is not purchased after m's value is learned.

Case 2: $l < m$ (i.e., l is read before m is read). Because

$$\mathcal{I}^R(m) \geq d_m > d_l \equiv \min\{\mathcal{I}^C(\xi(l)), (1-\delta)v_l, \mathcal{I}^R(l)\},$$

and because $\mathcal{I}^R(m) \leq \mathcal{I}^R(l)$, it must be that $d_l = \min\{\mathcal{I}^C(\xi(l)), (1-\delta)v_l\}$ and hence

$$\min\{\mathcal{I}^C(\xi(l)), (1-\delta)v_l\} < d_m \leq \min\{\mathcal{I}^C(\xi(m)), (1-\delta)v_m\}. \tag{8}$$

Furthermore, because $\mathcal{I}^R(\cdot)$ is non-increasing, $\mathcal{I}^R(l+1) \geq ... \geq \mathcal{I}^R(m-1) \geq \mathcal{I}^R(m)$. Along with the fact that $d_l = \min\{\mathcal{I}^C(\xi(l)), (1-\delta)v_l\} < d_m \leq \mathcal{I}^R(m)$, this implies that $\min\{\mathcal{I}^C(\xi(l)), (1-\delta)v_l\} < \mathcal{I}^R(k)$ for all $(l+1) \leq k \leq m$. This last property in turn implies that either clicking on l, or purchasing l, is dominated by reading any product k, with $(l+1) \leq k \leq m$. If m is read, then (8) implies that l will not be purchased (the arguments are similar to those for case 1). If, instead, m is not read, it must be that another product $k \neq l, m$ is purchased. In either case, product l is not purchased. ■

Proof of Proposition 3. The proof is in two steps. Step 1 shows that $\mathcal{I}^R(m) \geq \max_{l<m}\{w_l\}$ is necessary for product m to be read and that $\mathcal{I}^R(m) > \max_{l<m}\{w_l\}$ implies that product m is read. Step 2 shows that product m is clicked only if

$$\mathcal{I}^R(m) \geq \max_{l<m}\{w_l\} \text{ and } \mathcal{I}^C(\xi(m)) \geq \max\{\max_{l>m}\{d_l\}, \max_{l<m}\{w_l\}\} \tag{9}$$

and that, when both of the above inequalities are strict, product m is necessarily clicked. The result in the proposition then follows directly from the above properties.

Step 1. To see that $\mathcal{I}^R(m) \geq \max_{l<m}\{w_l\}$ is necessary for product m to be read, suppose

that, for some $l < m$, $w_l > \mathcal{I}^R(m)$. That is, both the clicking index $\mathcal{I}^C(\xi(l))$ for any product $l < m$ and product l's flow value $(1-\delta)v_l$ are strictly greater than $\mathcal{I}^R(m)$. Because product l is read before product m is read, m is never read.

Next, we show that, when $\mathcal{I}^R(m) > \max_{l<m}\{w_l\}$, product m is always read. To see this, note that, because $\mathcal{I}^R(\cdot)$ is non-increasing, $\mathcal{I}^R(1) \geq ... \geq \mathcal{I}^R(m-1) \geq \mathcal{I}^R(m)$. Therefore, $\mathcal{I}^R(m) > \max_{l<m}\{w_l\}$ implies that, for any $1 \leq l \leq m$, $\mathcal{I}^R(l) > w_{l-1} = \min\{\mathcal{I}^C(\xi(l-1)), (1-\delta)v_{l-1}\}$. Hence, for any $1 \leq l \leq m$, it cannot be that product $l-1$ is purchased before product l is read. Repeatedly applying this argument for all $1 \leq l \leq m$ implies product m must be read before any product $l < m$ is purchased.

