Discussion Paper Series – CRC TR 224

Discussion Paper No. 021
Project B 05

Consumer Search With and Without Tracking

Marcel Preuss [1]

May 2018

[1] Graduate School of Economic and Social Sciences at the University of Mannheim, Germany.
E-mail: mpreuss@mail.uni-mannheim.de

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

Consumer Search with and without Tracking*

Marcel Preuss†

May 2018

Abstract

In this paper, I develop a tractable framework with sequential consumer search to address the effect of tracking on market outcomes. Tracking search histories is informative about consumers' valuations because different consumer types have different stopping probabilities. With tracking, the unique equilibrium price path is increasing whereas without tracking, an average uniform price prevails. Welfare effects largely depend on how tracking affects consumers' search persistence. For intermediate search costs, tracking based price discrimination exacerbates the hold-up problem and leads to inefficiently low search persistence. For high search costs instead, tracking prevents a market breakdown as low prices conditional on short search histories secure consumers a positive surplus from search. Tracking prevails endogenously when consumers can dynamically opt out from tracking. This holds since disclosing their search history is always individually rational for consumers, irrespective of the overall effect on consumer surplus.

Keywords: consumer search, privacy, dynamic price discrimination.

JEL Classification: D11, D18, D83, L13, L86

*I am extremely grateful to Martin Peitz for his support and guidance. I am also very grateful to Alessandro Lizzeri, Henrik Orzen, Andrew Rhodes, Nicolas Schutz, Sandro Sheliga, Thomas Tröger, Chengsi Wang and Asher Wolinsky for valuable and helpful comments.
Financial support from Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 is gratefully acknowledged.

†Graduate School of Economic and Social Sciences at the University of Mannheim, Germany. E-mail: mpreuss@mail.uni-mannheim.de

1 Introduction

In many markets, consumers learn about products and their prices only by searching different sellers sequentially. Often, the expected number of searches varies greatly across consumers as tastes and preferences are rarely homogeneous. Hence, observing a consumer's search history might convey relevant information to sellers. For instance, think of two consumers A and B (Alice and Bob), and suppose that both are looking for a suit for the job market. Since it is the first suit they buy for a long time, neither of them has a particular preference before they search and they visit the stores of different brands in a random order. While searching, Alice realizes that she is fine with almost any cut and color and thus does not need to search long. Instead, Bob finds that most cuts and colors do not suit him well, requiring him to search longer. When finally encountering his ideal suit, Bob's willingness to pay for it is, most likely, higher than Alice'. This is because Bob not only obtains utility from getting a new suit, but from having the right cut and color as well. In contrast, none of those features matter to Alice, implying that she is willing to pay less. In this environment with niche consumers like Bob and mass consumers like Alice, observing search histories may inform sellers about consumers' preferences and is thus going to provoke sellers' attention. Evidently, tracking a consumer's search process has become a widely used practice both for online and brick and mortar businesses and it will most likely be even more prevalent in the future due to exponentially improving technologies.[1] However, progress in understanding even the most general implications of tracking has been hindered by the lack of tractable models.

The questions I address in this paper are the following: What is the effect of tracking on market outcomes such as search behavior and prices? Does tracking always raise profits or can it - perhaps contrary to common wisdom - also benefit consumers while making sellers worse off? Do the welfare effects depend on the level of search costs? Finally, can tracking prevail in equilibrium if consumers possess measures to prevent it?

To address these questions, I propose a tractable framework of consumer search with tracking. Moreover, I account for consumer heterogeneity with respect to the nicheness of their taste as laid out in the introductory example. Since search with tracking compares with ordered search, the framework provides the first model of ordered search with heterogeneous consumers. Tracking search histories enables sellers to receive imperfect signals about a consumer's type and thus to learn about their preferences because stopping probabilities are type-dependent. In the baseline version of the model, there always exists a unique equilibrium with a price path that is strictly increasing in the order of

[1] For example, Google places cookies on a user's computer if the retailer's website visited uses Google-Analytics for customer management, or if the retailer has joined one of the Google owned ad-networks *Doubleclick* or *Adwords*. Indeed, Mikians et al. (2012) find that both *Google-Analytics* and *Doubleclick* but also other online services providers and advertising networks such as those powered by *Facebook* or *Yahoo* are prevalent on the majority of the 200 most popular shopping websites. Mikians et al. (2012) also used automated bots to mimic different consumer types. Evidently, the bots' browsing histories had been tracked as searches for the same keyword yielded different search results and prices.

search.

I evaluate the welfare consequences of tracking by comparing the tracking equilibrium with the equilibrium when tracking is not available. In general, niche types like Bob are made worse off from tracking while mass types like Alice are made better off because they are more likely to benefit from low prices at early sellers. The welfare consequences of tracking largely depend on its effect on consumers' search persistence, which is the number of sellers they are at most willing to sample if they do not encounter a sufficient match at earlier sellers. In general, tracking raises welfare if it leads to weakly higher search persistence and reduces it otherwise. For a wide range of intermediate search costs, search persistence decreases due to tracking since higher prices conditional on long search histories reduce the incentive for consumers to continue searching. However, search persistence can be lower without tracking if search costs are high. This happens when the market breaks down without tracking because the no-tracking price is inevitably too high.

Overall, it seems without any doubt that tracking fundamentally changes market outcomes, irrespectively of the model used. Consequently, one of the most essential questions appears to be whether we should expect to encounter tracking in markets if it is not imposed exogeneously. In fact, online tracking often requires a consumer's (silent) consent. For example, a consumer must not delete her cookies to enable online retailers to observe her search history.[2] I therefore apply the novel framework of sequential consumer search to investigate whether tracking can arise endogenously. In addition to choosing a stopping strategy, consumers are able to opt out from tracking and thereby prevent sellers from observing their search history at any stage during the search process in this extension of the model. Surprisingly, the unique equilibrium outcome always exhibits full disclosure. The intuition behind this result goes back to Milgrom and Roberts (1986) and their striking unraveling argument. For any alternative equilibrium candidate in which a subset of possible search histories is not disclosed, there always exists a consumer whose search history belongs to the depicted subset and who is better off from allowing tracking.

The full disclosure prediction provides a rational explanation for why only few people delete their cookies or select the 'do not track"-request option provided by their web browsers.[3] Moreover, the analysis provides a useful benchmark for thinking about the regulation of personal data processing. Although both sellers and consumers individually

[2] In the aftermath of several bills being introduced in the US to regulate tracking, all major web browsers integrated the option to send a "do not track"-request into their software. In addition, the European general data protection regulation law (GDPR) mandates to inform consumers when personal information is being processed. As browsing data qualifies as personal information, it requires websites to explicitly ask consumers to agree to the use of cookies. For more information about the interpretation and application of the law, also refer to the "Article 29 Data Protection Working Party" by the European Commission or the paper by Borgesius and Poort (2017).

[3] A study of German internet users from 2013 ("Maßnahmen der Internetnutzer: Digitaler Selbstschutz und Verzicht, conducted by the GfK) shows that while 70% are worried about their privacy, only 29% regularly delete their cookies.

prefer tracking, it may make particularly sellers worse off when it leads to a lower search persistence, which is typically the case for intermediate search costs. As even welfare might decrease due to the forgone matching surplus, there is potential for welfare-increasing regulation when the level of search costs lies in the depicted range.

The model builds on the following assumptions. Consumers must sample sellers sequentially at a cost $s > 0$ to learn about prices set by sellers and match values, which are independently distributed random variables. Consumers are ex ante heterogeneous as they draw match values from different distributions. To keep the model tractable, those distributions are simplified to two-point distributions. While one of the values is normalized to zero for all types, consumers differ with respect to their positive match value. Besides, the probability of drawing a positive match value is assumed to be a function of the match value itself. In the main part of this paper, I assume that the matching probability is decreasing in the "conditional" (positive) match value. That is, high conditional match values coincide with low matching probabilities and vice versa. The assumption seems reasonable in markets consisting of mass and niche consumers as illustrated in the introductory example. Niche consumers have a particular taste hampering their willingness to consider most products suitable. However, once they encounter a product meeting their individual requirements, their utility from the product is relatively high. In contrast, mass consumers find most products satisfactory but only have an average willingness to pay for them.

When sellers learn about their position in a consumer's search process through tracking, search becomes perfectly ordered from their perspective. The analysis shows that consumers must then expect increasing prices in any equilibrium, leading to a simple stopping rule which lets only those consumers without previous matches continue search. Due to the interplay of consumer heterogeneity and the optimal stopping rule, expected demand from consumers with longer search histories is less elastic and, thus, prices indeed increase in the order of search. Consumer heterogeneity also leads to novel predictions regarding the effect of search costs on prices. As intuition suggests, consumers' search persistence is weakly decreasing in the level of search costs. However, the fact that consumers sample fewer sellers does not imply reduced competition and higher prices. From a seller's perspective, the probability of facing a consumer with a longer search history decreases while the probability of facing a consumer with a short search history increases when consumers sample fewer sellers. As demand from the latter group is more elastic, the equilibrium price is decreasing in search costs. This counterintuitive result provides a theoretical explanation for the empirical finding that low search cost environments like the internet sometimes lead to higher prices (Ellison and Ellison, 2014).

The dynamics described above help to understand the welfare implications of tracking. Generally, reduced asymmetric information due to tracking has diametrically opposed effects on consumer surplus. On the one hand, lower prices for short search histories have a positive market expansion effect. On the other hand, improved price discrimination reduces surplus from niche consumers who search longer. Importantly, tracking is not detrimental to every niche consumer per se due to the *order effect.* In other words, both a mass type like Alice and a niche type like Bob can benefit from lower prices at the beginning of search under tracking. Hence, the detrimental effect of price discrimination on consumer surplus is mitigated because learning about a consumer's types does not take place instantaneously but sequentially. If search costs are negligible such that consumers without a match never stop their search before they have sampled all sellers, tracking can be favorable both for consumers and sellers. Otherwise, tracking affects consumers' search persistence and changes consumers' and sellers' surplus in diametrically opposed ways. The longer a consumer's search history, the larger the share of the matching surplus sellers can extract with tracking. Therefore, the expected surplus from sampling an additional seller necessarily falls below search costs beyond some fixed search history. Because the hold-up problem prevents an equilibrium with lower prices, search persistence decreases due to tracking. This effect is most pronounced for an intermediate level of search costs and implies less matches from consumers like Bob with a high willingness-to-pay, leaving sellers with less profits and reducing total welfare while the average consumer is still better off from tracking due to lower initial prices.

In contrast, tracking may also raise everyone's surplus for high search costs. Unless consumers sample only a single seller under search without tracking, the no-tracking price always exceeds the price set by the first seller under search with tracking. Yet, sampling only a single seller in equilibrium may not be consistent with the stopping rule for any level of search costs. Then, there is a range of high search costs where the market with tracking is still active whereas it breaks down without tracking, thus implying significant welfare losses from unrealized matches. The reason why a no-tracking equilibrium with a search persistence of one or few sellers may not exist is because of the adverse effects of search persistence on prices. When the expected surplus from search is negative given the level of search costs and the no-tracking price, a lower search persistence could reduce the price and make sampling the first seller worthwhile again. However, stopping search as early as implied by a low search persistence is not sequentially rational given the reduced price. Moreover, randomizing over sampling the first seller cannot change sellers' belief about the average search history in the market, and, thus, cannot restore the equilibrium with active search if this inconsistency problem prevails at the beginning of the search process.

In an extension, I study the complementary case of a matching probability function which is weakly increasing in the conditional match value. Such a positive relationship is likely to prevail in markets where consumer heterogeneity is mainly determined by heterogeneous budget sets rather than differences in taste. Hence, the two distinct cases of an in- and decreasing matching probability function refer to markets that are likely to be inherently different from one another. Also, market outcomes stand in stark contrast to the baseline version of the model since the price is now weakly decreasing in the order of search. The comparative statics of tracking compare with those for a decreasing matching probability, except for the market breakdown result, which cannot obtain in this specification. Besides, the result that tracking arises endogenously is robust to this extension of the model.

