Search Platforms: Big Data and Sponsored Positions*

Maarten Janssen[†] Thomas Jungbauer[‡] Marcel Preuss[§] Cole Williams[¶]

March 14, 2023

PRELIMINARY DRAFT – PLEASE DO NOT CIRCULATE

Abstract

Search platforms that possess abundant consumer-specific information are ubiquitous in today's economy. We study a search platform's incentives to rank products on their website in response to a consumer query, taking into account how consumers and firms optimally react. Rankings are important as they as they help consumers direct their search efforts and affect firms' sales. In particular, we compare how a platform's optimal ranking depends on its objectives and whether or not they have sponsored slots in addition to organic slots. We describe the platform's incentive to fully obfuscate organic slots, not only to increase its revenue from sponsored positions, but likewise to increase sales commissions. The welfare effect of sponsored positions crucially varies with the platform's other objectives. For example, if the platform maximizes sales commissions, then the consumer is better off with sponsored positions.

1 Introduction

Search platforms such as Google, Tripadvisor or Yelp assist consumers with their online searches for products and services (henceforth firms). A consumer types a keyword into the search engine (e.g., Italian restaurants in Vienna), the platform provides a list of results, and the consumer inspects these results in whatever order she prefers. In producing this list, the platform draws on its information about the consumer (personal information, past searches, order histories, etc.) as well as on choices of other consumers who search for similar keywords. Depending on its objectives, it uses this information to determine which ranking of firms it wants to present to the consumer. Platforms have different ways to monetize the ranking. While they typically sell advertising space and earn commission when facilitating a

*We thank seminar participants at the Workshop on Digital Platforms (Barcelona, December 2022) for their helpful comments.

[†]Department of Economics, University of Vienna. maarten.janssen@univie.ac.at

[‡]SC Johnson Graduate School of Business, Cornell University. jungbauer@cornell.edu

[§]SC Johnson Graduate School of Business, Cornell University. mp2222@cornell.edu

[¶]Department of Economics, University of Vienna. cole.williams@univie.ac.at. Williams acknowledges financial support from the Austrian Science Fund FWF under project number FG 6-G.

sale, platforms also allocate desirable positions, i.e., so-called *sponsored* slots or positions, via auctions.[1]

In this paper, we analyze how the incentives of search platforms, firms and consumers interact by addressing several questions: (i) How does the introduction of sponsored positions affect the informational content of the platform's ranking of search results? (ii) How do sponsored positions affect the social welfare generated by the platform? (iii) How do these incentives to rank items depend on the market power and the data platforms have? And (iv) to which extent does the platform's revenue model determine its optimal choice? Analyzing the role of sponsored slots and how platforms leverage their data does not only improve the understanding of online markets in general, but is of particular relevance for policy makers and regulators of these markets.

Our analysis highlights the crucial and dual role of *obfuscation* in the platforms' strategy, as well as in addressing all the above questions. First, obfuscation of organic slots increases the revenue from sponsored positions. Firms are willing to pay for these sponsored slots only, if it increases their expected sales. Consequently, the platform chooses its algorithm such that consumers do not want to skip the "sponsored positions," i.e., these positions have to be relevant for the consumer's search query. Obfuscating the query relevance of organic slots gives consumers an incentive (i) to first inspect the sponsored slot *and* (ii) to continue searching organic slots less frequently. Second, when maximizing sales commission, a platform obfuscates the information content of organic slots to maximize the number of slots a consumer inspects, (thereby) increasing the likelihood a consumer (eventually) buys.

The model we use to address these questions builds on the seminal papers by Wolinsky (1986) and Anderson and Renault (1999), to which we add a platform that ranks the alternative firms. Using big data techniques, the platform has some understanding of each consumer's preferences over firms, resulting in a *match score* for each firm that is informative of a consumer's *match value*. Additionally, the platform introduces a sponsored slot which it puts up for auction. The platform's strategy then allocates firms to the organic slots and the sponsored position to maximize its revenue. The auction to allocate the sponsored slot uses a combination of bids and match scores.

There is a finite but large number of keyword relevant firms. In our baseline model, these firms are ex ante symmetric and do not have specific knowledge about the consumer's preferences. Each firm submits a bid for the sponsored slot. Given the platform's ranking of sponsored and organic slots, consumers decide how to sequentially search among them. Every time a consumer inspects a firm, she must pay a search cost to learn its match value. If the platform's ranking is informative, consumers may also learn about the likely match values of firms that are not yet inspected while searching.

[1] Throughout the paper we refer to slots or positions that are not sponsored as *organic*.

Our main result is as follows: If the platform wants to maximize the revenue from auctioning the sponsored position, then the platform fully obfuscates the organic positions, but allocates the sponsored position to the firm with the highest match score. In order to see why full obfuscation, i.e., the provision of the least possible amount of information, of the organic slots is optimal when the number of firms is large, consider the winner of the auction for the sponsored slot. Its expected sales under full obfuscation of organic slots necessarily dominates the expected sales under any other ranking of the organic slots. Under full obfuscation, there is minimal learning by consumers once they decide to search beyond the sponsored slot. Moreover, their optimal search pattern among organic slots is random. Any other ranking, by contrast, conveys some information about the match scores of other slots. Therefore, consumers would direct their search among the organic slots toward those firms more likely to generate a higher match value. As a result, the consumer is more likely to continue to search after observing the match value at the sponsored slot relative to full obfuscation. We show that this reduces the value of the sponsored slot, and, thus, causes revenue loss when compared to full obfuscation.

To see why the platform allocates the sponsored position to the firm with the highest match score, note that in the absence of private information, all firms submit identical bids. As a consequence, if the platform commits to assign the firm with the highest match score to the sponsored slot, the sponsored position contains informational value for both consumers *and* firms. Consumers optimally examine the sponsored slot first before inspecting any of the organic ones. Firms, on their part, learn by winning the auction that they are the best possible match for the consumer, increasing the likelihood of the consumer to buy their product/service. In sum, to maximize profits from the auction for its sponsored position, the platform allocates the sponsored slot to the firm with the highest bid, and uses its estimated match value (score) for each firm to break ties. Moreover, it fully obfuscates all organic slots.

One important difficulty that arises is that as the platform's ranking conveys relevant information, consumers update their expectations about the match values they may discover at slots not yet inspected while searching on the platform.[2] It is well-known that learning in consumer search models leads to complications (see, e.g., Garcia and Shelegia (2018)), and that a reservation price may fail to exist (see, e.g., Rothschild (2021) and Janssen et al. (2017)). Thus, to prove our main result, we employ and adapt the so-called mixing property of stochastic processes that has found recent application to the literature on social learning (see Mossel et al. (2020)). The mixing property implies that—due to independence of match values across firms—the match value of the firm in the sponsored slot can be

[2] For example, suppose a platform puts the firm with the highest match score at the sponsored position, and the consumer observes a relatively low match value upon inspection. As a result, the consumer adjust her expectation of the match values at organic slots downwards.

strongly correlated with the match values of at most a few organic firms. Thus, when the number of firms is large, a low realised match value does not reveal much information about the continuation value of search. We indeed show by means of a counterexample that full obfuscation of organic slots may fail to be optimal for a small number of firms.

An implication of our main result is that platforms have an incentive to collect and use increasingly larger amounts of data to predict consumers' preferences. This is because firms find it more valuable to win the sponsored position if the platform can predict matches more accurately. Higher accuracy means that winning the sponsored slot leads to more demand. Importantly, the platform only uses its knowledge to allocate the sponsored position and strategically fully ignores it when allocating the organic slots.

We also show that our main result continues to hold if it is the firms rather than the platform that possess consumer-specific information. In this case, there is an equilibrium in which the firms' bids for the sponsored slot fully reveal their information to the platform. This is because firms with more favourable information are willing to bid more. Again, the platform uses this inferred knowledge to allocate the sponsored slot to the firm with the highest bid (which also is most likely to serve the consumer interest best), but ignores the inferred knowledge when allocating the organic slots. In this way, the platform maximizes the firms' incentives to win the sponsored slot and, as a result, their bids.

To address the welfare effects of introducing sponsored positions, we investigate how the platform will rank products in case it does not offer sponsored slots. We show that welfare effects strongly depend on the platform's alternative objectives. If, in the absence of sponsored positions, the platform cares about its reputation via consumers' ex post utility (so that the platform's and consumer interests are aligned), it would not not want to use uniform obfuscation, even though for a small number of firms it may also not find it optimal to rank the alternatives in the order of its match scores. Consequently, the obfuscation of organic slots that comes with having a sponsored position harms consumers and social welfare for two reasons: (i) the average match value of consumers who buy at the top slot is lower (as consumers are more ready to stop searching even if their match value is not very high), and (ii) consumers who continue to search will spend on average more time inspecting other products.

If, on the other hand, the platform maximizes sales commission, as is common practice, then (perhaps surprisingly) consumers are better off with sponsored positions. With and without a sponsored position, the platform obfuscates the organic slots. When there is a sponsored position, however, the platform allocates it to the firm with the best match score, thus providing the consumer with information. Evaluating these welfare results, it follows that a platforms such as Yelp or Booking.com which charge commission fees even for their organic slots, the sponsored positions help consumers find a good match more quickly.

By contrast, the reverse is true for search platforms such as Google that do not charge commissions.

Starting with Athey and Ellison (2011), Chen and He (2011) and Eliaz and Spiegler (2011) there is a growing literature on position auctions that explicitly take into account that the value of a position depends on how consumers search. A key difference with this literature is that they assume it is firms, not the platform, that have information about consumer preferences related to the keywords they use.[3] Thus, unlike this previous literature, our paper focuses on the important policy question of how online search platforms use the big data they have access to and the role they play in steering consumer search. Recently, Anderson and Renault (2021) develop a setting where firms' prices are endogenized. In the context of a firm directly selling to consumers, and in line with the second role of obfuscation we mentioned above, Nocke and Rey (2023) find that "garbling" of information may be optimal for the firm as it induces a consumer to inspect a larger number of items before terminating search. Unlike many of these papers, our paper does not make the simplifying assumption that a firm-consumer pair either has a match or no match. Another important difference is that we study the interaction between sponsored and organic links, whereas most of the existing literature focuses on one type of slots only.

Ghose et al. (2014), Ursu (2018) and Donnelly et al. (2022) empirically show that personalised rankings affect consumer choices and also have important positive welfare effects, but they only consider organic slots and do not study the effects of search platforms selling their top position. Moreover, they do not consider the incentives platforms have in providing certain types of rankings. Recommendation systems have also recently been studied by e.g. Che and Hörner (2018) and Glazer et al. (2021). Their focus is, however, very different from ours as they focus on how the recommender may choose their recommendations strategically such that it optimally learns from the choices different agents make. In addition, they study settings where agents value goods identically and there are no firms whose interests should also be taken into account. Armstrong and Zhou (2022) also consider how information provision affects consumers, but they do so in a setting where consumers cannot actively search.

Our paper is also related to Janssen and Williams (2022) who study how the recommendation of a social influencer affects market outcomes through its effect on the consumer search order and where the preferences of the influencer are correlated with those of their followers. An influencer, however, does not have to provide a complete ranking and typically only recommends one product. Our current paper focuses on the interaction between the choice of firm for the sponsored slot and the complete ranking of organic slots.

[3] A notable exception is a recent paper by Ke et al. (2022) that also considers the information platforms have about consumer preferences.

The rest of the paper is organized as follows. The next section describes the model in more technical detail, while apart from our main result Section 3 also shows how it extends to firms having private information or a small number of firms. Section 4 contains the analysis of platform rankings under alternatives objectives. We conclude with a discussion.