Step 2. To see that both inequalities in (9) must hold for product m to be clicked, first observe that we already established in Step 1 that the first inequality in (9) is necessary for product m to be read. Thus assume that such inequality holds. To see that the second inequality in (9) must also hold, suppose that $\mathcal{I}^C(\xi(m)) < \max\{\max_{l>m}\{d_l\}, \max_{l<m}\{w_l\}\}$. Then either there exists a product $l < m$ such that $w_l > \mathcal{I}^C(\xi(m))$, or a product $l > m$ such that $d_l > \mathcal{I}^C(\xi(m))$, or both. Suppose there is a product $l < m$ such that $w_l > \mathcal{I}^C(\xi(m))$. Then product m cannot be clicked, because product l is necessarily read before m and, because both $\mathcal{I}^C(\xi(l))$ and $(1-\delta)v_l$ are strictly greater than $\mathcal{I}^C(\xi(m))$, product l is purchased before m is clicked. Next, suppose that there exists a product $l > m$ such that $d_l \equiv \min\{\mathcal{I}^R(l), \mathcal{I}^C(\xi(l)), (1-\delta)v_l\} > \mathcal{I}^C(\xi(m))$. By the monotonicity of the search indexes, $\mathcal{I}^R(m) \geq \mathcal{I}^R(m+1) \geq ... \geq \mathcal{I}^R(l)$. That $\mathcal{I}^R(l) > \mathcal{I}^C(\xi(m))$, then implies that $\mathcal{I}^R(k) > \mathcal{I}^C(\xi(m))$ for any $k = m, m+1, ..., l$. In turn, this last property implies that clicking on m is dominated by reading product k, for any $k = m+1, ..., l$. If product l is read, because both $\mathcal{I}^C(\xi(l))$ and $(1-\delta)v_l$ are strictly greater than $\mathcal{I}^C(\xi(m))$, product m is not clicked. If, instead, product l is not read, it must be that another product $k \neq l, m$, with $k \in \{m+1, ..., l-1\}$, is purchased. In either case, product m is not clicked. Hence, the two inequalities in (9) are necessary for product m to be clicked.

Next, we show that when the two inequalities in (9) are strict, product m is necessarily clicked. We already established in Step 1 that, when the first inequality in (9) is strict, product m is read. Now suppose that the second inequality is also strict. That $\mathcal{I}^C(\xi(m)) > \max_{l<m}\{w_l\}$ implies that, for each product $l < m$, either $\mathcal{I}^C(\xi(l))$ or $(1-\delta)v_l$ are strictly smaller than $\mathcal{I}^C(\xi(m))$. Because product m is read, it cannot be that any product $l < m$ is purchased before product m is clicked. Similarly, that $\mathcal{I}^C(\xi(m)) > \max_{l>m}\{d_l\}$ implies that, for each $l > m$, either $\mathcal{I}^R(l)$, or $\mathcal{I}^C(\xi(l))$, or $(1-\delta)v_l$ are strictly smaller than $\mathcal{I}^C(\xi(m))$, which again guarantees that no product $l > m$ can be purchased before product m is clicked. Because one of the products is necessarily purchased (product 0 representing the outside option), it must be that product m is clicked. Hence, we conclude that when the two inequalities in (9) are both strict, product m is necessarily clicked. ■

Proof of Proposition 4. The result follows from the fact that $\mathcal{I}^R(m)$ declines with m and that, for any m, $\mathcal{I}^R(m) \leq \mathcal{I}^C(1)$. These properties imply that, when m is small, after reading the m-th ad and finding that it is of type $\xi > 1$, the consumer may prefer not to click on the ad (if $\mathcal{I}^R(m+1) > \mathcal{I}^C(\xi)$) and instead read the next ad. Later in the process, though, say after reading the m'-th ad, with $m' > m$ such that $\mathcal{I}^R(m'+1) < \mathcal{I}^C(\xi)$, the consumer may start clicking on some of the ads encountered at positions m'', with $m < m'' < m'$. These ads are clicked in the order of their clicking indexes $\mathcal{I}^C(\xi)$. Because, with positive probability, $\xi(n) > \xi(n+j)$ for $m < n < n+j < m'$, the clicking over the ads displayed between positions m and m' can be non-sequential-non-cascading. ■

Proof of Proposition 5. Consider any state of the world in which the attractiveness of the two firms' products differs, (that is, $\xi_1 \neq \xi_2$), and, without loss of generality, assume that firm 1's product is the most attractive one. Note that this property also implies that the two firms' clicking indexes satisfy $\mathcal{I}^C(\xi_1) \geq \mathcal{I}^C(\xi_2)$. As explained in the main text, when $\mathcal{I}^C(\xi_2) \leq \mathcal{I}^R(2)$, each firm drops out instantaneously because it expects the consumer to click on the ad of the most attractive firm first, irrespective of the position at which such an ad is displayed. Thus consider the case in which $\mathcal{I}^C(\xi_2) > \mathcal{I}^R(2)$. That the two firms drop out at the same price then follows from the arguments in the main text after the proposition by observing that, given any assignment $(\xi(1), \xi(2))$, $P(2; \xi(1), \xi(2)) = 1 - P(1; \xi(1), \xi(2))$, which in turn is due to the assumption that the payoff the consumer expects from discovering her value for each firm's product exceeds her outside option. Finally, that the firms' bidding strategies are symmetric in states in which their attractiveness coincides is obvious (and hence the proof is omitted). ■