2 Literature

The paper relates to two broad strands in the literature. On the one hand, the search framework developed contributes to the literature on consumer search by embedding ex ante consumer heterogeneity into a model of ordered search, two areas, which so far have been studied only in separation. On the other hand, this paper studies a consumer's privacy data disclosure problem in a sequential search framework and thus relates to the literature on the economics of privacy. More precisely, it complements existing research on the consumer's data protection problem which has mostly been addressed using static models with exogenous data as opposed to a dynamic search environment with endogenous data contained in search histories. I first review the literature on consumer search before providing an overview about the paper's relation to the economics of privacy literature.

In the seminal paper by Diamond (1971), consumers search for prices in a random order. Despite multiple sellers producing a homogeneous good, sellers charge the monopoly price because demand is completely inelastic for any price below the monopoly price due to the hold-up problem. As a consequence, consumers, rationally expecting monopoly prices, are better off from not searching at all. Wolinsky (1986) shows that this counter-intuitive result, often referred to as the *Diamond Paradox*, disappears when products are differentiated and consumers thus search not only for prices but product fit as well. Anderson and Renault (1999) complement Wolinsky (1986) by showing how both the *Diamond Paradox* and the Bertrand outcome arise in the limit as either the degree of product differentiation or the level of search costs vanishes. The model by Wolinsky (1986) with the extension of Anderson and Renault (1999) (henceforth WAR) has since then become the workhorse model of consumer search for many researchers. As in their model, I assume that consumers' preferences are heterogeneous by modeling match values as independently and identically distributed shocks. As opposed to WAR however, I assume that these shocks are identically distributed only for a particular consumer type but differently distributed across types. That is, consumer heterogeneity is revealed not only ex post after sampling sellers, but already prevails *ex ante* even before search begins, thus leading to type-dependent search behavior.

Even though consumers sample sellers in no particular order in my model, tracking enables sellers to learn about their position in a consumer's search proces, which is the defining assumption in the research on ordered search and thus perhaps the closest literature this paper relates to. Arbatskaya (2007) shows that search cost heterogeneity can explain why prices might increase in the order of search even when consumers search for a homogeneous product. Zhou (2011) considers ordered search in the WAR model with differentiated products. He also finds that prices increase in the order of search in equilibrium for the reason that later sellers possess a larger monopoly power over remaining consumers. As prices depend on a seller's position, which in turn, can be inferred perfectly from a consumer's search history, Zhou (2011), in fact, also studies search history-based price discrimination. In a related work, Armstrong and Zhou (2015) analyze a seller's optimal strategy to discriminate between "fresh" and returning consumers. In contrast to this paper, which allows for search histories of arbitrary length, they restrict attention to a duopoly version of the WAR model and mainly focus on a seller's incentive to deter consumer search by offering buy-now discounts.

Besides, even though Zhou (2011) provides a solution to the WAR model with ordered search for a specific distribution of match values, the model is not tractable enough to account for additional consumer heterogeneity that might create potentially countervailing effects regarding price discrimination and search behavior. Indeed, I find that tracking often leads to more efficient search by raising consumers' search persistence, which stands in stark contrast to Zhou's finding that ordered search leads to inefficiently low search. Moreover, his very intuitive price dispersion result disappears if the number of sellers grows and the difference in monopoly power becomes arbitrarily small due to the infinite number of remaining sellers. That is, ordered search or tracking plays no role when the number of sellers is large, suggesting that the WAR model might not capture all important aspects of search markets.

Several papers building on the WAR model and focusing on particular applications of ordered search share this property with the work by Zhou (2011). Among others, Armstrong et al. (2009), Haan and Moraga-González (2011) and Moraga-González and Petrikaitė (2013) have studied how higher quality, more advertising, or merging with a competitor can make a subset of sellers more salient and thereby lead to partially ordered search. Importantly, search in these models is ordered only with respect to the first, salient seller as it would otherwise become intractable.

Though quite different from this paper, some authors have also explored in detail applications of ordered search that are more closely related to internet search and tracking. In Chen and He (2011), a monopoly platform uses an auction to determine the order in which sellers are shown to consumers. As they assume that some sellers' products are more relevant to consumers than others, the mechanism places those sellers at the top consumers have the highest valuations for. De Cornière (2016) studies a platform using a keyword matching mechanism to determine a consumer's search order. If sellers choose to be associated with a keyword the consumer entered and pay an advertsising fee, they obtain a prominent position in the search list. In both papers, the authors start from the fact that sellers vary in terms of their relevance to particular consumers. While it is modeled explicitly only in De Cornière (2016), they thus presuppose that consumers are somewhat ex ante heterogeneous. However, tracking occurs at a single instant prior to actual search in their papers whereas it is modeled as a dynamic process enabling sellers to learn gradually about consumers from search histories in my framework.

Other papers study ex ante consumer heterogeneity within the the WAR model more explicitly but (have to) restrict attention to random search. Moraga-González et al. (2017) study price formation when consumers have different search costs. As changes in the distribution of search costs affect both the extensive and intensive search margin, they find that lower search costs can increase the price charged from actively searching consumers.[4] To understand the effect of targetability or in other words, the quality of search, Yang (2013) studies a seller's choice whether to serve a mass product many consumers like or a niche product appealing only to few. By assuming that consumers draw positive match values only from sellers serving their preferred category and that the probability of encountering a respective seller depends both on the quality of search and the category's coverage in the market, he shows how the long-tail effect is driven by the quality of search. Bar-Isaac et al. (2012) also study the long-tail effect, but do not model ex ante heterogeneity on the side of consumers. Though conceptually similar, the approach taken by Yang (2013) to model mass versus niche consumers is different from mine. Instead of introducing different product categories for mass and niche consumers, I assume that mass consumers are more likely to find any product suitable due to their less restrictive taste compared to niche consumers.

While *search history* based pricing has received relatively little attention in the literature, *purchase history* based price discrimination has been studied extensively. In the absence of online shopping and related privacy concerns, this literature with early works by Hart and Tirole (1988), Fudenberg and Tirole (2000) and Villas-Boas (1999) deals with consumers making purchase decisions in multiple periods. As a consumer's revealed choice for a particular seller is a signal of her willingness to pay, it affects prices she obtains in future periods. A common prediction in this literature is *consumer poaching*: a seller's strategy to offer low prices to those consumers who have revealed their preference for the competitor's product in the previous period. As a consequence of low prices in later periods, prices also become competitive in the initial period. However, when accounting for the possibilty of strategic waiting, Chen and Zhang (2009) identify a novel and opposing incentive to price high in the initial period as it allows to better learn about the willingness to pay of those consumers who still make a purchase.

More closely related to the economics privacy, Taylor (2004) considers purchase history-based customer lists as valuable information one firm would want to sell to another firm if a consumer's valuations for different products are correlated. He finds that privacy protection policies are necessary if consumers are naive, but not so otherwise as the willingness to pay of sophisticated consumers in the first period decreases when anticipating exploitation in future periods. Acquisti and Varian (2005) reaches similar conclusions in a monopoly model where they also model the consumer's decision to remain anonymous as in my model. Conitzer et al. (2012) also study the consumer's privacy choice in a model similar to Acquisti and Varian (2005) but introduce a cost to maintain anonymity.

Motivated by the rising concern for internet privacy, a number of authors have revisited the consumer's data protection problem and extended the analysis of Conitzer et al. (2012). Taylor and Wagman (2014) compare the welfare implications of maintaining privacy across different oligopoly models and find ambiguous effects. In Montes et al. (2017), competing firms can buy data containing consumers' private information from an

[4] I discuss the effect of search cost heterogeity in my framework in section 7.

intermediary unless consumers pay a "privacy cost" to remain anonymous. Since paying the cost also reveals some private information, a higher "privacy cost" may in fact raise consumer surplus. Belleflamme and Vergote (2016) study an environment where a monopolist is able to detect private information with some probability unless consumers use a costly technology to maintain privacy. Similarly to the previous authors, they find that the availability of such a technology makes consumers worse off. The results I obtain are quite different. Although more privacy protection would yield a higher surplus in some cases, consumers never use the technology despite its availability at no cost. Besides, the papers mentioned above apply a static model where information about consumers exists from an exogenous source. If sellers can rely on the availability of informative big data, that approach may be quite accurate. However, if informative data is rare, sellers might pay more attention to a consumer's search history for the product they sell, as studied in this paper.

3 Search Model

There is a continuum of consumers $i \in [0, 1]$ and a finite number of firms N. Firms are selling a horizontally differentiated product. The goods can be produced at constant marginal cost, which is normalized to zero. Sellers set their prices and can condition them on different search histories, if these are observed.

Consumers. Consumers sample sellers sequentially in no particular order and with free recall. They search for both prices and product fitness and pay a sampling cost $s > 0$ for each seller. Consumer i obtains utility $u_{ik} = v_{ik} - p_k$ if she purchases from seller $k \in \mathbb{N}$, where p_k is the price and v_{ik} captures her seller specific match utility. Match utilities for sellers' products $(v_{i1}, v_{i2}, ..., v_{iN})$ are random draws from the set $v_{ik} \in \{0, x_i\}$ and independently and identically distributed across sellers. The conditional match value x_i defines consumer i's type. A type is randomly drawn ex ante from the compact set X with $\underline{v} = \inf(X)$ and $\bar{v} = \sup(X)$ from the log-concave distribution $F(x)$. To avoid corner solutions, assume that $\underline{v} > 0$ is sufficiently small. Denote by $g(x) = Prob(v = x)$ the matching probability, which is type x_i's probability of drawing a positive match value $v_{ik} = x_i > 0$ at a random seller k. The matching probability function $g(x)$ is asumed to be log-concave and monotone decreasing, implying that high conditional match values correspond to low matching probabilities and vice versa. This assumptions seems to fit well many markets where the main difference between consumers is the extent to which they care about all features of a product. In the introductory example, Bob cares both about the cut and the color of his new suit while Alice does not. Consequently, Bob has a lower matching probability than Alice. However, his conditional match value is higher because he obtains utility from all features of the suit. A decreasing matching probability function extends this heterogeneity between picky niche consumers and accepting mass consumers to a continuous type space. Section 5 contains an extension of the baseline model to the case of a weakly increasing matching probability function.

Additionally, both consumers and sellers know only the distribution of types $F(\cdot)$ and a particular consumer's type x_i must be learned during search. Notice that this assumption holds in the example concerning Alice and Bob looking for a new suit. Neither of them knows about their preferences over suits ex ante. Instead, they find out about how picky they are while searching. While this simplifying assumption may seem stark, it actually renders the purpose of search more realistic. In other words, because consumers do not know their conditional match value, they truly search for both price and product fit as in the WAR model. Technically, consumers must not have perfect information about their type x_i ex ante to prevent a market breakdown result as in the *Diamond Paradox.* [5] If all consumers knew their types, sellers would always have an incentive to deviate from any price leaving a strictly positive surplus to actively searching consumers. This is because sellers know that any actively searching consumer's conditional match value must exceed the expected price as the consumer would not have incurred the search costs otherwise. Hence, in the only equilibrium with perfectly informed consumers, no consumer would search.

While the random match value framework I use is more stylized than in WAR, it allows me to handle the complexity arising from incorporating consumer heterogeneity into a model of search with tracking. Recently, several authors have passed on the continuous match value distribution (for individual consumers) as well in order to gain tractability, see for example Chen and He (2011), Anderson and Renault (2015) or Armstrong and Zhou (2011).

Search history. The search history h is what other sellers observe from an arriving consumer under search with tracking. I assume that other sellers cannot observe the price of previously sampled products. Further, nothing can be learned from knowing a particular sellers' identity since match values are uncorrelated across sellers. Consequently, the total number of past sellers a consumer has sampled is a sufficient statistic to update the posterior beliefs about her type. Hence, $h \in \mathbb{N}$.