2 Model

The market comprises n firms, a representative consumer, and a platform. The platform's information regarding the consumer and firms is summarized by a real vector of **scores** $\theta = (\theta_1, \ldots, \theta_n) \in \Theta$ where θ_i is a signal of the consumer's match value with firm i. Scores are independent and identically distributed (IID) across firms according to a distribution $F(\theta_i)$ supported on the compact interval $[\underline{\theta}, \bar{\theta}]$. The consumer demands one unit of the good and has match values $v = (v_1, \ldots, v_n)$ where v_i is the consumer's value for the good sold by firm i. Match values are independent across firms with the match value with firm i drawn from the distribution $G(v_i|\theta_i)$. We assume that each distribution $G(\cdot|\theta_i)$ shares a common support $[\underline{v}, \bar{v}] \subset \mathbb{R}$ which may be unbounded. Higher scores indicate "good news" in the sense that $G(\cdot|\theta_i)$ has likelihood ratio dominance over $G(\cdot|\theta_i')$ if $\theta_i \geq \theta_i'$. We also include an independent nonatomic random variable z with support Z for the platform to use as a randomization device to play a mixed strategy. The probability measure on $\Theta \times V \times Z$ is denoted by $\mathbb{P}[\cdot]$ and expectations taken with respect to this measure are denoted by $\mathbb{E}[\cdot]$.

The platform displays a ranking of the firms to the consumer. It does so by running an auction in which firms submit bids, the winner of the auction is placed at the top in the "sponsored" position, and all other firms are arranged into the remaining "organic" positions. Denote firm i's bid by $b_i \geq 0$ and the vector of bids submitted to the platform by $b = (b_1, \ldots, b_n) \in B$. Given the bids and scores, the platform's algorithm determines which firm wins the sponsored position and how the remaining firms are arranged in the organic positions. We denote the set of firm permutations by X,[4] so that the platform's **algorithm** is a function $a : B \times \Theta \times Z \to X$. We let $\mathcal{A}_n$ be the set of algorithms in the game with n firms for which $a \in \mathcal{A}_n$ implies that $a(b, \cdot)$ is measurable for every bid vector.

The platform allocates the sponsored position according to a **weighted second-price auction**. The term "weighted" refers to the fact that the winner of the auction can be influenced by the scores as is implicit in the definition of an algorithm. If the winner of the auction submitted the highest bid, then it pays the platform the value of the second-highest bid. If instead the winner did not submit the highest bid, then it pays the value of its own bid to the platform.

Consumers are initially uninformed of their match values with firms and can only find

[4]That is, X is the set of bijections from $\{1, \ldots, n\}$ to itself.

this out through costly sequential search. Consumers face a non-negligible inspection cost $s > 0$ to learn the match value with an individual firm. In principle, the consumer search problem is Pandora's box problem as studied in Weitzman (1979). However, knowledge that the platform utilizes a ranking algorithm creates interdependence between match values so that inspecting the goods of one firm provides the consumer with information about other firms. Consumers have perfect recall when searching.

The timing of the interaction is as follows. The platform begins by committing to an algorithm which is observed by all players. Firms privately submit their bids to the platform. Nature determines the scores and match values. The platform receives the firms' scores and bids and the algorithm determines the position each firm takes in the list. The consumer receives the list and then proceeds with their search.

The consumer's payoff is equal to the match value minus the price of a good it purchases net the search costs. A firm's profit equals the revenue minus product cost and any fee paid for the sponsored position. The platform's expected profit corresponds to the expected revenue from the sponsored search auction. We focus on symmetric Perfect Bayesian Equilibria where the platform ranks the products to maximize expected profits, consumers choose an optimal sequential search strategy. In terms of strategies and strategy spaces, the firms' and consumers' strategies are standard from the auction and search literatures.

We end this section with a few comments on the model. First, we assume that the platform commits to its ranking algorithm. We see this as a reasonable approximation of the real world situation where platforms submit a ranking of alternatives within a split second after the consumer has types its key words. Consumers typically frequently use one and the same platform over and over and see the resulting rankings and how they satisfy their needs. Platforms may, of course, work on different algorithms to improve their functioning, but will implement new algorithms only once in a while.

Second, we implicitly assume that the platform has to rank all firms and that it cannot present a truncated ranking. We see this as a short-hand for a platform not having unlimited market power. In particular, if consumers would not continue to use the platform if they see that a firm they expected on the list is not ranked at all, then platforms may prefer to rank all firms and not run the risk of consumers not using their services in the future.

Third, the model treats firms' prices as exogenous and in particular that they do not depend on whether or not a firm is recommended. We think that this is realistic in many cases where the revenue a firm makes is only to a limited extend dependent on the sales via the search platform. Implicitly, we also assume that all firms charge identical prices, but that turns out to be inessential as we will explain in the next section after stating our main result.

3 Maximizing Revenues from the Sponsored Slot

in this section we consider that the platform has a sponsored slot and is interested in maximizing the revenue it gets from it, as explained in the previous section. our main result is that as the number of firms grows large, it is optimal to fully obfuscate the organic slots so that they dot not contain any information regarding the platform's match scores. In addition, we show that the platform allocates the sponsored slot to the firm with the highest match score.

We first introduce two definitions. Let $\Pi(a)$ denote the platform's expected profit when the consumer and firms play an equilibrium of the subgame following the selection of a. In general, an algorithm could induce multiple equilibria that differ in the amounts firms bid for the sponsored slots and in that case $\Pi(a)$ simply selects the pay-off of an arbitrary equilibrium in that set. In what follows, a sequence of algorithms $\{a_n\}_{n\in\mathbb{N}}$ assumes $a_n \in \mathcal{A}_n$ for all $n \in \mathbb{N}$.

Definition 1. *A sequence of algorithms $\{a_n\}_{n\in\mathbb{N}}$ is **asymptotically optimal** if for every sequence $\{a'_n\}_{n\in\mathbb{N}}$ and $\epsilon > 0$ there exists an n^* such that $n \geq n^*$ implies $\Pi(a_n) + \epsilon > \Pi(a'_n)$.*

Definition 2. *An algorithm $a \in \mathcal{A}_n$ exhibits **uniform obfuscation** if the firms that lose the auction are assigned to the organic positions with uniform probability.*

Using these definitions, we can now state our main result.

Theorem 1. *There is a sequence of uniformly obfuscating algorithms that is asymptotically optimal. The optimal algorithm allocates the sponsored slot to the firm with the highest match score.*

It is clear that if the platform obfuscates the match scores of firms at the organic slots and allocates the sponsored slot to the firm with the highest match score, consumers will start their search at the sponsored slot. from the literature on prominence in search markets (see, e.g., Armstrong et al. 2009) by itself creates a value to firms of occupying the sponsored slot. This value gets larger the less attractive it is for consumers to continue searching after having observed the match value of the firm at the sponsored slot. Obfuscating the organic slots has the additional advantage to the platform in that it becomes very unattractive to firms not to win the sponsored slot as this implies they end up on a completely random organic slot. Finally, as firms do not have private information on which to condition their bid, the platform receives identical bids from all firms and allocates the sponsored slot to the firm with the highest match score. This creates an additional informational value of winning the sponsored slot: firms know that if they are chosen they have the highest match score making it even more likely that the consumer ends up buying from them.

The informational value of winning the sponsored slot creates multiple equilibria: given the symmetric equilibrium bids, firms do not want to outbid their competitors as outbidding implies that they win the auction because of their higher bid, and not because of their higher match score. As a consequence, the informational value of winning the auction gets lost. the platform can, however, easily resolves this issue of equilibrium multiplicity as it could set a reserve price for auction bids that is equal to (or close to) the highest bid equilibrium.

The argument about equilibrium multiplicity also shows that the result stated in the theorem continues to hold if different from what we have assumed so far) firms have different margins, for example because they set different prices. Without the informational value to firms of winning the auction, firms with higher margins nay have more incentives to be ranked first (as this will increase the probability they sell) and are therefore willing to bid more. However, in our context and starting from an equilibrium with equal bids, firms with higher margins may not want to bid more as this will remove the informational value from winning the auction.

The conclusion of Theorem 1 abstracts from the presence of competitors the platform faces. In practice, a consumer typically has the ability to leave the platform and continue searching elsewhere if they become dissatisfied. Interestingly, as long as the outside option does not preclude the consumer from visiting the platform for every algorithm, because the algorithm detailed in Theorem 1 places the best firm in the sponsored position, it continues to lure the consumers to first examine the sponsored firm even if they then continue their search elsewhere. Formally, we refer to the outside option $\eta \in \mathbb{R}$ as the payoff received by the consumer if they decide to exit the platform at any point.[5]

Corollary 1. *The sequence of uniformly obfuscating algorithms identified in Theorem 1 remain asymptotically optimal for every value of the outside option* $\eta \in \mathbb{R}$.

Given the above explanation, the reader may get the impression that the result should also hold when n is finite. The statement of the theorem is, however, about uniform obfuscation only being asymptotically optimal and as the extensive counterexample in the next subsection shows it may not hold when n is finite. With finite n learning may affect the optimal consumer search rules and to prove the result in the appendix we therefore had to rely on results for social learning in general, and the mixing property in particular.

[5]The proof of the Corollary 1 follows from the exactly the same argument as the argument for Theorem 1 where we replace the upper bound on sales by the sponsored firm with $1 - G(\max\{\bar{r}, \eta\}|\bar{\theta})$.

3.1 Uniform Obfuscation may be Suboptimal with a small Number of Firms

In this subsection, we discuss an example illustrating how uniform obfuscation can fail to be optimal when there are a small number of firms, because of non-monotonicities in the inference consumers draw from observing their match value at the sponsored slot. Many search queries result in many firms being ranked and in that case our main theorem applies. the current subsection serves the purpose of showing why the result may fail to hold when only a small number of firms are ranked.

Suppose there are three firms $i = 1, 2, 3$. The consumer's match value is either low ℓ, medium m, or high h and a good is only worth purchasing if it provides at least a medium value. A firm's match score is L when the value is low and H when the value is either medium or high, i.e., the platform can distinguish firms with low match scores from other firms, but cannot distinguish firms with medium and high match scores. Let p_H and p_L denote the marginal probability that a firm's score is high and low, respectively.[6]

Suppose the platform employs the following algorithm. The firm with the highest bid is placed in sponsored positions, ties are broken in favor of the firm with the highest match score, further ties are broken with equal probability. For the two nonsponsored firms, if only one of them has a high signal it is placed in the second position with probability $\alpha \geq \frac{1}{2}$, otherwise they are arranged in the organic positions with equal probability.

Given the algorithm, the consumer's optimal search proceeds in the following manner. If the sponsored firm's value is high h, then the consumer buys it immediately since there is no advantage from continuing. If instead the sponsored firm's value is low ℓ, then given the algorithm, the consumer learns that all remaining firms must likewise have low match values and so the consumer might as well exit the market. If, however, the consumer observes m in the sponsored slot, then it might still be prudent to continue searching as some remaining firm might deliver a higher match value. To describe the consumer's learning over the course of search, let subscripts denote the index of the position so that the list of possible events are $\{(H_1, H_2, H_3), (H_1, L_2, H_3), (H_1, H_2, L_3), (H_1, L_2, L_3)(L_1, L_2, L_3)\}$ which occur with corresponding probabilities $\{p_H^3, 3(1-\alpha)p_H^2 p_L, 3\alpha p_H^2 p_L, 3p_H p_L^2, p_L^3\}$. The probability that slot two has a high score given that the first does is

$$\mathbb{P}(H_2|H_1;\alpha) = \frac{p_H^3 + 3\alpha p_H^2 p_L}{1 - p_L^3}. \tag{1}$$

[6]This example departs from the assumptions of our model in that the distribution of match values conditional on the match scores do not share the same support. This is insignificant to the particular example since we could modify the distributions to $\mathbb{P}(\{\ell\}|L) = \mathbb{P}(\{m, h\}|H) = 1 - \varepsilon$ so that the conclusion continues to hold for $\varepsilon > 0$ sufficiently small.