Proof of Corollary 2. The result follows directly from the arguments in the main text. ■

Proof of Proposition 6. We consider separately the case of an exogenous and endogenous CS.

Exogenous CS. Suppose the consumer learns the type of the product advertised online before starting the exploration of the products. The consumer's optimal policy is the same as in Weitzman, 1979 and consists in inspecting the products in descending order of their clicking indexes, stopping when the remaining indexes are all smaller than the maximal value among the inspected products. Clearly, in this environment, each firm (weakly) benefits from an increase in the probability the consumer's CS contains an additional product from the firm and an equivalent reduction in the probability it contains an additional products from one of its rivals, as both effects contribute to a higher ex-ante probability the consumer purchases one of the firm's products.

Endogenous consideration set. To see that the conclusion above need not extend to an endogenous CS, suppose, for simplicity, that each F_ξ is a Bernoulli distribution assigning probability q_ξ to $v = \hat{v}_\xi$ and $1 - q_\xi$ to $v = 0$, with $\hat{v}_\xi \in \mathbb{R}_{++}$.[12] Further assume the consumer incurs no cost for inspecting any of the products and for reading the ad displayed by the platform (in other words, the only cost is the time cost of postponing the final purchase); that is, $\lambda_\xi = 0$ for $\xi = 1, 2, 3$, and $c = 0$. We then have that each firm's product has a clicking index $\mathcal{I}^C(\xi)$ given by the solution to the equation in (1) which, when applied to the specific structure under consideration, reduces to

$$\mathcal{I}^C(\xi) = \frac{(1-\delta)\delta q_\xi \hat{v}_\xi}{1 - \delta + \delta q_\xi}.$$

[12]One can think of $\hat{v}_\xi$ as the value (net of price) the consumer obtains from the product, provided it matches her needs, which happens with probability q_ξ.

The result in the Proposition then follows from Claim 1 below.

Claim 1. There exist parameter values such that an increase in ρ_2 by ζ, together with a reduction by the same amount in ρ_1, leads to a decrease in the ex-ante probability that one of firm 2's products is sold (and hence in its ex-ante expected profits).

Proof of Claim 1. We establish the claim by showing that such a change may reduce the reading index $\mathcal{I}^R$ inducing the consumer to inspect firm 3's product before reading the ad revealing the type of the forth product. We then show that this effect may imply a drop in firm 2's ex-ante profits.

In this example, the reading index $\mathcal{I}^R$ solving equation (3) is given by

$$\mathcal{I}^R = \delta^2 \max_{k\in\{1,2,3\}} \left\{ \frac{\sum_{\xi\in\{\xi'\in\Xi:\mathcal{I}^C(\xi')\geq\mathcal{I}^C(k)\}} \rho_\xi q_\xi \hat{v}_\xi}{1+\sum_{\xi\in\{\xi'\in\Xi:\mathcal{I}^C(\xi')\geq\mathcal{I}^C(k)\}} \rho_\xi \delta\left(1+\frac{q_\xi\delta}{1-\delta}\right)} \right\}. \tag{10}$$

For concreteness, let $\delta = 0.9$ and suppose that $(\hat{v}_1, q_1) = (10, \frac{1}{10})$, $(\hat{v}_2, q_2) = (3, \frac{1}{3})$, and $(\hat{v}_3, q_3) = (2, \frac{1}{2})$. Note that the corresponding distributions F_ξ are mean-preserving spreads of one another, and that $\mathcal{I}^C(1) > \mathcal{I}^C(2) > \mathcal{I}^C(3)$. Suppose that, initially, $\rho_1 = \rho_2 = \frac{1}{4}$, and $\rho_3 = \frac{1}{2}$. It can be verified that, in this case, the max in the formula for $\mathcal{I}^R$ in (10) is attained for $k = 2$ and that $\mathcal{I}^C(1) \approx 0.473$, $\mathcal{I}^C(2) \approx 0.225$, $\mathcal{I}^C(3) \approx 0.163$, and $\mathcal{I}^R \approx 0.174$, so that $\mathcal{I}^C(3) < \mathcal{I}^R < \mathcal{I}^C(2) < \mathcal{I}^C(1)$. In other words, $\mathcal{I}^R$ does not account for the benefits of bringing firm 3's product to the CS.