Timing. Players move in the following order. First, sellers set prices conditional on any feasible search history. Under search without tracking, they set an unconditional price. Prior to searching, nature draws each consumer's type. Next, consumers search by sampling sellers sequentially at a cost s per seller. Under search with tracking, sellers observe the consumer's search history when being sampled. Consumers observe the respective price and match value and decide whether to purchase, to return to a previous seller, or to stop search.

Equilibrium concept. The equilibrium notion I consider is perfect Bayesian Nash Equilibrium (PBE). Sellers choose pricing strategies to maximize expected profits given other sellers' prices and consumers' stopping rule. Consumers maximize surplus by choosing an optimal (possibly non-stationary) stopping rule. Both sellers' beliefs about a consumer's type and consumers' expectations regarding prices need to be consistent

[5] If information is sufficiently imperfect, the hold-up problem has no bite and consumers will find search worthwhile. Thus, no information about x_i is a simplifying but not strictly necessary assumption.

with the equilibrium stopping rule and equilibrium pricing. Since equilibrium pricing strategies will be deterministic, I assume that consumers have passive beliefs if they observe a non-equilibrium price.[6] I restrict attention to symmetric equilibria.

3.1 Search With Tracking

I begin by analyzing the equilibrium under search with tracking and search history based price discrimination and then turn to the analysis of search without tracking.

When consumers sample sellers sequentially, their decision about when to stop searching not only depends on the available match values and prices, but also on prices they expect at forthcoming sellers. The main result of this section is that that there is a unique PBE in which prices satisfy $p_1 < p_2 < ... < p_N$. Therefore, I seek to construct such an equilibrium first and assume that consumers expect $p_1^e \leq p_2^e \leq ...p_N^e$. Second, I show that these beliefs are the only beliefs permissible under rational expectations.

In the following analysis, it facilitates notation to write "seller k" when referring to the k's seller a consumer has sampled. The index k thus does not denote a specific seller for all consumers. Moreover, note that prior to sampling seller k, the consumer's history h equals $k-1$ while it is $h = k$ thereafter. Besides, I omit the subscript i for brevity when it does not lead to ambiguous statements.

Optimal stopping. If $v_{ik} = x_i$ at some seller k, a consumer has no incentive to continue to search as she expects at most to obtain x_i again but to pay a higher price at any forthcoming seller. Therefore, consumers encountering a match $v_{ik} = x_i$ buy if $x_i \geq p_k$ and stop searching without making a purchase otherwise. Whether consumers prefer continuing to search after a history of h unsuccessful matches depends on the continuation value from sampling seller $h+1$, denoted by V_{h+1}. Then, a consumer with history h samples seller $h+1$ if both $v_h = 0$ and $V_{h+1} > 0$.[7]

Lemma 1. *After inspecting seller k, the following non-stationary stopping rule, denoted by $\mathcal{R}^*$, is optimal: Buy if $v_k = x \geq p_k$ and continue to search if $v_k = 0$ and $V_{k+1} > 0$. Otherwise, end search.*

Learning. A consumer perfectly learns her type only upon encountering a match but can learn from the length of her search history otherwise. The optimal stopping rule $\mathcal{R}^*$ implies that a consumer who has a history h and who is about to sample seller $h+1$ must have received only $v_{ik} = 0$ at any seller k, $k \leq h$. As the probability of not encountering

[6]This restriction seems reasonable as an individual consumer is of mass zero here. Hence, an off-equilibrium discriminatory price for a single consumer followed by an out of equilibrium action by this very consumer does not change expected demand at any other seller and thus gives no rise to expect subsequent prices to be different from the equilibrium prices.

[7]Since V_{h+1} has no effect on equilibrium prices, its derivation is postponed to the end of this section. For now, it is sufficient to note that V_{h+1} depends on h but not on an individual's type x_i because consumers do not know their type perfectly but must learn about it from searching.

a match at h previous sellers varies with x, the search history h is informative about one's type. The expected probability of no match at a single seller is given by:

$$\int_{\underline{v}}^{\bar{v}} \big(1 - g(t)\big) f_k(t)\mathrm{d}t,$$

where f_k is seller k's posterior belief about a consumer's type conditional on $h = k - 1$. Since consumers observe their own history by construction, f_k represents their belief prior to sampling seller k as well. Note that in the following analysis, the initial prior $f(x)$ without subscript refers to the distribution of types expected by the first seller a consumer samples. By repeated use of Bayes' rule, I obtain the posterior belief $f_k()$ for any seller k:

$$f_k(x) = \frac{\big(1 - g(x)\big)^{k-1} f(x)}{\int_{\underline{v}}^{\bar{v}} \big(1 - g(t)\big)^{k-1} f(t)\mathrm{d}t}. \tag{1}$$

Consumer demand. The consumer's optimal stopping rule implies that conditional on a match, a consumer always buys the product from seller k immediately if $x_i \geq p_k$ and $p_k \leq p_j^e \ \forall\ j > k$. If the latter constraint is not binding in equilibrium, expected demand from a consumer who is known to be visiting her first firm writes:

$$D_1(p) = P(v \geq p) = \int_p^{\bar{v}} g(x) f(x)\mathrm{d}x,$$

Based on the posterior $f_k()$, a general expression of seller k's demand, denoted by $D_k(p)$, obtains:

$$D_k(p) = \int_p^{\bar{v}} \frac{\big(1 - g(x)\big)^{k-1} g(x)}{\int_{\underline{v}}^{\bar{v}} \big(1 - g(t)\big)^{k-1} f(t)\mathrm{d}t} f(x)\mathrm{d}x \tag{2}$$

Pricing. Seller k's demand at price p is given by (2) if $p \leq p_j^e \ \forall\ j > k$ since consumers might follow an alternative stopping rule otherwise. The following analysis will show that the constraint is not binding at the equilibrium price and that deviating to a price $p_k > p_j^e$ cannot be profitable neither. That is, the profit-maximizing price is independent of all competitors' prices and seller k's problem is equivalent to a monopolist's pricing decision:

$$p_k \in \arg\max_p D_k(p)p \tag{3}$$

The reason why monopoly prices prevail in the presence of competing sellers is similar to Diamond (1971), even though he considers consumer search for homogeneous products. Despite product differentiation however, expectations about future prices suppress any form of price competition between sellers. This holds for any strictly positive search friction. The solution to (3) yields:

Lemma 2. *The profit-maximizing price is uniquely defined for every seller k. The sequence of profit-maximizing prices $\{p_k^+\}_{k=1,...,N}$ satisfies $p_1^+ < p_2^+ < .... < p_N^+$.*

All omitted proofs are presented in appendix 9. Intuitively, prices increase because consumers with a higher conditional match values need to sample more sellers on average than consumers with a low conditional match value until they encounter the first match. Therefore, the relative share of consumers with a high conditional match value is larger for longer search histories. Consequently, expected demand becomes more inelastic and profit-maximizing prices increase in a consumer's search history. The uniqueness of prices is due to the fact that the RHS of the FOC $p = -\frac{D(p)}{D'(p)}$ is decreasing in p, which follows from the log-concavity assumption about $f_k(x)$ and $g(x)$.[8]

The sequence of increasing prices $\{p_k^+\}_{k=1,...,N}$ obtained in lemma 2 is optimal conditional on consumers expecting an increasing price path. Consequently, consumers follow the stopping rule $\mathcal{R}^*$ by lemma 1. Since lemma 2 shows that $\{p_k^+\}_{k=1,...,N}$ is the profit-maximizing sequence of prices given $\mathcal{R}^*$, an equilibrium with prices $\{p_k^+\}_{k=1,...,N}$ and consumer stopping characterized by $\mathcal{R}^*$ indeed exists. Nevertheless, there might be other equilibria. Instead of expecting an increasing price path, consumers might initially expect a decreasing or non-monotonic price path, leading to a different stopping rule and thus to different prices. However, even when allowing for arbitrary consumer beliefs about prices along their search path, the only beliefs consistent with equilibrium pricing of sellers are those of an increasing price sequence given by $\{p_k^+\}_{k=1,...,N}$.[9] This can be shown by means of contradiction. Suppose that consumer expectations $\{p_k^e\}_{k=1,2,...,N}$ are not increasing in k. First, note that:

Lemma 3. *In any PBE, consumer expectations satisfy $p_k^e \geq p_k^+ \ \forall \ k \leq K^*$.*

To see intuitively why lemma 3 holds, consider seller j^* where j^* denotes the seller closest to the end of the search process whose expected price $p_{j^*}^e$ lies below $p_{j^*}^+$. The fact that all remaining sellers are expected to charge higher prices by construction has important implications for seller j^* when deviating to a price in the neighborhood of $p_{j^*}^e$. Seller j^*'s expected demand does not depend on whether arriving consumers have available matches from previous sellers since sampling j^* is only worthwhile if they are in fact willing to buy at $p_{j^*}^e$. That is, seller j^* can sell to all consumers whose match value exceeds the price and thus has full monopoly power over its demand, implying that his problem can be characterized by (3). Nevertheless, the profit maximizing price need not be equal to $p_{j^*}^+$ due to changes in the distribution of arriving consumers. This is because consumers may have applied a stopping rule different from $\mathcal{R}^*$at previous sellers.

Changes in the distribution of arriving consumers must be of the following kind: If a

[8]Bagnoli and Bergstrom (2005) discuss properties of log-concave functions and show that the FOC for demand functions of the type $D(p) = \int_p^{\bar{v}} h(x)\mathrm{d}x$ is decreasing in p if $h(x)$ is log-concave. Hence, decreasingness follows from log-concavity of both $g(x)$ and $f_k(x) \ \forall \ k$.

[9]As in other models of consumer search, there always exists an uninteresting equilibrium where consumers expect prices to be larger than their expected surplus from search. In such an equilibrium, no consumer searches and setting such high prices indeed constitutes an equilibrium strategy for sellers.

consumer samples j^* despite an available match, her conditional match value must satisfy $x_i > p^e_{j^*}$ as she already knows her conditional match value and would otherwise be better off from not sampling j^*. Hence, if due to any alternative stopping rule demand at seller j^* changes, it is due to an increase in demand from types $x_i > p^e_{j^*}$. Notably, demand of these additional consumers attracted by the lower expected price is completely inelastic in the neighborhood of $p^e_{j^*}$. Moreover, note that lemma 2 implies that any (local) upward deviation to $p'_{j^*} > p^e_{j^*}$ is profitable for seller j^* even in the absence of those additional consumers. Hence, setting a price above $p^e_{j^*}$ is for sure profitable in the presence of this additional, perfectly inelastic demand. Since this is true for any $p^e_{j^*} < p^+_{j^*}$, the argument can be applied repeatedly from the last to the first seller, yielding lemma 3.

In addition, notice that consumers apply $\mathcal{R}^*$ at the first seller if he sets a price p_1 in the neighborhood of p^+_1 since $p^e_k > p^+_1$ by lemma 3 $\forall\ k > 1$. By lemma 2, p^+_1 maximizes the first seller's profits under the stopping rule $\mathcal{R}^*$. Hence, any alternative price $p'_1 > p^+_1 + \delta$ ($\delta > 0$) inducing an alternative stopping rule for some types must yield strictly lower profits. This is because under any alternative stopping rule, there are some types who continue searching despite a match $v_{i1} > p_1$ and return only with some probability less than one. Consequently, demand and thus profits must be strictly lower than when consumers apply $\mathcal{R}^*$. It follows that consumers must expect $p_1 = p^+_1$ in any PBE. Further, the same argument applies to the second seller a consumer visits and so forth. Thus, no equilibrium exists, in which consumers do not expect an increasing price path.