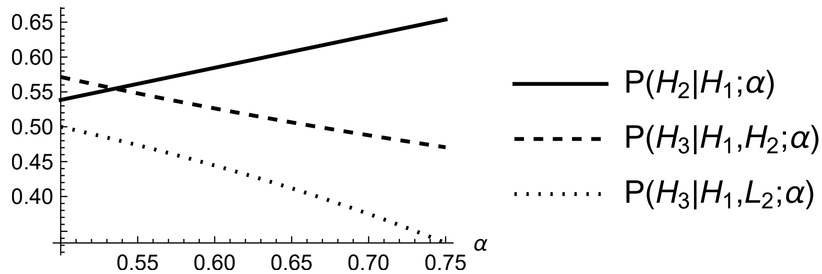

Figure 1: The figure plots the conditional probabilities as a function of α given that match values each occur with equal probability.

The probability that slot three has a high score given that the first two also do is

$$\mathbb{P}(H_3|H_1, H_2; \alpha) = \frac{p_H^3}{p_H^3 + 3\alpha p_H^2 p_L}. \tag{2}$$

The probability that slot three has a high score given that the first does and the second has a low score is

$$\mathbb{P}(H_3|H_1, L_2; \alpha) = \frac{3(1-\alpha)p_H^2 p_L}{3(1-\alpha)p_H^2 p_L + 3p_H p_L^2}. \tag{3}$$

Inspecting the above expressions, we find that (2)>(3) which naturally implies that the consumer is more optimistic about the third slot upon observing a medium match value in the first two slots than if she were to observe a medium in the first and a low in the second. Also, (1)>(3) holds true, implying the consumer is more optimistic about the second slot after observing a medium in the first than she is about the third slot upon observing a medium and low value in the first two. The comparison between (1) and (2) depends on α. As Figure 1 illustrates, increasing α makes proceeding to the second firm more attractive, but continuing to the third less so.

Suppose the parameters are such that, under uniform obfuscation (i.e. when $\alpha = 1/2$), the consumer continues searching when observing m in the first firm and m in the second firm, but halts otherwise. To be concrete, assume that the parameters satisfy $\frac{1}{2}\mathbb{P}(H_2|H_1; \alpha = \frac{1}{2})(h - m) = s$ which leads the consumer to follow the desired search pattern. We compare uniform obfuscation against a nonuniformly obfuscating algorithm with $\alpha = \alpha^* > \frac{1}{2}$ whereby α^* is large enough to ensure that a consumer will inspect the second slot if the first provides a medium match value, but will search no further. For example, in Figure 1, given our assumptions on parameters, setting $\alpha = 0.6$ guarantees that $\mathbb{P}(H_3|H_1, H_2; \alpha^*) < \mathbb{P}(H_2|H_1; \frac{1}{2})$ and thus it is never optimal for the consumer to inspect the third firm. Notice that by providing some information in the organic slots, the nonuniformly obfuscating algorithm makes inspecting the second firm more desirable, but increases a sponsored firm's return

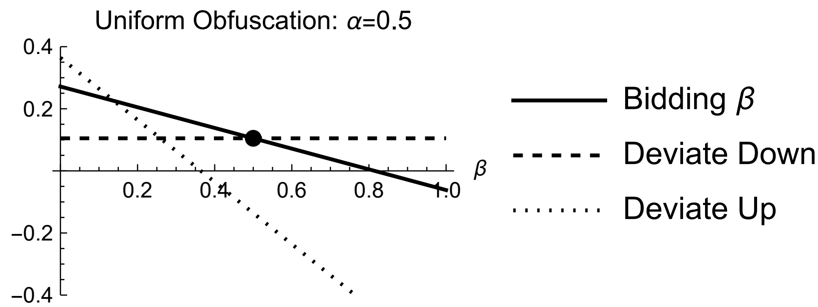

(a)

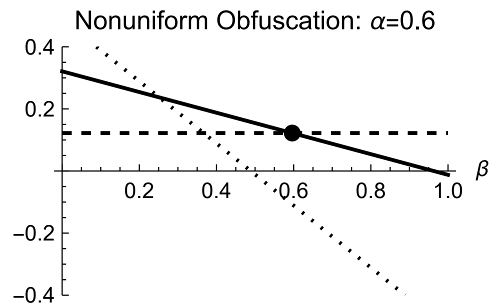

(b)

Figure 2: The figures plot an individual firm's expected profit in a tentative symmetric equilibrium in which all firms bid β under the respective algorithms from likewise bidding β, deviating to a lower bid, and deviating to a higher bid (given that match values each occur with equal probability).

demand as consumers will not inspect the third firm.

For each algorithm, consider a symmetric equilibrium in which each firm bids β. Figure 2 plots the tentative expected equilibrium and deviation profits under the two proposed algorithms for different values of the bid β. Naturally, the expected profit from playing the tentative equilibrium strategy is decreasing in the bid with a slope of $-\frac{1}{3}$ as each of the three firms win with equal probability. Placing a higher bid, no matter how high, ensures that a firm wins the auction, but also makes winning uninformative as it does not provide information about the firm's match score. The slope of the expected profit given an upward deviation is thus -1. Thus, this deviation risks winning in case the firm has a medium match score and consumers continue searching after visiting the sponsored slot. The deviation is optimal for low values of the candidate equilibrium bid, but not for higher values. Offering a bid less than β ensures that a firm does not win the sponsored slot; hence, the deviation profit is simply the expected profit in an organic slot and does not depend on the bid.

We can use Figure 2 to compare the profitability of the two algorithms by comparing the range of bids firms are willing to place. The figure shows that, in principle, there can be a continuum of equilibrium bids. However, we only need to compare the equilibria with the highest bids as the platform can secure a profit equal to this bid by setting a reserve price equal to the highest possible equilibrium bid. As illustrated by Figure 2, using a uniform obfuscation algorithm, the platform can achieve a profit of exactly $\frac{1}{2}$, which is identified by finding the highest value for the bid at which no firm wishes to deviate from placing that bid. On the other hand, using the nonuniformly obfuscating algorithm, the platform can secure a profit of approximately 0.6. Thus, due to the learning from match values at the sponsored slots, the platform is better off choosing the nonuniformly obfuscating algorithm.

3.2 Improved Platform Information

Now that we better understand for which purposes the platform may (not) use the information it possesses, we can also answer the question how the market is affected by the platform having more accurate information.

So far, we have written the joint density over firm i's match score and value as the product of the marginal score density and the conditional value density: $g(v_i|\theta_i)f(\theta_i)$. To model improved information for the platform, it is helpful to rewrite the joint density as $g(v_i)f(\theta_i|v_i)$ where $g(v_i)$ is the marginal value density and $f(\theta_i|v_i)$ is the conditional score density capturing the quality of the platform's information. Let $F(\theta_i|v_i)$ denote the corresponding conditional score distribution.

Consider an improvement in the quality of the platform's information in the sense of Lehmann (1988). That is, if $F(\theta_i|v_i)$ is the platform's initial score distribution, then the platform's new score distribution $\tilde{F}(\theta_i|v_i)$ is such that

$$\tilde{F}^{-1}\left(F(\theta_i|v_i)|v_i\right)$$

is nondecreasing in v_i for all θ_i.[7]

Proposition 1. *An improvement in the quality of the platform's information yields higher limiting profits under uniform obfuscation.*

Consequently, even though the platform's ranking is uninformative beyond the sponsored slot, it still has an incentive to acquire information even if doing so is costly. Firms value winning the sponsored slot more, the more certain they can be that the consumer will buy from them conditional on winning. If match scores were barely informative about match values, the chance that the consumer buys from the firm placed in the sponsored slot is barely greater than the chance that she would buy from any randomly inspected firm. But the more accurate the information about match values that the platform uses to decide

[7] Dewatripont et al. (1999) and Persico (2000) discuss economic applications of Lehmann information.

which firm is placed in the sponsored slot, the higher is the probability that consumers buy from the sponsored slot and, thus, the higher is the firms' willingness to pay for this event.

This is consistent with the interpretation of search platforms being a critical gatekeeper in online markets. Basically, the platform sells "preferred access" to consumers (the probability that the consumer inspects any given firm is very small if n is large, unless a firm wins the sponsored slot) to firms. This access is, of course, more valuable to a firm if the likelihood that the consumer likes the firm's product is higher. This is why more accurate information increases the platform's profits.

3.3 Privately Informed Firms

We have so far assumed that only the platform has access to relevant information regarding consumer preferences, and firms do not possess such information. We think this is relevant in many instances where firms are relatively small and do not have the relevant technology in place to digest large amounts of information. However, there are other instances where firms also do have relevant information that the platform does not have access to. In this subsection, we therefore consider the reverse situation, namely that firms have superior information about how well their product fits a particular search query.

To model those instances, we assume that the match score θ_i is firm i's private information. As in the base model, scores are independent and identically distributed (IID) across firms according to a compactly supported distribution $F(\theta_i)$. Unlike in the base model, however, the platform does not know the vector of all match scores $\theta = (\theta_1, \ldots, \theta_n)$.

We show that a similar algorithm as we used so far is asymptotically optimal ex post: The sponsored slot is assigned to the highest bidder who pays the second highest bid; and the platform uniformly obfuscates the organic slots. The only difference is that now firms place different bids, depending on their private information, and that the platform cannot discriminate between identical bids as it does not have information about match scores.

Thus, the gist of Theorem 1 continues to hold if firms know their match scores and the platform does not. As the equilibrium will be separating, the platform can infer match scores from bids. However, the platform consciously decides not to use this information to allocate the organic slots in order to boost the bids on the sponsored slots firm make.

Proposition 2. *When firms are privately informed of their match scores, there is a sequence of uniformly obfuscating algorithms that is asymptotically optimal for the platform.*

It is not difficult to see that for the same match scores and when n is large, the platform earns the same profits, whether it knows the match scores ex ante or not. If firms do not know their match scores, they know that the platform uses match scores to allocate the sponsored slots and take this already into account when making the bid. If, on the other hand, firms know the match score they realize that the firm with the highest match score will win the auction and that the platform ignores this information when allocating organic slots. Thus, for a given match score of the winner, the sales will be identical. Finally, if n is large the expected value of the second highest bid becomes arbitrarily close to the match value of the winning firm. In other words, auctioning a sponsored slot may have an additional advantage of having firms reveal their private information.

4 Welfare effects of Sponsored Slots

Sponsored positions are ubiquitous in online markets, but a more recent phenomenon, and their effects are, as a result, not yet well understood. In this section, we aim to shed light on the welfare consequences of sponsored positions. In practice, platforms have typically been and continue to be compensated for their intermediary services in different ways. As a result, many platforms face a challenging problem, that is, how to present search results to the consumer when the revenue stream is multi-dimensional. Besides selling sponsored positions, platforms typically generate revenue from sales commission, i.e., payments from firms in the case of a transaction. Another common revenue stream for platforms is advertising.

In this section, we account for this complex reality, and ask how the introduction of a sponsored slot affects the consumer if multiple objectives determine the platform's decision making. In fact, we consider two potential alternative objectives, and analyze the effects of introducing a sponsored slot for each. First, a platform may directly care about consumer welfare. This is the case, for example, if unsatisfied consumers discontinue using the platform and the platform's ad revenue is directly related to its number of consumers. Second, a platform may accrue revenue through sales commissions it receives when consumers buy products from firms.