Now suppose ρ_2 is increased by $\zeta \in [0, 0.25]$ while ρ_1 is reduced by the same amount. Let $\phi(\zeta)$ denote the ex-ante probability that one of firm 2's products is purchased when the probability that online query brings an additional product from firm 2 is $\rho_2 + \zeta$. Figure 1 depicts the change $\phi(\zeta) - \phi(0)$ in the ex-ante probability that one of firm 2's products is purchased as a function of ζ, where $\phi(0) = (1-q_1)(q_2 + (1-q_2)\rho_2 q_2) = 0.35$. The horizontal gray lines correspond to the clicking indexes $\mathcal{I}^C(1)$, $\mathcal{I}^C(2)$, and $\mathcal{I}^c(3)$, whereas the dark gray curve depicts $\mathcal{I}^R$, with all variables as functions of ζ. Note that $\mathcal{I}^R$ is decreasing in ζ, since $\mathcal{I}^c(1) > \mathcal{I}^c(2)$. Hence, an increase in ζ implies a lower reading index. The reading index $\mathcal{I}^R$ starts out above $\mathcal{I}^C(3)$, and intersects $\mathcal{I}^C(3)$ at an interior ζ (smaller than 0.25),

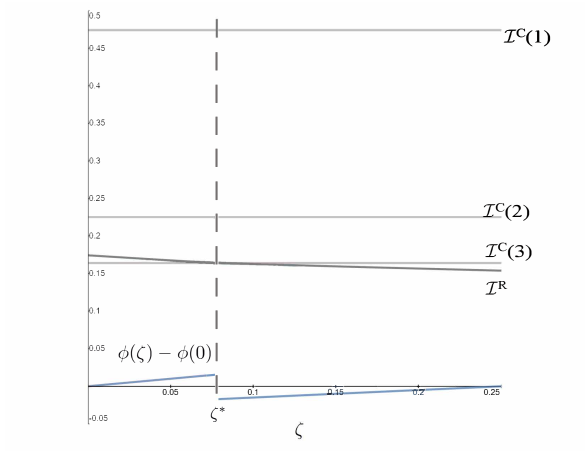

Figure 1: Change in the ex-ante probability the consumer purchases one of firm 2's products, as a function of the increase (by ζ) in the probability the online query reveals an additional product by firm 2: the change in probability is measured by $\phi(\zeta) - \phi(0)$.

denoted by ζ^* (the vertical dashed line). For $\zeta < \zeta^*$, $\mathcal{I}^C(3) < \mathcal{I}^R < \mathcal{I}^C(2)$, whereas for $\zeta > \zeta^*$, $\mathcal{I}^R < \mathcal{I}^C(3)$. The reading index $\mathcal{I}^R$ has a kink at $\zeta = \zeta^*$. For $\zeta \in [0, \zeta^*)$, the consumer reads the ad before inspecting firm 3's product, whereas for $\zeta \in (\zeta^*, 0.25]$ the opposite is true. Therefore, the ex-ante probability that one of firm 2's products is chosen is equal to $\phi(\zeta) = (1 - q_1)\left(q_2 + (1 - q_2)(\rho_2 + \zeta)q_2\right)$ for $\zeta \in [0, \zeta^*)$ and is equal to $\phi(\zeta) = (1 - q_1)\left(q_2 + (1 - q_2)(1 - q_3)(\rho_2 + \zeta)q_2\right)$ for $\zeta \in (\zeta^*, 0.25]$, with a downward discontinuity at $\zeta = \zeta^*$ equal to $(1 - q_1)(1 - q_2)q_2 q_3(\rho_2 + \zeta^*)$. Furthermore, the downward drop in $\phi(\zeta)$ at $\zeta = \zeta^*$ makes $\phi(\zeta) - \phi(0)$ negative over $(\zeta^*, 0.25]$, thus establishing the claim. ■