Proposition 1. *With tracking, in the unique equilibrium, prices increase in the order of search and consumers follow the stopping-rule $\mathcal{R}^*$. The equilibrium always exists.*

Search persistence. While search costs have no effect on equilibrium prices under search with tracking, they matter for consumers' search persistence: how long to continue search if no match occurs. Consumers' search persistence is captured by $K^* \in \mathbb{N}$. First, note that K^* depends on a consumer's continuation value V_{h+1} at any history h. Its recursive formulation writes:

$$V_{h+1} = \mathbb{E}\big[\max\big\{[v_{h+1} - p_{h+1}], 0, V_{h+2}\big\}\big] - s.$$

To derive an explicit expression for V_{h+1}, one needs to account for K^* since the continuation value from sampling seller k must contain the option value from continuing to search at least seller $k+1$, which is feasible only if seller k is sampled first. For any K^* and history $h < K^*$, the continuation value writes:

$$V_{h+1}(K^*, s, p_1, p_2, ..., p_{K^*}) = \frac{1}{\int_{\underline{v}}^{\bar{v}} \big(1 - g(x)\big)^h f(x)\mathrm{d}x} \cdot \Bigg\{ ...$$

$$\sum_{j=h+1}^{K^*} \Bigg(\underbrace{\int_{p_j}^{\bar{v}} g(x)(x - p_j)\big(1 - g(x)\big)^{j-1} f(x)\mathrm{d}x}_{\textit{weighted matching surplus}} - s \underbrace{\int_{\underline{v}}^{\bar{v}} \big(1 - g(x)\big)^{j-1} f(x)\mathrm{d}x}_{\textit{expected search attmepts}} \Bigg) \Bigg\}. \quad (4)$$

Even though $V_{h+1}(K^*, s, p_1, ..)$ always depends on all sellers' prices, equilibrium search persistence K^* and search costs s, I omit those arguments for brevity when it does not affect comprehensibility. The term before the curly brackets is part of the belief updating regarding the consumer's type and for $h = 0$, the term disappears. It equals the inverse of the total probability of not finding a match after sampling h sellers and thus normalizes the probability of encountering a match at sellers $h + 1$, $h + 2$, and so forth. The term within the brackets sums over the surpluses from additional search attempts weighted by the updated consumer's type after history $h + j$. Notice that the sum of additional search attempts goes from any history h to K^*.

For the optimal search persistence K^*, it must hold that $V_k(K^*) \geq 0 \;\forall\; k \leq K^*$ and $V_{K^*+1}(K^* + 1) < 0$. In words, sampling any seller prior to K^* must be rational. Moreover, there cannot be another $K' > K^*$ satisfying the first the first condition and rendering $V_{K^*+1}(K') \geq 0$. Since continuation values are affected by search costs, K^* depends on search costs as well. Formally, define $\hat{K}(s) := \{K : V_k(K, s) \geq 0 \;\forall\; k \leq K \in \mathbb{N}\}$. Then, equilibrium search persistence satisfies $K^* \in \mathbb{K}(s)$, where

$$\mathbb{K}(s) := \big\{K : V_{K+1}(K', s) < 0 \;\forall\; K' \in \big(\hat{K}(s) \cap \{k : K' > K\}\big)\big\}. \tag{5}$$

By construction of $\mathbb{K}(s)$, there always exists exactly one K^* for any given s. Besides, equation (4) shows that for a fixed K^*, $V_{h+1}(K^*, s)$ is decreasing in s for any history h. Hence, $\hat{K}(s') \subseteq \hat{K}(s)$ for any $s' > s$. Consequently, $K^*(s)$ as defined in equation (5) must be weakly decreasing in s. This result is, of course, very intuitive. As prices are independent of search costs, an increase in search costs reduces the continuation value for any history h. That is, an increase in search costs can only make consumers switch from continuing to search given a particular history h to stopping to search given the same history, but will never induce a change in the other direction. As a consequence, search persistence cannot be increasing in search costs.

3.2 Search Without Tracking

Without tracking search histories, sellers cannot price discriminate. Therefore, sellers expect the same demand (elasticity) from any newly arriving consumer. Since given those expectations, only one price maximizes profits, search without tracking implies a uniform equilibrium price set by all sellers. As a consequence, consumers have no incentive to defer the purchase decision after encountering a match and thus their optimal stopping rule equals $\mathcal{R}^*$. That is, i either buys if $x_i > p_{ij}$ or leaves the market without a purchase which is identical to the uniquely optimal stopping rule when $g(x)$ is decreasing. Consequently, sellers set prices monopolistically. Due to the simple optimal stopping rule, the distribution of consumer types $f_k(x)$ at any seller k is correctly specified by the posterior belief derived in (1). Further, demand functions - if sellers could discriminate - are given by (2), i.e. they are identical to those under search with tracking. However, only consumers but not sellers can update the beliefs conditional on different search histories. Hence, the sellers' expected demand is composed of the expected demand for each possible search history, weighted by the respective probabilities. Based on the common prior $F(x)$ and the matching probability function $g(x)$, sellers can compute the probability that a consumer has a history of $h = 0, 1, 2, ..., K^* - 1$ previous sellers. Note that consumers' search persistence K^* depends only on equilibrium prices but does not change in response to any deviating price. I analyze the equilibrium level of K^* after discussing the equilibrium pricing. The probability ϕ_k of being in position $k = h + 1$ in a consumer's search process writes:

$$\phi_k = \tag{6}$$

$$\underbrace{\frac{1}{N} \cdot \prod_{j=1}^{k-1} \int_{\underline{v}}^{\bar{v}} \big(1 - g(t)\big) f_j(t) \mathrm{d}t}_{no\ match\ up\ to\ k-1} = \frac{1}{N} \cdot \prod_{j=1}^{k-1} \int_{\underline{v}}^{1} \frac{\big(1 - g(t)\big)^j f(t)}{\int_{\underline{v}}^{\bar{v}} \big(1 - g(t)\big)^{j-1} f(t) \mathrm{d}t} \mathrm{d}t$$

$$= \frac{1}{N} \int_{\underline{v}}^{1} \big(1 - g(t)\big)^{k-1} f(t) \mathrm{d}t \ \forall\ k \leq K^*.$$

and $\phi_k = 0\ \forall\ k > K^*$.

Notice that ϕ_k is the unconditional probability for being in a particular position. While this is the actual probability of being sampled by a consumer with history $h = k - 1$, sellers can condition the probability on the fact that the consumer is still searching. However, normalizing by $1/\sum_k^{K^*} \phi_k$ has no effect on a seller's first order condition. Expected demand is composed of the expected demand functions for each possible search history, weighted by the respective probabilities. Denote by $D(p)$ the expected demand weighing the individual demand functions from $D_1(p)$ to $D_{K^*}(p)$ at the seller's non-discriminatory unit price p, i.e. $D(p) = \sum_{k=1}^{K^*} \phi_k D_k(p)$. Sellers maximize:

$$\Pi(p) = p \sum_{i=1}^{K^*} \phi_i D_i(p) = \sum_{i=1}^{K^*} \phi_i D_i(p) p = \sum_{i=1}^{K^*} \phi_i \pi_k(p) \tag{7}$$

where $\pi_k(p)$ equals the profit function of a seller at the k's position if discrimination was feasible and where ϕ_i is given by (6) (see the appendix) and $D_k(p)$ by (2). As log-concavity of the individual demand functions is preserved in this weighted demand function, I obtain the following:

Proposition 2. *Without tracking, the unique equilibrium has a uniform price.*

Since sellers cannot observe consumers' strategies before consumers make a purchase, they must have symmetric beliefs in equilibrium. Given these beliefs, there exists a unique optimal price, set by all sellers. Thus, consumers must believe that prices are constant in any PBE. Consequently, the equilibrium with the price maximizing (7) is unique even when allowing for arbitrary expectations ex ante.

Without tracking, there is no price discrimination. Yet, consumer heterogeneity has another effect on the comparative statics of equilibrium price, which depends on consumers' search persistence. Intuitively, the more sellers consumers are at most willing to sample, the smaller the share of consumers with short search histories each seller can expect because probability mass shifts to longer search histories. Since demand from consumers with longer search histories is less price elastic, the equilibrium price increases. Let $p(K^*)$ be the uniform random search price if a consumer's search persistence equals K^*, then:

Lemma 4. *For $K_2^* > K_1^*$, it holds that $p(K_2^*) > p(K_1^*)$.*

By lemma 4, more persistent search behavior by consumers leads to higher prices. Thus, consumer heterogeneity implies a novel, and perhaps surprising, effect of search persistence on the equilibrium price, which is not present in the WAR model.

3.3 Search Persistence under Search Without Tracking

An increase in K^* as discussed in lemma 4 can be the result either of an increase in the number of sellers in the market or of a reduction in search costs. The former is true if search costs are sufficiently low such that $K^* = N$. Intuitively, the latter might be true if $0 < K^* < N$ and K^* increases due to a decrease in search costs. While this turns out to be correct, it does not follow immediately from the construction of $K^*(s)$ given in equation (5) due to the reverse effect of K^* on prices.

The full characterization of equilibrium search persistence as a function of search costs is provided by the lemmata 13, 14 and 15 in appendix 9. To summarize, there exist two disjoint sets of search cost intervals that give rise to a different characterization of search persistence. By lemma 13, consumers' search persistence K^* equals a fixed number of sellers for a set of relatively large intervals. Lemma 14 and 15 characterize intervals where consumers continue sampling seller k only with some probability less than one. That is, only a fraction of consumers without a match samples seller k while all remaining consumers stop search.[10] Between k and K^*, no further radomized stopping occurs since all continuation values are strictly positive.

To see intuitively why random stopping occurs, consider the following argument. By lemma 4, the optimal price without tracking is a function of K^*, which, in turn, depends on the continuation value and thus on the level of search costs. As can be seen from equation (4) by substituting $p_j = p \;\forall\; j$, the continuation value V_{h+1} also depends on the no tracking price $p\big(K^*(s)\big)$, which in turn, depends on $K^*(s)$. That is, the continuation value under no tracking for any history h is given by $V_{h+1}\big(K^*, p\big(K^*(s)\big), s\big)$.

Ceteris paribus, a rise in search costs thus reduces the continuation value from search and leads to a lower search persistence K^*. However, if consumers sample fewer sellers

[10]Recall that k identifies a seller's position in any consumer's search process and not a unique seller.

in total, lemma 4 implies that the profit-maximizing price decreases, thus increasing consumer surplus. When search costs are such that a consumer is just indifferent between sampling seller k and stopping search, a marginal increase in s has only a marginal direct effect on the continuation value from search while the indirect effect through a lower k is large. Without random stopping, the discontinuity in K^* results in an inconsistency that rules out an equilibrium in pure strategies. Instead, consumers shift sellers' beliefs towards expecting more consumers with shorter search histories by sampling later sellers only with some probability less than one. This prevents price jumps and restores the equilibrium.

It remains to analyze the effect of search costs on prices in the absence of tracking. Considering only the intervals where consumers do not use mixed strategies, higher search costs imply a lower K^* by lemma 13 and thus a decrease in prices by lemma 4. When search costs are in a region where consumers follow a mixed stopping rule, prices must decrease as well. Consider an increase in search cost from $s = \hat{s}^*_K(K^*)$ to some $s > \hat{s}^*_K(K^*)$ requiring consumers to sample seller K^* only with some probability.[11] If search costs increase, the equilibrium probability of continuing to search decreases. This is because it reduces the mass of actively searching consumers with any history $h = K^* - 1$ such that sellers expect fewer consumers with long search histories and set lower prices. Hence, search persistence decreases smoothly in search costs.

Proposition 3. *Without tracking, the uniquely defined uniform price is weakly decreasing in search costs.*

The intuition behind this result follows immediately from lemma 4. Lower search costs increase consumers' search persistence, which reduces every seller's share of elastic demand and, thus, leads to higher prices. Proposition 3 provides a micro-founded theoretical explanation for some empirical papers suggesting that the internet does not always lead to lower prices despite reducing search costs. For example, Ellison and Ellison (2014) find that prices for used books are higher online than offline. In line with my model's predictions, the authors argue that higher prices obtain because sellers expect to sell mostly to consumers with high match values when consumers are willing to search longer due to lower search costs.