The following two subsections deal with the above laid out alternative platform objectives in turn.

4.1 Maximizing Consumer Welfare

We first focus on the case where in the absence of sponsored positions, the platform's and consumers' interests are aligned. Common sense may suggest that the platform will always choose a perfect ranking. This intuition, however, is wrong if the number of firms n is small and we first present a counterexample to make this point.

Example. Building on our earlier example is Section 3.1, continue to assume that match values are low, medium, or high: $L < M < H$. Now assume that the platform observes when the match value is low for certain, and otherwise receives a noisy signal of the match values. In particular, there are three possible match scores $\{\theta_L, \theta_M, \theta_H\}$ with $\mathbb{P}(L|\theta_L) = 1$, $\mathbb{P}(H|\theta_H) = \mathbb{P}(M|\theta_M) = q > \frac{1}{2}$, and $\mathbb{P}(H|\theta_M) = \mathbb{P}(M|\theta_H) = 1 - q$.

Suppose the consumer is only willing to buy a high-valued item but also, if they knew the match scores, would only be willing to inspect firms with high scores: $r_H > \eta > r_M > r_L$. With this, there exists some number k such that searching one firm with a low score and then being able to turn to the outside option would be preferable to searching k firms with a medium score before being able to turn to the outside option.

Under a perfect ranking, by taking the prior probability of a high match value to be large enough, we can ensure that optimal search involves continuing on after observing arbitrarily many medium values. In particular, we can ensure that the consumer searches at least k times after observing medium scores.

Hence, given these parameters, suppose the platform performs a perfect ranking, except that in the event that no firms have high scores, k or more firms have medium scores, and at least one firm has a low score, the platform places a firm with a low score in the first position. Then, the consumer's best reply is to always halt search after observing a low score. But this algorithm yields a higher expected payoff to the consumer.

The point of the counterexample is that a perfect ranking may not give the consumer the best information possible of whether or not it is optimal for the consumers to continue to search. Essentially, if the platform has information that is contained in the match scores that is not perfectly transmitted to the consumer via a ranking, then providing a perfect ranking may be suboptimal even if preferences are aligned.

There are two results that do hold, however, for a platform whose objective it is to maximize consumer welfare. First, for every n uniformly obfuscating the slots is not the optimal ranking for the platform. This essentially follows from Blackwell informativeness of an algorithm. As a consumer can obtain the outcome of a more informative algorithm by randomizing in a specific way over the outcome of another algorithm, potentially disregarding information, the consumer can always make herself at least as well off under a more informative algorithm than under uniform obfuscation. Second, as n grows large, the concern expressed in the counterexample that a consumer may search for too long diminishes as the chance that there are not enough firms that are worth inspecting becomes small. It follows that consumers have an incentive to follow a perfect ranking if n is large.

Proposition 3. *If a platform allocates firms to slots to maximize consumer surplus, then it is asymptotically optimal to choose a perfect ranking.*

Thus, when n is large a perfect ranking delivers two advantages to consumers. First, the expected number of searches needed to find a satisfactory product is smaller than under a less informative ranking. Second, consumers more quickly learn when it is optimal to abort their search altogether as the chance they find a valuable object diminishes. Corollary 2 compares the consumer-optimal ranking with the main result of our base model.

Corollary 2. *If, in the absence of a sponsored slot, the platform maximizes consumer welfare, then introducing a sponsored slot reduces consumers welfare if n is large.*

Note that Corollary 2 directly follows from the remarks on Blackwell informativeness. What is interesting is that the change in objectives does not affect the allocation of the top slot. Under both objectives this slot is allocated to the firm with the highest match score. It is rather the allocation of the remainder of firms across organic slots that lowers consumer welfare when the platform maximizes its auction revenue. Moreover, the platform's optimal ranking is far from obvious if its objective considers consumer welfare and auction revenue at the same time. Nevertheless, it remains true that the consumer is better off without a sponsored slot.

4.2 Revenues from Sales Commissions

We now amend the base model and analyze how the platform allocates firms across slots if it earns commission fees whenever the consumer buys a product. First, we consider a model in which the platform maximizes sales commission only in the absence of sponsored positions. In the base model with equal prices, this boils down to the platform maximizing the probability of sale of any one product.

Note, that under the absence of a sponsored slot, the consumer does not learn about the match value of the firms in the remaining slots if the platform commits to full obfuscation. This directly follows from the independence of match values across firms. Thus, the posterior match value distribution at any stage of the search process is given by G. Denoting the value of the consumer's outside option by $\eta \in \mathbb{R}$, it follows that under uniform obfuscation, the consumer opts for the outside option only if she inspected every slot if

$$\int_{\eta}^{\infty} (v - \eta) \mathrm{d}G(v) > s, \tag{4}$$

because the option value of sampling any slot is positive if this inequality holds. Conversely, if inequality (4) fails, the consumer does not even begin to search, implying a sales probability of zero if the platform uniformly obfuscates.

Recall that the consumer's reservation value $\bar{r}$ if the platform uniformly obfuscates satisfies

$$\int_{\bar{r}}^{\infty} (v - \eta) \mathrm{d}G(v) = s. \tag{5}$$

Thus, inequality (4) is equivalent to $\bar{r} \geq \eta$ so that our first result can be stated as follows:

Proposition 4. *Uniform obfuscation maximizes sales commission revenue for any n, if and only if, $\bar{r} \geq \eta$.*

It is worth highlighting that the probability of a sale is not necessarily zero if $\bar{r} < \eta$. Instead of uniformly obfuscating over slots, the platform could, for example, use a perfect ranking and in that case the consumer's posterior regarding the first slot(s) first-order dominates G. Consequently, the value of opening some slots would exceed $\bar{r}$ and, thus, potentially η as well.

Given that as far as the organic slots are concerned this result is in line with Theorem 1, we can now also consider that the platform derives revenue from both a sponsored slot as well as commission fees. By Theorem 1, an algorithm that uniformly obfuscates the organic slots and allocates the sponsored slot via a second prize auction (and where ties are broken in favor of higher match scores) is asymptotically optimal if the sole objective is to maximize profits from selling the sponsored slot.

Proposition 4 shows that under uniform obfuscation, if $\eta < \bar{r}$, consumers should explore all firms if their match values are independently distributed. We cannot directly apply this result, however, when the platform maximizes revenue from both sources. If the firm with the highest match score is allocated the sponsored slot, independence across slots fails to persist. Nevertheless, as the mixing property implies that an event can only be strongly related to a finite number of IID random variables, it is true that, if $\eta < \bar{r}$, consumers should inspect more firms as long as there is still a large number of slots remaining. Since the probability that consumers inspect more than m firms is arbitrarily small for large enough m (see Lemma A.2), the following result holds:

Proposition 5. *Suppose the platform sells the sponsored slot and earns a commission fee when selling a product. Then, if $\eta < \bar{r}$, there is a sequence of uniformly obfuscating algorithms allocating the sponsored slot using a second-prize auction that is asymptotically optimal.*

That is, if the platform earns a commission fee when the consumer buys from any of the firms in its list, the following is true. In the absence of a sponsored slot, the platform has an incentive to uniformly obfuscate all slots. With a sponsored slot, however, the sponsored slot contains relevant information to the consumer. From the increased informativeness of the ranking, it immediately follows that introducing a sponsored slot benefits consumers.

Corollary 3. *If, in the absence of a sponsored slot, the platform maximizes sales commission revenue, then introducing a sponsored slot increases consumer welfare.*

Intuitively, the additional information provided to the consumer via the sponsored slot has two effects. First, it ensures that the consumer samples the firm that is most likely to have a high match value. This raises the expected match value of the product the consumer eventually chooses. Second, sampling the best match in expectation first reduces the number of slots the consumer expects to inspect, thereby lowering expected search costs.

5 Discussion and Conclusion

In this paper, we analyze how selling a sponsored slot affects a search platform's ranking of products depending on its objective in the absence of a sponsored slot. When deciding on its ranking, the platform takes into account that consumers are free to choose their search rule, i.e., the order in which to inspect different slots, when to buy, and when to abort search altogether. Obfuscation of organic slots plays an important and dual role: (i) it increases the incentive for firms to acquire the sponsored slot by lowering the consumer's best alternative, thereby boosting the revenues from auctioning the sponsored position. Moreover, (ii) it increases the probability of a consumer buying via the platform by making as many products as possible worthwhile to inspect, thus increasing the platform's sales commission revenue. In addition, we have established the robustness of this result to firms holding private information about the match value of their products for a specific consumer. In fact, platforms may use sponsored positions as a means to acquire consumer-specific information from firms.

Introducing a sponsored position harms consumers, if otherwise the platform's ranking maximizes consumer welfare. This is because with sponsored positions, the platform has an incentive to reduce the informational value of the organic slots. If, in the absence of sponsored positions, the platform maximizes sales commissions, then introducing a sponsored slot benefits consumers by inducing a strictly more informative ranking of firms. Finally, we have shown that, in order for these results to universally apply, the number of firms selling via the platform needs to be sufficiently large. With a small number of keyword relevant firms, counterexamples can be created where consumer learning is an offsetting force.

This paper addresses the relation between search platforms' strategies and the consumer-specific information they possess. We see several fruitful directions for future research in this area. First, throughout this paper, firms do not charge different prices depending on whether or not they are recommended by a platform. How do the platform's incentives change if this condition is violated? Second, in our model, the consumer prefers to inspect all slots in the absence of learning over the outside option. How do platforms react if this is not the case, for example, under competition? Third, platforms tend to choose different numbers of sponsored positions while we restrict the platform to feature a single one. Under which conditions it is optimal for the platform to create an additional slot?

References

S. P. Anderson and R. Renault. Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics*, pages 719–735, 1999.

S. P. Anderson and R. Renault. Search direction: Position externalities and position auction bias. 2021.

M. Armstrong and J. Zhou. Consumer information and the limits to competition. *American Economic Review*, 112(2):534–77, 2022.

S. Athey and G. Ellison. Position auctions with consumer search. *The Quarterly Journal of Economics*, 126(3):1213–1270, 2011.

Y.-K. Che and J. Hörner. Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, 133(2):871–925, 2018.

Y. Chen and C. He. Paid placement: Advertising and search on the internet. *The Economic Journal*, 121(556):F309–F328, 2011.

M. Dewatripont, I. Jewitt, and J. Tirole. The economics of career concerns, part i: Comparing information structures. *The Review of Economic Studies*, 66(1):183–198, 1999.

R. Donnelly, A. Kanodia, and I. Morozov. Welfare effects of personalized rankings. *Available at SSRN 3649342*, 2022.

K. Eliaz and R. Spiegler. A simple model of search engine pricing. *The Economic Journal*, 121(556):F329–F339, 2011.

D. Garcia and S. Shelegia. Consumer search with observational learning. *The RAND Journal of Economics*, 49(1):224–253, 2018.

A. Ghose, P. G. Ipeirotis, and B. Li. Examining the impact of ranking on consumer behavior and search engine revenue. *Management Science*, 60(7):1632–1654, 2014.

J. Glazer, I. Kremer, and M. Perry. The wisdom of the crowd when acquiring information is costly. *Management Science*, 67(10):6443–6456, 2021.

M. Janssen and C. Williams. Influencing search. 2022.

M. C. Janssen, A. Parakhonyak, and A. Parakhonyak. Non-reservation price equilibria and consumer search. *Journal of Economic Theory*, 172:120–162, 2017.