It is important to note that a mixed stopping rule that can continuously decrease a seller's belief about the average search history need not exist always. That is, consumers' search persistence may not decrease gradually but may immediately fall from $K^* > 1$ to zero. Then, the *no search* equilibrium, which always exists, is the only equilibrium. Formally, define by $\hat{s}_k(K^*) \in \{s : V_k\big(K^*, p\big(K^*(s)\big), s\big) = 0\}$ the threshold level of search costs such that sampling seller k conditional on $p(K^*)$ and K^* is worthwhile if and only if $s \leq \hat{s}_k$. Then,

Proposition 4. *If there exists an equilibrium with $K^* = \underline{K}^* \leq N$ for some level of*

[11] The notation is explained in the appendix. For the argument however, it is sufficient to treat $\hat{s}^*_K(K^*)$ as some fixed threshold.

search costs such that $\hat{s}_1(\underline{K}^*) < \hat{s}_k(\underline{K}^*)$ $\forall$ $1 < k \leq \underline{K}^*$, no consumer searches and the market breaks down if $s > \hat{s}_1(\underline{K}^*)$.

Notably, the no search equilibrium prevails even though the surplus from search would be positive for consumers if they were able to commit to sampling less than a certain number of sellers. However, for any price making the first search worthwhile, the search persistence, resulting from consumer's sequentially optimal stopping decision, exceeds the search persistence for which the assumed price is maximizing profit. Instead, sellers anticipate consumers' search persistence conditional on initiating search and want to set a higher price, making sampling even the first seller not worthwhile for consumers. Further, the conditions from proposition 4 rule out any mixed stopping rule as shifting sellers beliefs to a lower average search persistence is not feasible when the problem occurs at the first seller.

Proposition 4 applies if threshold search cost levels are not decreasing in the order of search. Importantly, this always happens for some level of search costs if the continuation value is not only decreasing with longer search histories. Technically, this depends on $g'(x)$ as well as $x \cdot g(x)$. In reality, a situation where the continuation value increases may in fact be quite common. Without knowing well what he is looking for, Bob might not be too enthusiastic about getting a new suit prior to searching. While sampling the first sellers however, Bob might learn that he likes particular kind of buttons and becomes excited about finding a suitable shirt. As a consequence, his interim continuation value from search might well exceed his expected surplus prior to search.

4 Comparative Statics of Tracking

As the subsequent analysis will show, the implications of tracking for overall welfare depend on the level of search costs and thus on consumers' search persistence as well. Because obtaining predictions that depend on the model's fundamentals requires additional structure, I first discuss how the effects of tracking vary with the search persistence parameter K^*.

4.1 General Analysis

The first result concerns equilibrium prices and immediately follows from the sellers' first order conditions.

Proposition 5. *For any $K^* > 1$, prices with and without tracking satisfy:*

$$p_1 < p(K^*) < p_{K^*}.$$

By proposition 5, the uniform no-tracking price exceeds the price consumers face at their first seller when searching with tracking but is strictly below the last seller's price.

Moreover, a general result regarding profits is available when K^* under search with tracking is at least as high as under search without tracking:

Lemma 5. *If K^* is weakly larger under search with tracking, sellers' profits are strictly larger under search with than under search without tracking.*

The reason is fairly intuitive. If consumers sample the same number of sellers under both regimes, the aggregated distribution of types from all consumers is identical. Then, the only difference between tracking and no tracking is that in the former case, sellers can condition the optimal price on private information about consumers. By proposition 1, prices are increasing in the order of search under tracking and thus different from a uniform price. Hence, the uniform price is not profit-maximizing if better information is available, implying that tracking yields higher profits. Importantly, K^* is indeed weakly larger under search with tracking in many cases. If search costs are sufficiently low and the number of sellers not too large, $K^* = N$ irrespective of tracking. Consequently, it follows generally that sellers always benefit from tracking if there are not too many of them or search costs are sufficiently low.

Moreover, the market breakdown result stated in proposition 4 implies that if the continuation value from search is not decreasing in the consumer's search history, there exists a threshold search cost level $\hat{s}_1(\underline{K}^*)$ such that $K^* = 0$ for any $s > \hat{s}_1(\underline{K}^*)$ under search without tracking. Thus,

Proposition 6. *If there exists a $\underline{K}^* \leq N$ such that $\hat{s}_1(\underline{K}^*) < \hat{s}_k(\underline{K}^*)\ \forall\ 1 < k \leq \underline{K}^*$, tracking leads to strictly higher consumer surplus and profits for $\hat{s}_1 > s > \hat{s}_1(\underline{K}^*)$.*

Proposition 6 holds irrespective of how much surplus sellers would extract from consumers via search history-based price discrimination. For $s \in [\hat{s}_1(\underline{K}^*), \hat{s}_1)$, it holds that $V_1(p_1^+, s) \geq 0$ under search with tracking. That is, the market under search with tracking is active at a level of search costs where the market without tracking is not. Consequently, tracking raises everyone's surplus. The reason is that tracking reduces the information asymmetry between consumers and sellers and thereby leads to sufficiently low prices, making initiating search worthwhile for consumers. In contrast, an equilibrium with low search persistence and low prices does not exist without tracking under the conditions of proposition 6. Proposition 6 also stands in contrast to the results derived by Zhou (2011). Without ex ante consumer heterogeneity, he finds that consumer search is inefficiently low when search is ordered and leads to a lower overall surplus.

As opposed to standard monopolistic group pricing, the following analysis suggests that tracking may raise overall consumer surplus even when there is no market breakdown in the absence of tracking. Similarly to group price discrimination, it has a market expansion effect since prices are lower for consumers with a lower expected willingness to pay.[12] However, discrimination based on search histories is likely to benefit consumers

[12] In fact, Belleflamme and Peitz (2015) demonstrate how partitioning the demand into smaller intervals ("groups") has a non monotonic effect on consumer surplus due to two opposing factors. Due to better

even more than standard group price discrimination. Rather than being part of mutually exclusive groups, consumers arriving at some seller k have already had the opportunity to buy at any seller $j < k$, implying that the different "groups" of consumers facing different prices are in fact subsets of one another. Moreover, the market price without tracking is above the first and below the last seller's price with tracking by proposition 5. Hence, there must exist a threshold history $\bar{h}$ such that only consumers with a history $h > \bar{h}$ pay discriminatory prices exceeding $p(K^*)$. Note that consumers with a niche taste search longer on average than consumers with a mass taste do. Hence, in expectation, some types are made better off from tracking while others are made worth off.

Corollary 1. *There exists a cut-off type $\tilde{x} \in X$ such that a consumer's expected surplus is reduced due to tracking if $x > \tilde{x}$.*

4.2 Linear Matching Probability

In this section, I impose further structure on the model to derive additional analytic results. In particular, I am interested in observing when the reduction of asymmetric information via tracking can lead to welfare and consumer surplus improvements, whether tracking always raises industry profit and, how these effects depend on search costs. Consider a linear matching probability function $g(x) = 1 - x$ and let the type x be uniformly distributed on $[0, 1]$. The total number of sellers is held constant at $N = 10$ but consumers may sample only $K^* \leq N$. Proposition 7 summarizes the findings from this section. Importantly, additional computations show that qualitatively identical results obtain from any linear matching probability function.

Proposition 7. *There exist two thresholds $s_1^p < s_2^p$ such that profits are strictly larger under search without tracking if $s \in (s_1^p, s_2^p)$. Besides, there exist two thresholds $s_1^w < s_2^w$ such that welfare is strictly larger without tracking if $s \in (s_1^w, s_2^w)$. Consumers surplus is always higher under tracking. The market breaks down without tracking for $s > s_2^w$.*

The prices under search with tracking can be obtained from solving the maximization problem as specified in (3). Under search without tracking, I observe that prices increase in K^*.

The continuation value for every history and feasible K^* is given by equation (4). Under search with tracking where continuation values and threshold search cost levels are decreasing due to increasing prices, K^* can be derived from the set $\mathbb{K}$ as defined in equation (5). For search without tracking, the dependence of the continuation value on K^* via its effect on prices imposes an additional constraint on the optimal K^* as explained in section 3.3. Beyond $V_k(K^*, p(K^*)) \geq 0 \ \forall \ k < K^*$, it requires that $V_{K^*}(K^*, p(K^*)) \geq 0$ for the price that is optimal conditional on Kj. Lemma 13 specifies the search cost intervals

information about the willingness to pay sellers charge lower prices from groups with a lower willingness to pay - eventually leading to an expansion of the market as consumers are served that would not have bought under the uniform price. However, as this information becomes more precise, the surplus left to each group is decreasing and approaches zero for infinitesimally small intervals.

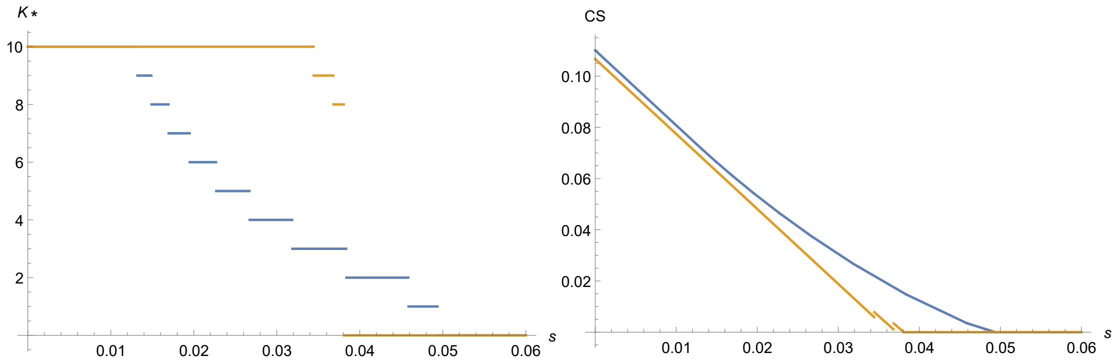

(a) Search persistence

(b) Consumer surplus

Figure a) shows the search persistence K^* as a function of search costs. Figure b) shows consumer surplus as a function of search costs. Orange represents search without tracking, blue search with tracking.

Figure 1: Search costs, search persistence and consumer surplus

for every $K^* \leq N$. Pure stopping strategies sometimes lead to dynamic inconsistencies due to the adverse effect of K^* on the price p. By lemma 14, consumers with the maximum possible search history in the market randomize over sampling the "last" seller for some levels of search costs. However, the range of the intervals where consumers would choose mixed strategies is relatively small compared to those where they choose pure strategies under the model's specifications. Therefore, I omit the calculation of the mixed strategies.[13]

The sharp drop of K^* under search without tracking at around $s = 0.038$ as displayed in figure (1a) illustrates the market breakdown result from proposition 4. If $s = 0.038$, the continuation value from sampling the first seller (expected surplus from search) falls below zero while to any other interim continuation value would still be positive. Thus, the threshold search cost is lowest prior to sampling search and the dynamic search inconsistency problem cannot be prevented by a randomized strategy. Hence, the market shuts down entirely and K^* equals zero. As a consequence, consumer surplus is significantly lower under search without tracking when search costs hit the market-breakdown threshold as shown in figure (1b).

Moreover, it can be seen that tracking leads to higher consumer surplus even in the absence of the market breakdown. For search cost in the neighborhood of the cut-off level $s = 0.038$, this is not too surprising. Both consumer surplus with and without tracking are continuous functions of search costs. Since it is strictly higher with tracking at $s = 0.038$, it has to be higher for lower search costs as well. For the given linear specification of the model, tracking always raises consumer surplus. This holds irrespective of the fact

[13]The length of the intervals where consumers would mix can be inferred from figure (1b). The discontinuities between $s = 0.03$ and $s = 0.04$ would disappear if mixed strategies were accounted for. Importantly, neglecting the mixed stopping behavior does not affect the computation of consumer surplus, as $V_j = 0$ if consumers randomize over sampling j.

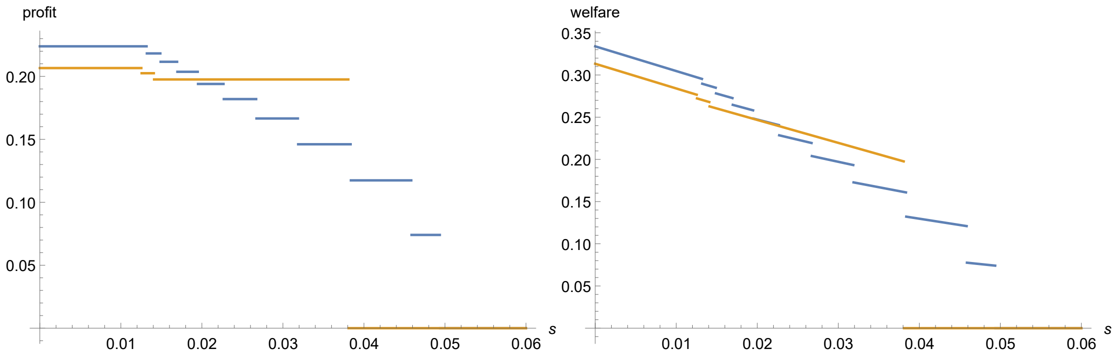

(a) Profits (b) Welfare

Figure a) shows total profits as a function of search costs. Figure b) shows welfare as a function of search costs. Orange represents search without tracking, blue search with tracking.

Figure 2: Search costs, profits and welfare

that for some levels of search costs, consumers sample fewer sellers with tracking. The foregone consumer surplus from sampling the last sellers is, however, relatively small due to higher prices conditional on longer search histories. Thus, consumers always benefit on average.

Note that, surprisingly, tracking does not always maximize industry profit as can be seen in figure (2a) for intermediate search costs. In particular, this happens if consumers' search persistence is significantly lower under search with tracking than without tracking. More precisely, the benefit of tracking for sellers comes from exploiting detailed information about consumers with long search histories. If search costs are too high however, consumers with long search histories (and high expected match values) anticipate that they will be left with almost no surplus and do not continue sampling additional sellers. While sellers would in general benefit from promising to leave a surplus to consumers compensating them for search costs, they cannot due to the hold-up problem.

Thus, the situation when facing a consumer with a long search history is comparable to Diamond (1971), where sellers have perfect information about a consumer's willingness to pay. In other words, the *Diamond Paradox* applies and consumers forego search, which mostly harms sellers who would have extracted most of the surplus. Figure (2a) shows that tracking raises sellers' profits otherwise and especially when the market would shut down without tracking for high search costs. Figure (2b) shows the effect of tracking on overall welfare, which depends on the level of search costs as well. There is a wide range of search costs for which foregone profits due to reduced search persistence cannot be offset by the increase in consumer surplus. For intermediate search costs, tracking thus reduces welfare. In fact, this is intuitive as the hold-up problem prevents the realization of matches especially from high value consumers.

5 Extension: Increasing Matching Probability

In markets where products are sufficiently complex and have a variety if differentiated features, the distinction between mass and niche consumers implying $g'(x) < 0$ seems a reasonable assumption. However, if products are more standardized, the major source of different match values may not lie in the nicheness of a consumer's taste. Rather, a consumer's match value will depend on her budget set. For instance, think of consumer electronics like flat screens. While not every brand's flat screen constitutes a match due to different preferences for diameter or energy consumption, the majority of consumers derive utility from all of its technical features. Consequently, their willingness to pay depends on mostly on the available income although they (still) buy the product only if the investigated features meet their individual preferences. Aguiar and Hurst (2007) find that low income consumers do less comparison shopping but still spend more time on shopping in total, suggesting that they spend more time than high income consumers on studying the products they buy. Given high income consumers inspect a product's features less carefully than low income consumers, they might be more likely to find a product suitable. In this section, I therefore analyze the case of a weakly increasing matching probability function $g(x)$ with $g'(x) \geq 0$.

Many of the arguments and results are either identical or simply a reversed version of those made in the previous section. In those cases, explanations and proofs are presented in an abbreviated form with complete references to the previous section. I first show under which conditions there exists an equilibrium that exhibits a decreasing price path. Second, I show that under those conditions, the equilibrium is unique. Moreover, I derive the equilibrium without tracking and compare market outcomes.

5.1 Search With Tracking

For $g'(x) \geq 0$, consumers with a high conditional match value are more likely to encounter a match early. Hence, intuition suggests that lower prices are charged from consumers with longer search histories in equilibrium. To construct such an equilibrium, suppose that consumers' expectations satisfy $p_1^e \geq p_2^e \geq ... \geq p_N^e$.

If consumers expect a decreasing price path, they might prefer to continue searching despite available matches at previous sellers. As the analysis of alternative expectations in the case of $g'(x) < 0$ has shown, this potentially leads to a new category of demand from consumers who might return to a match after sampling additional sellers. However, if potential gains from lower prices at forthcoming sellers are sufficiently smaller than search costs, continuing to search despite a match is not worthwhile. Then, consumers follow $\mathcal{R}^*$ and expected history-dependent demand can be characterized by equation (2) derived in section 3. In contrast to $g'(x) < 0$, demand from consumers with long search histories is more elastic if consumers follow $\mathcal{R}^*$ and types with a low conditional match value search longer on average. I thus obtain the opposite result of lemma 2.

Lemma 6. *Suppose that consumers always use the stopping rule $\mathcal{R}^*$. Then, the sequence of profit-maximizing prices is weakly decreasing and unique.*

As in the proof of lemma 2, uniqueness is due to log-concavity of the demand function, which is preserved under $\mathcal{R}^*$ for any seller along a consumer's search path. Lemma 6 does not yet complement proposition 1 for $g'(x) \geq 0$. The question remains under what conditions the continuation value from search after a successful match is always negative, thus rendering $\mathcal{R}^*$ indeed the optimal stopping rule?

Potential gains from lower prices depend on the changes in the elasticity of expected demand. Denote the decreasing sequence of optimal prices if consumers follow $\mathcal{R}^*$ by

$$\{p_k^*\}_{k=1,..,N}, \; p_k^* \geq p_{k+1}^* \; \forall \; k < N.$$

Note that in any PBE where $\mathcal{R}^*$ is the optimal stopping rule, consumers' expectations must be correct and thus satisfy $p_k^e = p_k^*$. The stopping rule $\mathcal{R}^*$ is optimal always only if the continuation value from search conditional on an available match is weakly negative for all types and for all possible search histories. Extending previous notation, denote the continuation value conditional from sampling seller k conditional on an existing match and known type x_i by $V_k(x_i)$. Formally, $V_k(x_i) \leq 0 \; \forall x_i \in X, k \leq N$ *iff*

$$g(x)\big(p_k^* - p_{k+1}^e\big) < s \; \forall x \in X, \; k \leq N. \tag{8}$$

Given any level of search costs $s > 0$, an upper bound $\hat{\Delta} > 0$ for the slope of $g(x)$ exists such that condition (8) holds for all matching probability functions with $g'(x) < \hat{\Delta}$. Consequently, $p_k^e = p_k^+ \; \forall k$ if $g'(x) < \hat{\Delta}$. Lemma 7 summarizes these findings.

Lemma 7. *There always exists a $\hat{\Delta} > 0$ such that under expectations satisfying $p_k^e = p_k^* \; \forall \; k \leq N$, the stopping rule $\mathcal{R}^*$ is optimal and $\{p_k^*\}_{k=1,..,N}$ constitute unique equilibrium prices for any matching probability function with $0 \leq g'(x) < \hat{\Delta} \; \forall \; x \in X$.*

The stopping rule $\mathcal{R}^*$ thus leads to the price sequence $\{p_k^*\}_{k=1,..,N}$, which implies that $p_k^e = p_k^*$, rendering $\mathcal{R}^*$ indeed optimal for $g'(x) < \hat{\Delta}$. However, other equilibria could exist for alternative consumer expectations, making some types adopt a stopping rule different from $\mathcal{R}^*$. However, it is possible to show that any expectations which are different from $\{p_k^*\}_{k=1,..,N}$ cannot constitute a PBE.[14] The steps towards this result are similar to those made in the previous section. First, observe that

Lemma 8. *In any PBE, expectations satisfy $p_k^e \geq p_k^* \; \forall \; k \leq N$ if $\hat{\Delta} > g'(x)$.*

Intuitively, the statement holds for the following reason. Begin with the last sellers whose price is expected to be below p_k^*, i.e. $p_k^e < p_k^*$. If despite alternative expectations, consumers' stopping behavior at previous sellers remains unchanged, seller k would maximize profits by deviating to $p_k = p_k^* > p_k^e$ as this price is profit-maximizing conditional on $\mathcal{R}^*$.

[14] As it is the case in most search models, there always exists an equilibrium in which consumers expect arbitrarily high prices and no consumer searches.

If instead, consumers' stopping behavior changes due to alternative expectations, it can affect only expected demand types with $x > \bar{x}$ for seller k, where $\bar{x} > p_k^e$. The threshold $\bar{x}$ characterizes the type whose matching probability $g(\bar{x})$ is too low to make sampling seller k despite available matches worthwhile. Consequently, expected demand from types $x_i > \bar{x}$ can at most in- but not decrease compared to demand that arises under $\mathcal{R}^*$. Since raising the price to $p_k^* > p_k^e$ leads to higher profits even in the absence of this extra demand by lemma 6, setting a price higher than p_k^e must constitute a profitable deviation when demand from types $x_i > \bar{x} > p_k^e$ as well. The same argument can be applied to any seller k' whose expected price is supposed to satisfy $p_{k'}^e < p_{k'}^*$, leading to the above lemma. As in the analysis in section 3, the lower bound on consumers' expectations leads to equilibrium uniqueness.

Proposition 8. *In any PBE, expectations must satisfy $p_k^e = p_k^* \ \forall\ k \le K^*$. Hence, the equilibrium with the increasing price path characterized by $\{p_k^*\}_{k=1,..,N}$ is unique.*

The proof proceeds along the same lines used in proposition 1. As the first seller knows that prices from all forthcoming sellers are expected to be higher than his own price if he charges a price in the neighborhood of p_1^*, he has local monopoly power over its demand because consumers would follow $\mathcal{R}^*$. Since p_1^* is the profit-maximizing price under $\mathcal{R}^*$, the first seller will always find it optimal to set a price equal to p_1^+. Given that the first seller sets $p_1 = p_1^*$, seller 2 can makes consumers follow $\mathcal{R}^*$ (and thus obtain local monopoly power) by setting $p_2 = p_2^*$ for the same reasoning. By repeatedly applying this argument for all forthcoming sellers, one obtains the above result.

Search persistence. As in the previous analysis, consumers sample up to K^* sellers if $V_{K^*}\big(s\big) \ge 0$ and $V_k\big(s\big) \ge 0 \ \forall\ k < K^*$ given equilibrium prices. In fact, the fact that prices are decreasing instead of increasing does not change the computation of K^*. That is, $K^* \in \mathbb{K}$, where $\mathbb{K}$ is given by (5). Since $\hat{K}(s') \subseteq \hat{K}(s)$ for $s' > s$ holds as well, K^* is decreasing in s, resulting in a set of disjoint search costs intervals for different K^*.