T. T. Ke, S. Lin, and M. Y. Lu. Information design of online platforms. *Available at SSRN*, 2022.

E. Lehmann. Comparing location experiments. *Annals of statistics*, 16(2):521–533, 1988.

E. Mossel, M. Mueller-Frank, A. Sly, and O. Tamuz. Social learning equilibria. *Econometrica*, 88(3):1235–1267, 2020.

V. Nocke and P. Rey. Consumer search and choice overload. *Journal of Political Economy*, forthcoming, 2023.

N. Persico. Information acquisition in auctions. *Econometrica*, 68(1):135–148, 2000.

M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. volume 82, pages 689–711. 2021.

R. M. Ursu. The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, 37(4):530–552, 2018.

M. L. Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.

A. Wolinsky. True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics*, 101(3):493–511, 1986.

A Appendix: Preliminaries

To prove our asymptotic results, we consider the framework with infinitely many firms $i \in \mathbb{N}$ and then embed the model with finitely many firms within it. Let $\theta = (\theta_1, \theta_2, \dots) \in \Theta$ be the vector of scores, $v = (v_1, v_2, \dots) \in V$ the vector of match values, and z the nonatomic random variable with support Z that the platform uses as a randomization device. We maintain the distributional assumptions of Section 2, denote the probability measure on $\Theta \times V \times Z$ by $\mathbb{P}[\cdot]$, and denote expectations taken with respect to this measure by $\mathbb{E}[\cdot]$. Letting $b = (b_1, b_2, \dots) \in B$ denote the bid vector and X the set of firm permutations, the set of algorithms $\mathcal{A}$ constitutes the set of functions $a : B \times \Theta \times Z \to X$ such that $a(b, \cdot)$ is measurable for all $b \in B$.

To embed the finite firm model into this framework, let $\mathcal{A}_n \subset \mathcal{A}$ be the subset of algorithms which only permute the first n firms, which is to say that $a(b, \theta, z) = x$ implies that $x(i) = i$ for all $i \geq n+1$. Thus, the game with n firms is obtained by requiring the platform to select an algorithm in $\mathcal{A}_n$ and restricting the consumer to only be able to inspect the first n positions.

After choosing the algorithm, what follows is a proper subgame played by the firms, the consumer, and nature. Let $\mathcal{A}_n^* \subset \mathcal{A}_n$ be the nonempty subset of algorithms for which an equilibrium of the subgame exists. For the remainder, assume some measurable selection of the set of equilibria of the subgames and define $\Pi : \bigcup_{n=1}^{\infty} \mathcal{A}_n^* \to \mathbb{R}$ to be the platform's expected profit and $D : \bigcup_{n=1}^{\infty} \mathcal{A}_n^* \to \mathbb{R}$ the demand for the sponsored firm given that the selected equilibria are played.

A.1 Consumer Search with Learning

We first present two lemmas that characterize the consumer's search problem in a general environment where the consumer learns while searching. We begin by establishing an algorithm-independent upper bound on the equilibrium expected match value acquired by a consumer from engaging in optimal search. Let $r_{\bar{\theta}}$ denote the reservation value corresponding to the distribution of match values given the maximal score, i.e. the value satisfying $s = \int_{r_{\bar{\theta}}}^{\bar{v}} (1 - G(v|\bar{\theta}))dv$.

Lemma A.1. *For any algorithm and any number of firms $n \geq 1$, the expected match value acquired by the consumer who has optimally searched m firms is less than $u^* \equiv \mathbb{E}[v_i|\theta_i = \bar{\theta}, v_i \geq r_{\bar{\theta}}]$, for all $m \leq n$.*

Proof. Suppose the consumer decides to conclude her search after inspecting m firms. We first argue that the match values at the previously $m-1$ inspected firms must be less than $r_{\bar{\theta}}$. At each of the $m-1$ decision nodes after inspecting each of the $m-1$ initial firms, the expected utility from searching an additional firm depends on the consumer's beliefs about the match scores of the remaining firms. Informing the consumer precisely of the remaining firms' match scores yields a Blackwell improvement and thus must weakly increase the expected utility from continued search. Also, setting the remaining firms' match scores to each equal the largest possible value weakly increases the expected utility from continued search. In this case, it is optimal to follow Weitzman's rule and thus the consumer will halt her search immediately when the match value at the current firm exceeds $r_{\bar{\theta}}$. Because the consumer has optimally searched m firms, it must be that the match values for the previous $m-1$ firms are less than $r_{\bar{\theta}}$. Thus, the expected match value for the consumer who searches m firms is bounded above by

$$\mathbb{E}\left[\mathbb{E}\left[v_i | v_i \geq r_{\bar{\theta}}, \theta_i\right]\right]$$

where the outer expectation is taken over the match score for the mth firm given that its good is selected, and the inner expectation is taken over the match value. This expression too is bounded above by the same expression where we suppose that the mth firm certainly has the maximal match score $u^* \equiv \mathbb{E}[v_i|\theta_i = \bar{\theta}, v_i \geq r_{\bar{\theta}}]$. Thus u^* provides an upper bound for the expected match value acquired by the consumer from engaging in optimal search. □

Using this upper bound, we can likewise bound the probability that a consumer engages in a lengthy search.

Lemma A.2. *For each $\varepsilon > 0$ there exists an $m \in \mathbb{N}$ such that the probability that a consumer searches beyond m firms is less than ε for every algorithm and number of firms.*

Proof. For a given algorithm a, let M denote the random variable equal to the number of organic firms inspected by a consumer who engages in optimal search. From the previous lemma, a consumer's expected match value given that she searches m' firms is less than u^* for all $m' \in \mathbb{N}$. Hence, a consumer's expected utility is bounded above by

$$\sum_{m' \in \mathbb{N}} (u^* - m' \cdot s)\mathbb{P}(M = m'|a).$$

For a given $m \in \mathbb{N}$, the above expression is less than or equal to

$$\mathbb{P}(M \leq m|a)u^* + \mathbb{P}(M > m|a)(u^* - m \cdot s).$$

As the consumer's expected utility from optimal search must be nonnegative, the above expression must likewise be nonnegative, implying

$$\mathbb{P}(M > m|a) \leq \frac{u^*}{m \cdot s}.$$

Thus, regardless of the algorithm or number of firms, if $m > \frac{u^*}{\varepsilon \cdot s}$, then the probability that the consumer searches beyond m firms is less than ε. □

The above arguments only make use of the fact that there is a maximal match score in the support. Employing two more of our assumptions, using Lemma A.1 to bound the probability of lengthy search is even more immediate. Having assumed that the match value distributions share a common support and that there is a minimal match score in the support then because a consumer ceases search after observing a match value larger than $r_{\bar{\theta}}$, the probability of searching at least m firms requires the first $m - 1$ match values to be less than $r_{\bar{\theta}}$

$$\begin{aligned}\mathbb{P}(M > m|a) &\leq \mathbb{P}(v_i \leq r_{\bar{\theta}} \text{ for each initial } m-1 \text{ firms inspected } |a) \\ &\leq \mathbb{P}(v_i \leq r_{\bar{\theta}}|\theta_i = \underline{\theta})^{m-1} = G(r_{\bar{\theta}}|\underline{\theta})^{m-1}.\end{aligned}$$

Thus, for all $\epsilon > 0$ if $m > 1 + \frac{\log \varepsilon}{\log G(r_{\bar{\theta}}|\underline{\theta})}$, then $\mathbb{P}(M > m|a) \leq \epsilon$.

A.2 Mixing

A key idea in our argument is simply that when there are many firms, then for *any* possible algorithm, when the consumer finds that his match value with the sponsored firm lies below

$\bar{r}$, he almost certainly has a better option for how to proceed with his search than to buy the sponsored firm's product. The tool we use to make this simple idea concrete is the fact that independently and identically distributed (IID) random variables have the property of *mixing* (Mossel, Mueller-Frank, Sly, and Tamuz, 2020, Lemma 1). Intuitively, mixing means that any event E defined on the same probability space of a sequence of IID random variables $\{Y_i\}_{i\in\mathbb{N}}$ can only be strongly related to a finite number of them. Formally, for every $\epsilon > 0$, except for a set $N \subset \mathbb{N}$ with $|N| < 1/\epsilon^2$, each $i \notin N$ has the property that for every event K_i only depending on Y_i

$$|\,\mathbb{P}(E \cap K_i) - \mathbb{P}(E)\,\mathbb{P}(K_i)| < \epsilon. \tag{6}$$

The random variables satisfying (6) are said to be *ϵ-independent of E*. In our model, scores $\{\theta_i\}_{i\in\mathbb{N}}$ form an IID sequence of random variables and hence have the mixing property.

B Proof of Theorem 1

To build up to the proof of Theorem 1, we first need to establish a few technical lemmas. To this end, let $F(\theta_i|E)$ denote the distribution over θ_i given the event E, define the *ex ante reservation price* to be the unique value $\bar{r}$ solving $s = \int_{\underline{\theta}}^{\bar{\theta}} \int_{\bar{r}}^{\bar{v}} (1 - G(v_i|\theta_i))dvdF(\theta_i)$, and define the *conditional reservation price* to be the unique value r_i satisfying $s = \int_{\underline{\theta}}^{\bar{\theta}} \int_{r_i}^{\bar{v}} (1 - G(v|\theta_i))dF(\theta_i|E)$.

For our first lemma, notice that if $h : \Delta[\underline{\theta}, \bar{\theta}] \to \mathbb{R}$ is weak* continuous and $\alpha \in (0, 1)$, then for every neighborhood U of $F(\theta_i)$, there exists an ϵ sufficiently small such that, if θ_i is ϵ-independent of an event E with $\mathbb{P}(E) \geq \alpha$, then $h\left(F(\theta_i|E)\right) \in U$. This follows from first recalling that if θ_i is ϵ-independent of E then $|\,\mathbb{P}(K_i|E) - \mathbb{P}(K_i)| < \epsilon/\,\mathbb{P}(E)$ for all θ_i-measurable events K_i and then applying the definition of weak* convergence. Noting that the reservation value is continuous in the distribution over match score, we thus obtain the following.

Lemma B.1. *For a given $\delta > 0$ and $\alpha \in (0, 1)$, there is an $\epsilon > 0$ small enough so that, if θ_i is ϵ-independent of the event E with $\mathbb{P}(E) > \alpha$, then $r_i > \bar{r} - \delta$.*

We have assumed that the consumer does not observe which firms lie in the various positions. Let $\hat{F}(\theta_j|E)$ be the conditional distribution over the match score of the firm in position j and $\hat{r}_j$ the reservation price satisfying $s = \int_{\mathbb{R}} \int_{\hat{r}_j}^{\bar{v}} (1 - G(v|\theta_j))d\hat{F}(\theta_j|E)$. Given that $\Delta[\underline{\theta}, \bar{\theta}]$ is compact, it follows that if ϵ is small enough to guarantee that $r_i > \bar{r} - \delta$ whenever θ_i is ϵ-independent of E, then if the probability a position j contains an ϵ-independent θ_i is q, then there is a q sufficiently large to ensure $\hat{r}_j > \bar{r} - \delta$.

From the mixing property, we know the scores for only finitely many firms fail to be ϵ-independent of an event E. For the consumer deciding which position to inspect next, it is uncertain which positions contain these firms. The next lemma argues that among the countably infinite positions, there exist those which the consumer is arbitrarily certain contains a firm that is ϵ-independent of the event.