If continuation values are not always decreasing for higher search histories, there might be jumps in K^* such that search persistence decreases by more than one seller at some search cost threshold. Using previous notation, let $\hat{s}_k(K) \in \{s : V_k(K, s) = 0\}$ define the search cost threshold above which sampling k is not worthwhile for a given search persistence K. For non-decreasing continuation values, there exists a K such that $\hat{s}_k(K) \neq \underline{\hat{s}}(K) \in \min\{\hat{s}_k(K)\}_{k=1,..,K}$. Then, $K^* = K$ only if $s \le \underline{\hat{s}}(K)$ and $K^* = j$ for $\hat{s}_{j'}(j) \ge s > \hat{s}_j(K)$ where seller j's threshold equals $\hat{s}_j(K) = \underline{\hat{s}}(K)$. The next threshold $\hat{s}_{j'}(j)$ is obtained from $\hat{s}_{j'}(j) \in \min\{\hat{s}_k(j)\}_{k=1,..,Kj}$ and specifies that $K^* = j'$ for $s > \hat{s}_{j'}(j)$ and so forth.

5.2 Search Without Tracking

Deriving uniform price equilibrium and its uniqueness under search without tracking is identical to the case of $g'(x) < 0$. The reason is that due to the uniform price, consumers follow the same stopping rule $\mathcal{R}^*$. Hence, the seller's problem is completely characterized by (7) with expected demand given by equation (2). Thus:

Lemma 9. *A unique uniform price equilibrium exists under random search.*

The only difference compared $g'(x) < 0$ is that demand is now more elastic for higher degrees of search persistence. That is:

Proposition 9. *Let $p(K^*)$ be the unit random search price if consumers maximum willingness to search is K^*. Then,*

$$p(K_2^*) < p(K_1^*) \text{ if } g'(x) > 0 \text{ and } K_2^* > K_1^*$$

Notice that the intuition for the result is a simple reversion of the statement before. A higher search persistence by consumers increases the probability that a consumer has a long history. Since probabilities for all feasible search histories must add up to one, sellers put less weight consumers with short histories and more weight on consumers with long search histories if K^* increases. Since demand from consumers with longer search histories is more price elastic for $g'(x) \geq 0$, the profit maximizing price decreases.

Search persistence. For a constant price, the continuation value from search decreases for longer search histories. This is because a consumer's expected type $\mathbb{E}_h[x]$ decreases in h and so does the instantaneous expected surplus from the next seller, $g(x)(x-p)$, if $g'(x) \geq 0$. Since consumers become increasingly pessimistic for higher search histories, continuation values are decreasing in h, implying that K^* decreases gradually. That is, consumers sample $K^* = K$ sellers if $s > \hat{s}(K)_K$ where $s_K(K) \in \{s : V_K(p(K), s) = 0\}$. Since $p(K)$ is decreasing in K, dynamic search inconsistencies as in the case of $g'(x) < 0$ cannot emerge. Hence, the market remains active for all $s < s_1(1)$. Moreover, consumers never have to choose mixed stopping rules.

5.3 Comparative Statics of Tracking

Since search persistence decreases gradually with search costs and $V_1(p(1), s) = V_1(p_1, s)$, it follows that consumers search both under tracking and no tracking *iff* $s < \hat{s}_1(1)$. Contrary to the case of $g'(x) < 0$, there thus exists no level of search costs for which tracking must lead to higher profits and consumer surplus due to a market breakdown without tracking. The comparison of prices is immediate after switching the sign of $g'(x)$ in the proof of proposition 5:

Corollary 2. *For any search persistence $K^* > 1$,*

$$p_1 \geq p(K^*) \geq p_{K^*}.$$

with strict inequality if $g'(x) > 0$.

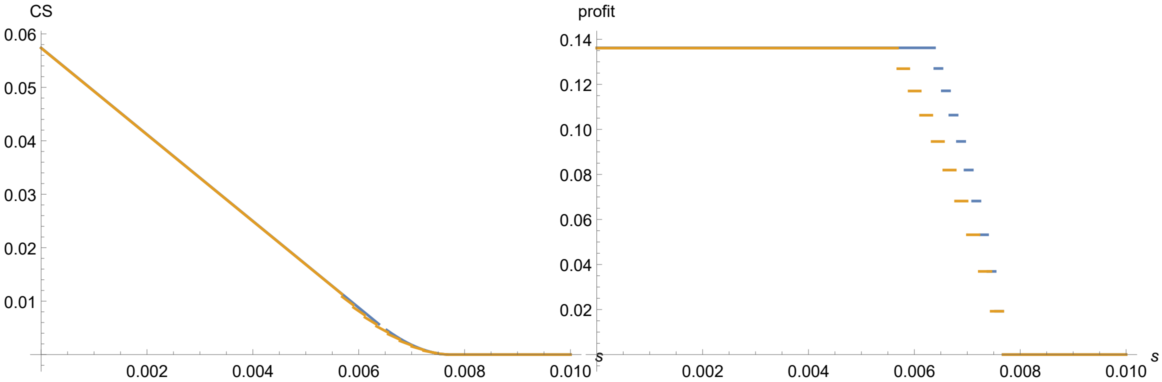

(a) Consumer surplus

(b) Profits

Figure a) shows consumer surplus as a function of search costs. Figure b) shows total profits as a function of search costs. Orange represents search without tracking, blue search with tracking.

Figure 3: Search costs, consumer surplus and profits

By corollary 2, sellers maximize profits by charging search history dependent prices, which differ from the uniform price without tracking. That is, reduced asymmetric information due to tracking enables sellers to extract more expected surplus from a consumer with a particular history h. Unless $g'(x) = 0$, I obtain:

Lemma 10. *If K^* is weakly higher under tracking than under no tracking, sellers' profits are strictly larger under search with than under search without tracking.*

To obtain specific results on the effect of search costs, I again impose additional structure. Consider a linear matching probability function$g(x) = 0.1x$, which is increasing in x. As before, let the type x be uniformly distributed on $[0, 1]$.[15] The total number of sellers is held constant at $N = 10$ but consumers may sample only $K^* \leq N$.

The computation proceeds as follows. First, I compute prices under search with tracking for any history h and under search without tracking for any search persistence K^*. Second, by calculating continuation values for every possible history, I obtain the optimal search persistence K^* as described in section 5.1 and 5.2 for every level of search costs. Recall that the decreasing price path under search with tracking may not constitute an equilibrium if search costs are too low or potential gains from price savings too high. While not displayed here, the maximum price difference a consumer can expect amounts to roughly $\Delta_p = 0.005$. Since the highest matching probability a consumer might have equals 0.1, condition 8, which ensures that a unique equilibrium exists, is satisfied for $s > 0.0005$. Figure (3a) shows that tracking hardly affects consumer surplus. That is, gains for longer searching consumers from lower prices are offset by higher starting prices for all consumers. In contrast, profits can be much larger due to tracking as shown in figure (3b). In fact, this is because K^* decreases faster if search is without tracking.

[15]The slope is chosen to be small to ensure equilibrium existence.

While under tracking, consumers with long search histories may find continuing to search worthwhile such that match values are realized, the price without tracking may prevent them from search. As the surplus left to consumers even if they continue searching is small, there is only a marginal effect on consumer surplus.

6 Application: Endogenous Tracking

In this section, I apply the consumer search framework to study whether tracking arises endogenously. For this purpose, I consider a consumer's dynamic choice about preventing tracking. Since the processing of personal data seems to be the default on the internet for its use goes far beyond price discrimination only, a seller's choice about tracking is only about deciding on whether to use the available data for price discrimination. Note that this implies a seller decides about using tracking at the price setting stage. Hence, the problem is simply part of his profit maximization. More precisely, if anything but a search history-independent price is optimal, it must hold that using tracking is the dominant strategy.

To model a consumer's tracking choice, I assume that every time before she samples a seller, she can either disclose her entire search history, i.e. allow tracking (T) or not disclose her search history and thus not allow tracking (NT). Search histories contain the number of all sellers previously visited, independently of whether the consumer had chosen NT or T when sampling previous sellers.[16] In reality, consumers usually also have the option to erase their histories, for example by deleting cookies. In the discussion section, I show that the predictions I obtain are robust to this extension if sellers can distinguish between a consumer who deletes her cookies and a consumer who just started searching for a particular product. Briefly, this assumption is motivated by the observation that by deleting cookies, consumers erase their entire search profile. However, a consumer who has not deleted cookies but just started searching for a product should still have an "unrelated" search history. The signal sent to sellers when deleting cookies is then identical to choosing NT and, thus, is redundant. Denote a consumer's tracking decision by $d \in \{T, NT\}$.

Timing. At first, sellers set prices for every possible search history and choice of d and nature draws each consumer's type. Consumers search by sampling sellers sequentially at a cost s per seller. Prior to each search attempt, consumers choose d. Sellers observes a consumer's search history h if $d = T$ and nothing but $d = NT$ otherwise. The consumer observe the price conditional on her disclosure strategy and search history as well as her match utility. Lastly, she decides whether to buy, to return to a previous seller, to

[16] Web-browsing with the "do not track" request option, where cookies are still stored on the consumer's device but simply not processed, fits this assumption fairly well. However, results are robust to the modified assumption that search histories contain only the number of sellers for which tracking had been enabled. This is because the search history a consumer can disclose would still be weakly increasing in the number of sellers and thus not affect the price path with tracking.

continue or to stop search.

Equilibrium concept. As before, the equilibrium concept is Perfect Bayesian Nash equilibrium. Extending the strategy space by the choice of d implies additional PBE conditions which are not present in the baseline mode. First, consumers choose d in order to maximize their expected surplus. Second, sellers' beliefs about a consumer's search history must be consistent with the consumers' disclosure strategy. Third, consumer's beliefs about prices must be consistent with sellers' equilibrium pricing strategies, and, thus, with their own disclosure strategy.

I analyze the cases of $g'(x) < 0$ and $g'(x) \geq 0$ separately and begin with $g'(x) < 0$. Denote by $p(NT)$ the uniform price set by sellers upon observing the choice NT. For prices under tracking, use the previous notation p_k with the index indicating the seller's position in the consumer's search process. Note that the choice about $d \in \{T, NT\}$ is without commitment and only affects the information revealed to the next seller while the prices from additional sellers remain unaffected. Hence, the consumer's decision to disclose her history is purely myopic as it only depends on the difference between the next price she can expect under tracking p^e_{h+1} and no tracking $p^e(NT)$. If $p^e_k = p^e(NT)$ for some consumer with history $h = k - 1$, I impose that a consumer stays with the default option, which is search with tracking. This tie-breaking rule is without loss of generality as the alternative rule would imply the same equilibrium outcome. Since consumers choose $d = T$ prior to sampling seller k if only if $p^e(NT) \geq p_k$, it is sufficient to restrict attention to single cut-off strategies, where such a strategy is defined as follows:

$$\textit{Choose } NT \;\; \forall\, h \geq \hat{h} \in \mathbb{N} \tag{9}$$

For brevity, denote the above defined single cut-off strategy by $\hat{h} \leq K^*$. To see why this restriction does not constrain equilibrium strategies, suppose that $p^e(NT) < p^e_k$ such that a consumer chooses $d = NT$ at seller k. Since her search history at any forthcoming seller will be $h > k$, forthcoming sellers' prices always satisfy $p^e(NT) < p^e_j \;\forall\, j > k$. Consequently, $d = NT$ must be optimal for all $h \geq k$ if it is optimal at k.

In any PBE, sellers anticipate the equilibrium strategy $\hat{h}$ and can thus condition $p(NT)$ on $h \geq \hat{h}$ when observing $d = NT$. Denote the profit-maximizing price conditional on the cut-off strategy $\hat{h}$ by $p(\hat{h})$. Then, $p(NT) = p(\hat{h})$ in equilibrium. Recall that by previous notation, $p_{\hat{h}+1}$ denotes the price set by a seller who observes a browsing history $\hat{h}$ and thus knows that his position in the consumer's search process is $\hat{h} + 1$. Without imposing equilibrium strategies yet, the following lemma compares prices with tracking and without tracking for arbitrary single cut-off strategies $\hat{h}$.