Lemma B.2. *Let E be an event with $\mathbb{P}(E) > 0$ and $N \subset \mathbb{N}$ the finite subset of firms who fail to be ϵ-independent of E. Then for every $q \in (0, 1)$ at most $\frac{1}{\epsilon^2 q}$ positions satisfy $\mathbb{P}(x^{-1}(j) \in N|E) \geq q$.*

Proof. Let $p_N(j) = \mathbb{P}(x^{-1}(j) \in N|E)$ be the probability position j contains a firm $i \in N$. If $p_N(j) \geq q$ at m positions, then $\sum_{j\in\mathbb{N}} p_N(j) \geq m \cdot q$. We also know that $\sum_{j\in\mathbb{N}} p_N(j) = \sum_{j\in\mathbb{N}} \sum_{i\in N} \mathbb{P}(x(i) = j|E) = \sum_{i\in N} \sum_{j\in\mathbb{N}} \mathbb{P}(x(i) = j|E) = \sum_{i\in N} 1 = |N|$. Combining these inequalities along with the mixing property provides that $m \cdot q \leq |N| \leq 1/\epsilon^2$ and thus the desired conclusion follows. □

The following lemma provides that, for a given $\delta > 0$, the probability that the consumer buys from the sponsored firm when it offers a match value less than $\bar{r} - \delta$ is bounded from above. Importantly, this bound vanishes as n grows large.

Lemma B.3. *Let $\delta > 0$ and denote by E_n the event in which the consumer buys from the sponsored firm and it offers a match value less than $\bar{r} - \delta$ when the algorithm is $a_n \in \mathcal{A}_n^*$. There exists a real sequence $\{n_k\}_{k\in\mathbb{N}}$ such that, $n > n_k$ implies $\mathbb{P}(E_n) < 1/k$.*

Proof. Lemma B.1 provides that for every $k \in \mathbb{N}$ we can specify $\epsilon_k > 0$ and $q_k < 1$ so that for every event with $\mathbb{P}(E) \geq 1/k$, when the probability that a position j contains a firm whose score is ϵ_k-independent of the event exceeds q_k, then $\hat{r}_j > \bar{r} - \delta$. Using this, define the sequence $n_k \equiv \frac{1}{q_k \cdot \epsilon_k^2}$ for $k \in \mathbb{N}$. For every $n > n_k$, it must be that $\mathbb{P}(E_n) < 1/k$. Otherwise, if $\mathbb{P}(E_n) \geq 1/k$, then there is a contradiction since the consumer would strictly benefit from inspecting some position j when the event E_n occurs. Thus, $n > n_k$ implies $\mathbb{P}(E_n) < 1/k$ for all k. An immediate corollary of this lemma is that $\limsup_{n\to\infty} \mathbb{P}(E_n) = 0$. □

Let us further elaborate on the critical step in the above proof. Given the construction of n_k, if both $n > n_k$ and $\mathbb{P}(E_n) \geq 1/k$ hold, then there must be a position j where the conditional reservation price[8] given E_n satisfies $\hat{r}_j > \bar{r} - \delta$. Supposing the consumer only knows his match value with the sponsored firm and that the event E_n has occurred, then Weitzman's rule provides that the consumer would be strictly better off to first inspect position j and then buy whichever product provides the highest value rather than to buy from the sponsored firm immediately.

[8]That is, $s = \int_{\mathbb{R}} \int_{\hat{r}_j}^{\bar{v}} (1 - G(v_j|\theta_j)) dv_j d\hat{F}(\theta_j|E_n)$.

Notice that whenever E_n occurs, the moment that the consumer decides to buy from the sponsored firm he knows that E_n has occurred. In fact, he has more information since he knows his precise match values at all of the firms he has inspected. Consider an outside observer who only knows whether or not E_n occurs. That is, the observer can only see the output of the indicator function $\mathbf{1}_{E_n}(\theta, v, z)$. From the mixing property, the outside observer who computes reservation prices for all of the positions when E_n occurs would find some position j with a reservation price satisfying $\hat{r}_j > \bar{r} - \delta$. But this means that the consumer's expected payoff would be strictly higher if, in the event E_n occurs, the outside observer forced him to inspect or possibly reinspect position j at the moment the consumer would have bought from the sponsored firm. But because the consumer has even more information than the outside observer does, he must have a profitable deviation that does at least as good as this.

We now argue that when the sponsored firm offers a match value below $\bar{r}$ its demand vanishes as n grows large for *any* sequence of algorithms $\{a_n\}$ with $a_n \in \mathcal{A}_n^*$ for all $n \in \mathbb{N}$. Throughout, we let v_S denote the consumer's match value with the sponsored firm.

Lemma B.4. *For a sequence of algorithms $\{a_n\}_{n\in\mathbb{N}}$ with $a_n \in \mathcal{A}_n^*$ for all n, $\limsup_n D(a_n) \leq 1 - G(\bar{r}|\bar{\theta})$.*

Proof. Notice that if $S_n \subset \Theta \times V \times Z$ is the event in which the consumer buys from the sponsored firm when the algorithm is a_n and the consumer can only inspect the first n firms, for a given $\delta > 0$ we have

$$D(a_n) = \mathbb{P}(S_n \cap \{v_S \geq \bar{r} - \delta\}) + \mathbb{P}(S_n \cap \{v_S < \bar{r} - \delta\}). \tag{7}$$

Lemma B.3 provides that the limit superior of the last term is zero. Thus, we have

$$\begin{aligned} \limsup_{n\to\infty} D(a_n) &= \limsup_{n\to\infty} \mathbb{P}(S_n \cap \{v_S \geq \bar{r} - \delta\}) \\ &\leq \limsup_{n\to\infty} \mathbb{P}(\{v_S \geq \bar{r} - \delta\}) \leq 1 - G(\bar{r} - \delta|\bar{\theta}). \end{aligned}$$

As the inequality holds for all $\delta > 0$, it follows that $\limsup_n D(a_n) \leq 1 - G(\bar{r}|\bar{\theta})$. □

Let $\mathcal{A}_n^* \subset \mathcal{A}_n$ be the subset of algorithms for which there exists an equilibrium among firms and consumers. Let $\Pi : \mathcal{A}_n^* \to \mathbb{R}$ be the platform's expected profit in the platform's preferred equilibrium. Let $\{a_n\}_{n\in\mathbb{N}}$ be a sequence of algorithms with $a_n \in \mathcal{A}_n^*$ for all $n \in \mathbb{N}$.

Lemma B.5. *For a sequence of algorithms $\{a_n\}_{n\in\mathbb{N}}$ with $a_n \in \mathcal{A}_n^*$ for all n, $\limsup_n \Pi(a_n) \leq p(1 - G(\bar{r}|\bar{\theta}))$.*

Proof. The platform's profit is less than or equal to the winning bid. Let β_n denote the expected winning bid given the algorithm a_n. Whenever a firm has a positive probability of winning the auction, its equilibrium bid cannot exceed the expected profit conditional on winning. That is, its bid is bounded above by $p \cdot D(a_n)$. From the preceding proposition, the limiting equilibrium bids satisfy $\limsup_{n\to\infty} \beta_n \leq p \limsup_{n\to\infty} D(a_n) \leq p(1 - G(\bar{r}|\bar{\theta}))$. □

Using uniform obfuscation, it becomes possible to approximately guarantee that consumers will buy from the sponsored firm whenever it provides a match value larger than $\bar{r}$.

Lemma B.6. *Suppose that the firm with the highest match score is placed in the sponsored slot and all other firms are uniformly and randomly placed in organic positions. Then* $\lim_{n\to\infty} \mathbb{P}(S_n|a_n) = 1 - G(\bar{r}|\bar{\theta})$

Proof. First, consider the problem of a "restricted" consumer who is able to search at most m organic firms. Under the proposed sequence of algorithms, the match value distribution for the sponsored slot, $\int_{\underline{\theta}}^{\bar{\theta}} G(v_i|\theta_i) dF(\theta_i)^n$, weakly converges to $G(v_i|\bar{\theta})$ and the match value distribution for each of the first m organic firms, $\int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{\theta}}^{\hat{\theta}} G(v_i|\theta_i) dF(\theta_i) dF(\hat{\theta})^n$, weakly converges to m independent draws from the marginal match value distribution $G(v_i) \equiv \int_{\underline{\theta}}^{\bar{\theta}} G(v|\theta_i) dF(\theta_i)$. At the limiting distribution, the restricted consumer's best response requires beginning search at the sponsored firm and buying immediately from the sponsored firm whenever $v_S > \bar{r}$. This means that the probability of buying the good offered by the sponsored firm without inspecting any organic firms converges to $1 - G(\bar{r}|\bar{\theta})$. Hence, for the restricted consumer, $\lim_{n\to\infty} \mathbb{P}(S_n|\{v_s > \bar{r}\}, a_n) = 1$.

Turning to the unrestricted consumer's problem (i.e. a consumer who is able to search all n firms), we now argue that for n sufficiently large, optimal search requires beginning search by inspecting the sponsored firm. To this end, it is useful to consider an alternative problem in which a consumer faces two firms, the first has the match value distribution conditional on the maximal score $G(v_1|\bar{\theta})$, the second has the marginal value distribution of $G(v_2)$, and the consumer's outside option is $\bar{r}$. Let $\delta > 0$ denote the difference in expected utility between optimally beginning search with Firm 1 and suboptimally beginning search with Firm 2. For the consumer restricted to inspecting at most m firms, for m and n sufficiently large, the difference in expected utility between beginning search with the sponsored firm U_m and beginning with an organic firm U'_m is approximately δ. For the unrestricted consumer, the expected utility from beginning search with an organic firm is bounded above by $(1 - \varepsilon) \cdot U'_m + \varepsilon \cdot u^*$ and beginning search with the sponsored firm yields an expected utility of at least U_m. Thus, the difference between beginning search with the sponsored firm and an organic firm is at least

$$U_m - (1 - \varepsilon) \cdot U'_m - \varepsilon \cdot u^* = U_m - U'_m - \varepsilon \cdot (u^* - U'_m)$$

which is positive whenever $U_m - U'_m > \varepsilon \cdot (u^* - U'_m)$. Thus, for $\varepsilon < \frac{\delta}{u^*}$, for m and n sufficiently large, we have $U_m - U'_m \approx \delta > \varepsilon \cdot u^* > \varepsilon \cdot (u^* - U'_m)$ and the unrestricted consumer commences search with the sponsored firm.

For a given $\varepsilon > 0$, let m be large enough so that the probability that the unrestricted consumers searches beyond m firms is less than ε. The expected utility from continuing search after visiting the sponsored firm can be given an upper bound by the random variable $U^* \equiv u_m \cdot (1-\varepsilon) + u^* \cdot \varepsilon$ where u_m is the random variable denoting the expected utility for the restricted consumer from searching at least one organic slot and then searching optimally thereafter. The probability that a consumer buys immediately from the sponsored firm exceeds

$$\begin{aligned}\mathbb{P}\left(\{v_S \geq U^*\}|a_n\right) &\geq \mathbb{P}\left(\{v_S \geq U^*\} \cap \{v_S \geq u_m\}|a_n\right) \\ &= \mathbb{P}(\{v_S \geq u_m\}|a_n) - \mathbb{P}\left(\{U^* \geq v_S \geq u_m\}|a_n\right).\end{aligned}$$

Thus, taking the limit of the above we find that

$$\begin{aligned}\lim_{n\to\infty} \mathbb{P}\left(\{v_S \geq U^*\}|a_n\right) &\geq \lim_{n\to\infty} \mathbb{P}(\{v_S \geq u_m\}|a_n) - \lim_{n\to\infty} \mathbb{P}\left(\{U^* \geq v_S \geq u_m\}|a_n\right) \\ &= 1 - G(\bar{r}|\bar{\theta}) - \left[G\left(\bar{r} + \varepsilon \cdot (u^* - \bar{r})\right) - G\left(\bar{r}\right)\right] \\ &\geq 1 - G(\bar{r}|\bar{\theta}) - \max_{z \in [\bar{r}, \bar{r}+\varepsilon\cdot(u^*-\bar{r})]} g(z) \cdot \varepsilon.\end{aligned}$$

As the above holds for all $\varepsilon > 0$, it follows that $\lim_{n\to\infty} \mathbb{P}\left(\{v_S \geq U^*\}|a_n\right) \geq 1 - G(\bar{r}|\bar{\theta})$. □

In a putative equilibrium in which all firms submit the same bid and the platform awards the sponsored slot to the firm with the highest match score, we find that optimal search approaches consumers buying immediately from the sponsored firm when it delivers a match value higher than $\bar{r}$ and never buying from the sponsored firm otherwise. Hence, uniform obfuscation reaches the upper bound of demand for the sponsored firm.

Notice that if a firm deviates to a higher bid then it wins the auction for sure, but loses information about whether it has the highest match score. It must be that such a deviation offers an expected profit of $1 - G(\bar{r})$ in the limit.

Lemma B.7. *Consider a strategy profile in which all firms place the same bid, the platform awards the sponsored position to the firm with the highest bid and breaks indifference in favor of the firm with the highest match score, and the consumer searches optimally given this strategy. Let S'_n be the event in which a firm which deviates to a higher bid makes a sale. Then $\lim_{n\to+\infty} \mathbb{P}(S'_n|a_n) \leq 1 - G(\bar{r})$.*

Proof. Due to the assumption that the conditional match value distributions $G(\cdot|\theta_i)$ share a common support for all θ_i, for each $k \in \mathbb{N}$ there is a compact subset $V^k \subset \operatorname{supp} G$ satisfying $\int_{V^k} dG(v_i) > 1 - \frac{1}{k}$ and for which there is a positive number $\delta_k > 0$ satisfying $\frac{g(v_i|\theta_i)}{g(v_i|\theta_j)} \geq \delta_k$ for all $v_i \in V^k$, θ_i, and θ_j. Let $\varepsilon > 0$ and let m be large enough to ensure that the probability of searching beyond m firms is less than ϵ. Because V^k is compact and the likelihood ratio is bounded away from zero for all $v_S \in V^k$, as n grows large, the consumer's beliefs and the true distribution of match values for the first m organic firms given $v_S \in V^k$ both converge to IID draws from the marginal distribution $G(v_i)$. [9] Hence, the probability of buying from the sponsored firm given that $v_S \in V^k$ and that the consumer searches at most m firms converges to $\frac{1}{\int_{V^k} dG(v_i)} \int_{\bar{r}}^{\infty} \mathbf{1}_{\{v_S \in V^k\}}(v_S) dG(v_S)$. This is because, in the limit, the consumer only searches m or fewer firms if one provides a match value larger than $\bar{r}$ and searches beyond the sponsored firm if and only if $v_S < \bar{r}$. Hence, for all $\varepsilon' > 0$ there is an $n_{\varepsilon'}$ sufficiently large such that $n \geq n_{\varepsilon'}$ implies that the probability of buying from the sponsored firm given that $v_S \in V^k$ and that the consumer searches fewer than m firms is less than $\frac{1}{1-\frac{1}{k}} \int_{\bar{r}}^{\infty} \mathbf{1}_{\{v_i \in V^k\}}(v_i) dG(v_i) + \epsilon'$. Thus, we obtain

$$\lim_{n \to +\infty} \mathbb{P}(S'_n | a_n) \leq \left(1 - \frac{1}{k}\right) \left[(1-\varepsilon) \left(\frac{1}{1-\frac{1}{k}} \int_{\bar{r}}^{\infty} \mathbf{1}_{\{v_i \in V^k\}}(v_i) dG(v_i) + \epsilon' \right) + \varepsilon \right] + \frac{1}{k}.$$

Taking the limit of the above as $k \to +\infty$, the above inequality becomes

$$\lim_{n \to +\infty} \mathbb{P}(S'_n | a_n) \leq (1-\varepsilon) \left(1 - G(\bar{r}) + \epsilon'\right) + \varepsilon.$$

As the above inequality holds for all $\varepsilon, \varepsilon' > 0$, it follows that $\lim_{n \to +\infty} \mathbb{P}(S'_n | a_n) \leq 1 - G(\bar{r})$. □

Proof of Theorem 1. Let $a_n^* \in \mathcal{A}_n$ be an algorithm which, for a given β_n: (1) if the highest bid exceeds β_n, assign a firm with the highest bid to the sponsored position, breaking ties in favor of a firm with the highest score and remaining ties with uniform probability; (2) if the highest bid is less than β_n, select a winner of the auction with uniform probability; and (3) assign those firms that do not win the auction to organic slots with uniform probability.

Let us provide the necessary conditions for there to be a symmetric equilibrium in which all firms submit the same bid equal to β_n. Suppose all firms bid β_n. A firm's expected profit from likewise bidding β_n is equal to

$$\frac{1}{n}(\pi(1, a_n) - \beta_n) + \sum_{m=2}^{n} \frac{1}{n} \pi(m, a_n) \tag{8}$$

where $\pi(m, a_n)$ is the firm's expected profit from playing the tentative equilibrium strategy and being relegated to the mth position. Deviating to a lower bid yields an expected profit

[9] For this step, it is necessary to bound how pessimistic a consumer can be made about the organic firms after observing his match value with the sponsored firm. This is why it is necessary for match value distributions to share a common support for all match scores. Absent this assumption, a firm may have a profitable deviation and benefit from winning the auction when it offers a low match value because it discourages the consumer from searching further. See Section B.1 for a concrete example where this occurs.

of

$$\sum_{m=2}^{n} \frac{1}{n-1} \tilde{\pi}(m, a_n) \tag{9}$$

where $\tilde{\pi}(m, a_n)$ is a firm's expected profit from being in a position $m \geq 2$ following the deviation. Deviating to a higher bid yields an expected profit of

$$\hat{\pi}(1, a_n) - \beta_n \tag{10}$$

where $\hat{\pi}(1, a_n)$ denotes the firm's expected profit in the sponsored position following the upward deviation. Combining these conditions, we find bidding β_n to be a best reply if and only if

$$\frac{n}{n-1}\hat{\pi}(1, a_n) - \frac{1}{n-1}\sum_{m=1}^{n} \pi(m, a_n) \leq \beta_n \leq \sum_{m=1}^{n} \pi(m, a_n) - \frac{n}{n-1}\sum_{m=2}^{n} \tilde{\pi}(m, a_n). \tag{11}$$

Denote the leftmost side by λ_n and the rightmost side by ρ_n so that (11) simplifies to $\lambda_n \leq \beta_n \leq \rho_n$. We want to show that as the number of firms grows large $\lambda_n < \rho_n$.

Given the algorithm and the consumer's optimal search, we establish the following. First, because total industry profit is bounded from above, the product $\frac{1}{n-1}\sum_{m=1}^{n} \pi(m, a_n)$ vanishes in the limit. Second, by Lemma B.7, the expected profit when deviating to a higher bid converges to $\hat{\pi}(1, a_n) \to p(1 - G(\bar{r}))$ since the consumer will only make a purchase if his match with the sponsored firm exceeds $\bar{r}$ and the distribution of the match value of the upward deviating firm is $G(v_i) = \int_{\underline{\theta}}^{\bar{\theta}} G(v_i|\theta_i)dF(\theta_i)$. Third, if we let $\tilde{\pi}^*(m, a_n)$ denote the expected profit for a firm deviating to a lower bid and being assigned to a position $m \geq 2$ given that it has the highest score, we have $\tilde{\pi}(m, a_n) = \frac{1}{n}\tilde{\pi}^*(m, a_n) + \frac{n-1}{n}\pi(m, a_n)$. Hence, we can write $\rho_n = \pi(1, a_n) - \frac{1}{n-1}\sum_{m=2}^{n} \tilde{\pi}^*(m, a_n)$ where the bound on industry profit guarantees that the rightmost term vanishes in the limit. Combining these three observations employing Lemma B.6, we find that $\lambda_n \to p(1 - G(\bar{r}))$ and $\rho_n \to p(1 - G(\bar{r}|\bar{\theta})) > p(1 - G(\bar{r}))$, where the inequality holds because $G(v|\theta)$ satisfies the MLRP. Thus, there exists an n^* such that $n \geq n^*$ implies $\lambda_n \leq \rho_n$.

Define $\beta_n = \rho_n$ whenever $\lambda_n \leq \rho_n$ and $\beta_n = 2p$ otherwise. Whenever $\lambda_n > \rho_n$, it is clear that each firm's best reply is to bid zero. Whenever $\lambda_n \leq \rho_n$, there is an equilibrium in which all firms bid β_n. Therefore, we obtain $\limsup_{n\to\infty} \beta_n \leq \limsup_{n\to\infty} \Pi(a_n^*)$ and also $\limsup_{n\to\infty} \beta_n = p(1 - G(\bar{r}|\bar{\theta}))$. Combining this with Lemma B.6, we can conclude that $\{a_n^*\}_{n\in\mathbb{N}}$ is asymptotically optimal if $\{a_n^*\}_{n\in\mathbb{N}}$ exhibits uniform obfuscation. □

B.1 Common Support

It is important to note that requiring the conditional distributions $G(v_i|\theta_i)$ to share a common support for each score θ_i is necessary for the argument of Theorem 1. To demonstrate, consider the example in Section 3.1, except now assume that the consumer is willing to purchase for any match value, though of course still prefers a larger value. With this modification, in a tentative equilibrium with uniform obfuscation, observing that the sponsored firm has a low value still discourages the consumer from inspecting any organic firms since she believes that they all must likewise have a low match value. However, unlike in the earlier example, the discouraged consumer now buys the low value good from the sponsored firm. If we suppose that a consumer is willing to continue searching organic firms when the sponsored firm's value is medium, then letting the number of firms grow large, the expected demand for a sponsored firm given that it has played the equilibrium strategy is $\frac{1}{2}$. That is, in the limit, when playing the tentative equilibrium strategy, the sponsored firm almost certainly has a high signal and thus only makes a sale if it provides a high value. However, if we consider the demand for a firm that deviates to a higher bid and secures the sponsored position, then its demand is $\frac{1}{2}p_H + p_L$ since it makes a sale if it offers either a high or low value. This contradicts Lemma B.7 and also shows that the demand for a lower-valued firm could exceed that of a higher-valued one.

B.2 Proof of Proposition 1

Proof of Proposition 1. As the maximal limiting profit under the information structures are $p\cdot(1-G(\bar{r}|\bar{\theta}))$ and $p\cdot(1-\tilde{G}(\bar{r}|\bar{\theta}))$ respectively, the result follows from showing $\tilde{G}(\bar{r}|\bar{\theta}) \leq G(\bar{r}|\bar{\theta})$. Drawing from the argument for Theorem 5.1 in Lehmann (1988), let $\{\alpha_m\}$ be a vanishing sequence of values in $(0,1)$, $\{t_m\}$ the sequence satisfying $F(t_m|\bar{r}) = 1-\alpha_m$, and $\{\tilde{t}_m\}$ the sequence satisfying $\tilde{F}(\tilde{t}_m|\bar{r}) = 1-\alpha_m$. Due to the fact that $\tilde{F}$ is more accurate than F, we have $F(t_m|v_i) \leq \tilde{F}(\tilde{t}_m|v_i)$ for all $v_i < \bar{r}$ and $F(t_m|v_i) \geq \tilde{F}(\tilde{t}_m|v_i)$ for all $v_i > \bar{r}$. Consider the two posterior probabilities

$$G(\bar{r}|\theta_i \geq t_m) = \frac{\int_{\underline{v}}^{\bar{r}} (1 - F(t_m|v_i))\, g(v_i)dv}{\int_{\underline{v}}^{\bar{v}} (1 - F(t_m|v_i))\, g(v_i)dv}$$

$$\tilde{G}(\bar{r}|\theta_i \geq \tilde{t}_m) = \frac{\int_{\underline{v}}^{\bar{r}} \left(1 - \tilde{F}(\tilde{t}_m|v_i)\right) g(v_i)dv}{\int_{\underline{v}}^{\bar{v}} \left(1 - \tilde{F}(\tilde{t}_m|v_i)\right) g(v_i)dv}.$$

By rearranging terms, we see that $\tilde{G}(\bar{r}|\theta_i \geq \tilde{t}_m) \leq G(\bar{r}|\theta_i \geq t_m)$ if and only if

$$\frac{\int_{\underline{v}}^{\bar{r}} \left(1 - \tilde{F}(\tilde{t}_m|v_i)\right) g(v_i)dv}{\int_{\underline{v}}^{\bar{r}} \left(1 - F(t_m|v_i)\right) g(v_i)dv} \leq \frac{\int_{\bar{r}}^{\bar{v}} \left(1 - \tilde{F}(\tilde{t}_m|v_i)\right) g(v_i)dv}{\int_{\bar{r}}^{\bar{v}} \left(1 - F(t_m|v_i)\right) g(v_i)dv} \tag{12}$$

which must hold as the left side is less than one while the right side exceeds one. Thus, $\tilde{G}(\bar{r}|\theta_i \geq \tilde{t}_m) \leq G(\bar{r}|\theta_i \geq t_m)$ for all m. At the same time, $G(\bar{r}|\theta_i \geq t_m) \to G(\bar{r}|\bar{\theta})$ and $\tilde{G}(\bar{r}|\theta_i \geq \tilde{t}_m) \to \tilde{G}(\bar{r}|\bar{\theta})$ as $m \to +\infty$, implying $\tilde{G}(\bar{r}|\bar{\theta}) \leq G(\bar{r}|\bar{\theta})$. □

C Match scores are private information

Before stating our main result, we must establish two interim results.

Lemma C.1. *For a sequence of algorithms $\{a_n\}_{n\in\mathbb{N}}$ with $a_n \in \mathcal{A}_n^*$ for all n, platform profits satisfy* $\limsup_n \Pi(a_n) \leq p(1 - G(\bar{r}|h))$.

Proof. Since θ follows the same distribution as the match scores in the main model, Lemma B.4 applies. That is, demand in the sponsored slot position satisfies

$$\limsup_n D(a_n) \leq 1 - G(\bar{r}|h). \tag{13}$$

By the argument presented in Lemma B.5, the equilibrium bid β_n thus cannot exceed $p(1 - G(\bar{r}|h))$, which bounds the platform's profits as stated in the lemma. □

For deriving next result, let $\hat{\pi}(m, \theta|\theta')$ denote the expected profit in position $m > 1$ of a firm with type θ, if the sponsored slot is occupied by a firm of type θ'. Similarly, let $\pi(1, \theta|\theta')$ denote the expected profit of a firm with type θ if it occupies the sponsored position ($m = 1$) and the highest type among all other firms equals θ'.

Lemma C.2. *If n is sufficiently large, there exists an equilibrium in which each firm bids*

$$\beta(\theta) = \pi(1, \theta|\theta) - \frac{1}{n-1} \sum_{m=2}^{n} \hat{\pi}(m, \theta|\theta).$$

Proof. We assume that β is monotone, which certainly holds for sufficiently large n. **First, consider downward deviations** to $\hat{\beta} = \beta(\hat{\theta}) < \beta(\theta)$, where $\hat{\theta} = \beta^{-1}(\hat{\beta}) < \theta$. Let θ' be the highest type among other firms and suppose $\theta' \geq \theta$ so that $\beta(\theta') > \beta(\theta)$. In this case, a downward deviation has no effect on i profits.[10] Likewise, the deviation has no effect if $\hat{\theta} > \theta'$ since i wins the sponsored slot regardless (and pays the same price because it is

[10] Even if $\theta = \theta'$, in which case firm i might have a 50% chance of winning, it is easy to verify using the definition of β that profits are $1/(n-1)\sum_{m=2}^{n} \hat{\pi}(m, \theta|\theta')$ regardless of whether i wins, which is equivalent to i's profits if it deviates downwards.

a second price auction). Thus, we can focus on the case that the highest type among all competitors satisfies $\theta' \in (\hat{\theta}, \theta)$. Then, equilibrium profits from bidding $\beta(\theta)$ are

$$\pi(1, \theta|\theta') - \beta(\theta')$$

for any $\theta' \in (\hat{\theta}, \theta)$. By contrast, the deviation to $\hat{\beta} = \beta(\hat{\theta})$ yields

$$\frac{1}{n-1}\sum_{m=2}^{n} \hat{\pi}(m, \theta|\theta').$$

Comparing both shows that a deviation is not profitable if and only if

$$\pi(1, \theta|\theta') - \pi(1, \theta'|\theta') \geq \frac{1}{n-1}\sum_{m=2}^{n} \left(\hat{\pi}(m, \theta|\theta') - \hat{\pi}(m, \theta'|\theta')\right). \tag{14}$$

Claim 1: Fix any (θ, θ') with $\theta' < \theta$. Then, for any $\varepsilon > 0$, there is an n large enough so that

$$\hat{\pi}(m, \theta|\theta') - \hat{\pi}(m, \theta'|\theta') \leq G^{m-1}(\bar{r}|\theta, \theta')\,(1+\varepsilon)\left(\pi(1, \theta|\theta') - \pi(1, \theta'|\theta')\right)$$

for all $m > 1$, where $G(\bar{r}|\theta, \theta')$ denotes the probability that a randomly drawn firm's match value (drawn from all but firm i) does not exceed $\bar{r}$, given the two known match scores θ' and θ.

To see why claim 1 holds, notice that for given (θ, θ'), profits in position m satisfy

$$\lim_{n\to\infty} \hat{\pi}(m, \theta|\theta') = G^{m-1}(\bar{r}|\theta, \theta')p[1 - G(\bar{r}|\theta)] \text{ if } m > 1 \text{ and} \tag{15}$$

$$\lim_{n\to\infty} \pi(1, \theta|\theta') = p[1 - G(\bar{r}|\theta)] \text{ if } m = 1. \tag{16}$$

Expression (16) is a corollary to Lemma B.6 because after the platform receives all truthful bids, the platform applies the uniform obfuscation algorithm as if it had the information ex ante as in the base model. Expression (15) follows from Lemma B.6 as well, subject to one additional observation. By Lemma B.6, the probability that the consumer buys from with match score θ is $[1 - G(\bar{r}|\theta)]$. Thus, to compute $\hat{\pi}$ in the limit, we only need to adjust for the probability that a consumer does not encounter a match value $v \geq \bar{r}$ at any previous firm (if $m > 1$), which gives $G^{m-1}(\bar{r}|\theta, \theta')$.

Using (15), we conclude

$$\lim_{n\to\infty} \left[\pi(m, \theta'|\theta') - \pi(m, \theta'|\theta')\right] = G^{m-1}(\bar{r}|\theta, \theta')p\left[\,(1 - G(\bar{r}|\theta)) - \left(1 - G(\bar{r}|\theta')\right)\right],$$

which is sufficient to prove Claim 1 because $G^{m-1}(\bar{r}|\theta, \theta') < 1$ for all $m > 1$.

Given Claim 1, we can find an integer n large enough so that

$$\frac{1}{n-1}\sum_{m=2}^{n}\left[\pi(m,\theta|\theta')-\pi(m,\theta'|\theta')\right]\leq(1+\varepsilon)\,\frac{1}{n-1}\sum_{m=2}^{n}G^{m-1}(\bar{r}|\theta,\theta')\left[\pi(1,\theta|\theta')-\pi(1,\theta'|\theta')\right]\tag{17}$$

for any (θ,θ') and $\epsilon>0$. In addition, note that $G(\bar{r}|\theta,\theta')<1$ for all $s>0$ since strictly positive search costs imply that $\bar{r}<\bar{v}$. Thus, $\sum_{m=2}^{n}G^{m-1}(\bar{r}|\theta,\theta')$ is a converging geometric series, allowing us to rewrite inequality (17) as

$$\frac{1}{n-1}\sum_{m=2}^{n}\left[\pi(m,\theta|\theta')-\pi(m,\theta'|\theta')\right]\leq \tag{18}$$
$$\frac{G(\bar{r}|\theta,\theta')\left(1-G^{n-1}(\bar{r}|\theta,\theta')\right)}{(1-G(\bar{r}|\theta,\theta'))}\,\frac{1+\varepsilon}{n-1}\left[\pi(1,\theta|\theta')-\pi(1,\theta'|\theta')\right].$$

For every $s>0$ there exists a $k>0$ such that for every pair (θ,θ'), $G(\bar{r}|\theta,\theta')<1-k$. Thus, there exists a large, but finite $\bar{n}$ such that for all pairs (θ,θ') $\frac{G(\bar{r}|\theta,\theta')\left[1-G^{n-1}(\bar{r}|\theta,\theta')\right]}{(n-1)[1-G(\bar{r}|\theta,\theta')]}\,(1+\varepsilon)<1$ for all $n>\bar{n}$. Consequently, downward deviations are not profitable for any pair (θ,θ') if n is large enough.

Second, consider upward deviations to $\hat{\beta}=\beta(\hat{\theta})>\beta(\theta)$. Analogously to studying downward deviations, we can now focus on the case that the highest type among all competitors satisfies $\theta'\in(\theta,\hat{\theta})$. In this case, equilibrium profits are

$$\frac{1}{n-1}\sum_{m=2}^{n}\hat{\pi}(m,\theta|\theta').$$

for $\theta'\in(\theta,\hat{\theta})$. By contrast, the deviation to $\hat{\beta}=\beta(\hat{\theta})$ yields

$$\pi(1,\theta|\theta')-\beta(\theta')$$

Comparing both shows that a deviation is not profitable if and only if

$$\pi(1,\theta'|\theta')-\pi(1,\theta|\theta')\geq\frac{1}{n-1}\sum_{m=2}^{n}\left[\pi(m,\theta'|\theta')-\pi(m,\theta|\theta')\right]\tag{19}$$

for some $\theta'\in(\hat{\theta},\theta)$. Inequality (19) is the same as (14) except for the arguments θ and θ' being swapped in the function π. However, $\theta'>\theta$ in inequality (19), whereas $\theta>\theta'$ in inequality (14). Thus, the exact same argument used to show that no downward deviation is profitable implies that (19) holds for large enough n. □

We are now ready to state and prove the main result.

Proof of Proposition 2. By lemma C.2, firms bid

$$\beta_n(\theta) = \pi(1, \theta|\theta) - \frac{1}{n-1} \sum_{m=2}^{n} \hat{\pi}(m, \theta|\theta)$$

in equilibrium for sufficiently large n. The platform's equilibrium profits are determined by the bid of the firm with the second highest match score. As $n \to \infty$, the second highest type approaches h. Moreover, $\lim_{n\to\infty} \beta(h) = (1 - G(\bar{r}|h))p$. Thus, the platform's profits when using a uniform obfuscation algorithm and a second price auction satisfy $\lim_{n\to\infty} \Pi(a_n^*) = (1 - G(\bar{r}|h))p$. In light of Lemma C.1, we can conclude that $\{a_n^*\}_{n\in\mathbb{N}}$ combined with the second price auction is asymptotically optimal if $\{a_n^*\}_{n\in\mathbb{N}}$ exhibits uniform obfuscation. □

D Welfare effects of Sponsored Slots

Proof of Proposition 3. To be completed. □

Proof of Proposition 4. Follows from the text above the Proposition. □