Lemma 11. *Given $\hat{h}$, there always exists a unique optimal price $p(\hat{h})$ with*

$$p(\hat{h}) > p_{\hat{h}+1} \;\forall\, \hat{h} < K^* - 1,$$

and $p(\hat{h}) = p_{\hat{h}+1}$ for $\hat{h} = K^*$.

The intuition behind lemma 11 is the following. When hiding the search history for all histories $h \geq \hat{h}$, sellers observing $d = NT$ attach positive probabilities on all histories $h \geq \hat{h}$ and zero probability on $h < \hat{h}$. Most importantly, $\mathbb{P}\big(h > \hat{h}\big) > 0 \ \forall \ \hat{h} < K^* - 1$, and hence the optimal price conditional on observing $d = NT$ is chosen with respect to a weighted demand function that is always less elastic than expected demand from a consumer disclosing $h = \hat{h}$ (implying $\mathbb{P}(h = \hat{h}) = 1$).

By lemma 11, a consumer whose search history equals the cut-off history is charged a lower price if she allows tracking ($d = T$) than if she does not ($d = NT$). However, by construction, a consumer with a history of $h = \hat{h}$ chooses $d = NT$, implying that profitable deviations exist at least for some consumers. The uniqueness result in the following proposition is an immediate consequence of this contradiction.

Proposition 10. *There always exists a unique PBE with the disclosure strategy $\hat{h} = K^*$ and a conditional no tracking-price $p(NT) = p_{K^*}$.*

Proposition 10 states that the search history is always disclosed in the unique equilibrium, leading to unrestricted tracking and price discrimination. Existence can be shown by means of an example. Simply consider an equilibrium with search history-based prices $p_1 < p_2 < ... < p_{K^*}$, a disclosure strategy $\hat{h} = K^*$ and a no tracking-price $p(NT) = p_{K^*}$. Since K^* is the maximum number of sellers a consumer is willing to sample, $\hat{h} = K^*$ means that no actively searching consumer chooses $d = NT$ and sellers should never observe NT. Denote by $\mu(h) \in [0, 1]$ a seller's out-of-equilibrium belief that the search history of a consumer having chosen NT equals h. Suppose it satisfies $\mu(h) = 0 \ \forall h < \hat{h}$ such that $\mu(K^*) = 1$. Since p_{K^*} maximizes profits conditional on $\mu(K^*) = 1$, sellers have no incentive to deviate from $p(NT) = p_{K^*}$. Besides, consumers have no incentive to deviate to another disclosure strategy $\hat{h}' < \hat{h}$ since $p(NT) \geq p_k \ \forall \ k \leq K^*$.

The uniqueness result is based on an unravveling mechanism similar to Milgrom and Roberts (1986). For any alternative cut-off strategy $\hat{h} < K^*$, sellers' beliefs must satisfy $\mu(h) = 0 \ \forall \ h < \hat{h}$. Hence, the optimal price conditional on observing NT satisfies $p(NT) = p(\hat{h}) > p_{\hat{h}+1}$ by lemma 11. Since consumers with a search history $h = \hat{h}$ can obtain the price $p_{\hat{h}+1}$ by allowing tracking prior to sampling seller $k = \hat{h}+1$, they always have an incentive to deviate from any cut-off strategy $\hat{h} < K^*$.

6.1 Increasing Matching Probability

If $g'(x) \geq 0$ but not too large, prices are monotone decreasing in search histories as shown in section 5. Thus, it follows that the optimal disclosure strategy belongs to the set of single cut-off strategies as well. However, the reverse pricing pattern requires to slightly adjust the notion of single-cut-off strategies, abbreviated by $\check{h}$:

$$\textit{Choose NT} \ \ \forall \ h \leq \check{h} \tag{10}$$

Denote by $p(\check{h})$ the seller's optimal price conditional NT and the consumer cut-off strategy $\check{h}$. Analogously to lemma 11, one can show that:

Lemma 12. *Given $\check{h}$, there always exists a unique optimal price $p(\check{h})$ which satisfies*

$$p(\check{h}) > p_{\check{h}+1} \ \forall \ \check{h} > 0,$$

and $p(\check{h}) = p_{\check{h}+1}$ for $\check{h} = 0$.

The distinction between a a search history h and the corresponding position $h+1$ in the search process of a seller observing h is again crucial to understand the implications of lemma 12. For illustration, suppose that $\check{h} = 1$ implying that consumers choose $d = NT$ if $h \in \{0, 1\}$. In any PBE, sellers would know that $h \in \{0, 1\}$ if $d = NT$ and $h \ngeq 1$. By lemma 12, the resulting optimal price $p(NT)$ then exceeds p_1, the price sellers would set if they knew that $h = 1$ with certainty. However, any consumer with a history of $h = 1$ can choose $d = T$ and convey their history to sellers, thereby obtaining a lower price than $p(NT)$. The example shows that $\check{h} = 1$ cannot be an equilibrium strategy for consumers. In fact, the unraveling argument applies again for any cut-off strategy $\check{h} > 0$ such that complete tracking remains as the unique equilibrium:

Proposition 11. *There always exists a unique PBE with the disclosure strategy $\check{h} = 0$ and a conditional no tracking-price $p(NT) = p_1$.*

While the unraveling mechanism illustrated in the above example is reversed compared to proposition 10, the proof of proposition 11 is otherwise identical. For any cut-off strategy $\check{h} > 0$ where NT is chosen from consumers with strictly positive search histories, there always exists consumers who are better off from tracking even though their history belongs to $h \leq \check{h}$. Hence, only $\check{h} = 0$ constitutes an equilibrium strategy.

The major implication of propositions 10 and 11 is that there is no privacy in the market because consumers rationally approve tracking at all stages during the search process. Since the consumers' best-response to price discrimination is a myopic decision based merely on the very next seller's price, the equilibrium outcome need not be consumer surplus maximizing.

7 Discussion

In this section, I discuss the robustness of my findings via several extensions to the model. Some extensions constitute an entirely new model for future research which cannot be discussed in every detail here. In this case, I only provide a brief discussion of what changes to expect. Lastly, I explain how the model contributes to the discussion of whether random search can be stable (Armstrong, 2017).

No commitment

Since consumers search with free recall, sellers might not only be interested in discriminating between consumers, but also to discriminate a consumer based on whether she arrives for the first time or whether she is returning. Indeed, Armstrong and Zhou (2010) focus on this issue. In my baseline model, the timing does not allow for within consumer price discrimination. Instead, I implicitely assume that sellers can commit to a price they will charge from returning consumers under the constraint that this price is equal to the price charged at the first encounter. Indeed, this is not without loss of generality. Due to free recall, a multiplicity of equlibria might arise without commitment since consumers could form any beliefs about the return-price. However, the commitment assumption could be relaxed by introducing a small but positive cost ϵ for returning to a previous seller. This is because when a consumer with history h decides to return to some seller $k < h$, she does so only if she expects $v_{ik} > p_k^R + \epsilon$, where p_k^R is the expected return-price. Hence, sellers have an incentive to raise their price at least by ϵ, making returning not worthwhile for some consumers. This argument holds for any $p_k^R < \bar{v}$ such that the arising hold-up problem prevents any consumer from returning if sellers have no commitment power. The reason why the *Diamond Paradox* arises in the market for returning consumers but not in the market for "fresh" consumers is because only in the former, consumers already know their willingness to pay, rendering their return decision informative for sellers.[17]

Heterogeneous search costs

Another dimension in which consumers naturally differ from one another might be the individual search cost. Should we expect countervailing effects from introducing search cost heterogeneity into the current framework, mitigating the search dynamics and welfare implications derived? Heterogeneous search costs imply that consumers differ with respect to their search persistence K^*. Under search with tracking, prices do not depend on K^*. Thus, there are no new qualitative effects of search cost heterogeneity, despite leading to heterogeneous stopping by consumers.

Under search without tracking, heterogeneous search costs mitigate the dynamic search inconsistency problem. Since for any marginal increase in search costs, there is only a marginal consumer reducing her search persistence, there are no jumps in K^* and prices react smoothly to changes in search costs. That is, there is again no effect of search cost heterogeneity except for making mixed stopping rules disappear. The inefficiency problem due to market inactivity as stated in proposition 6 also persists under search cost heterogeneity. While indeed some consumer will always search for reasonable levels of search costs, those with low search costs search particularly long, thereby driving up the price without tracking. Hence, consumers with high search costs abstain from

[17] This point is also made in Armstrong and Zhou (2010)

search entirely, driving up prices even more and exlcuding additional consumers from participating in search.

Deleting cookies

By deleting cookies consumers might be able to reset their search history to $h = 0$ and thereby trick sellers. However, it is possible to show that deleting cookies and not allowing tracking are equivalent with respect to their signal under some mild assumptions. These are: (1) cookies saved on a device cannot be manipulated but only erased completely or not at all. While a minority of internet users might be capable of deleting only particular cookies from their computer, the majority of users is restricted to the standard software which typically enforces this all or nothing property. And (2), sellers can distinguish between a consumer who deleted all her cookies and a consumer who only began searching for a product.

Note also that this second assumption can be derived from more fundamental assumptions about online browsing. A consumer who only begins searching for a particular product still has a non-empty browsing history including search for other product categories and various online services. Denote this "extended" browsing history by $\mathcal{H}$. The browsing history $\mathcal{H}$ is empty only if a consumer deletes her cookies and otherwise satisfies $\mathcal{H} \notin \emptyset$. In addition, note that $h = 0 \not\Rightarrow \mathcal{H} \in \emptyset$.

Denote the choice of deleting cookies by $t \in \{KC, DC\}$, where DC denotes the choice to delete cookies while K refers to keeping them. Focusing on sellers' belief about a consumer's (history-) type induced by the signal $t = DC$, it turns out there is no difference to the signal $d = NT$. Under NT, sellers cannot observe h. Under $t = DC$, sellers observe $h = 0$ but know that the consumer has deleted cookies since $\mathcal{H} \in \emptyset$. Hence, they know that they do not know the true h, which is equivalent to not knowing h at all.

Denote by $\{d, t\}$ the consumer's action tuple regarding the decision to allow tracking and keep her cookies. Due to identical signaling effects, it holds that sellers' beliefs satisfy:

$$\mu(h|T \wedge DH) = \mu\big(h|NT \wedge DH) = \mu(h|NT \wedge KH) \;\forall\; h.$$

Since the sellers' beliefs determine prices, these actions are thus payoff-equivalent from the perspective of consumers as well.

Now suppose that a fraction of consumers indeed searches with a new device and that her browsing history cannot be distinguished from a consumer who deletes her cookies. Equivalently, there might be a fraction of consumers who always disable tracking by default because they have a strong preference for privacy. If the fraction of those consumers is sufficiently large, the equilibrium from proposition 10 cannot be sustained since sellers must attach positive probability on shorter search histories conditional on observing NT (or equivalently, DC).

Notably, the first order condition of a seller observing NT shows that the optimal no tracking price $p(NT)$ is weakly decreasing in the cut-off strategy $\hat{h}$.[18] Besides, $p_{\hat{h}+1}$, the price a consumer with a history $h = \hat{h}$ would obtain if she chose T, is increasing as shown by lemma 2. Denote by α the fraction of consumers who disable tracking by default/ begin search with new devices. It follows that:

Corollary 3. *For $\alpha > 0$ but not too large, there exists a unique PBE with the disclosure strategy $0 < \hat{h} \leq K^*$ and the no tracking price $p(NT)$ solves*

$$p(NT) = \arg\max_p \big((1-\alpha) \sum_{k=\hat{h}+1}^{K^*} \phi_k D_k(p) + \alpha \sum_{k=1}^{K^*} \phi_k D_k(p)\big)p.$$

Stable random search

According to Armstrong (2017) random search in the WAR model is unstable. He argues that random search from the part of consumers depends crucially on consumers' expectations about other consumers' perfectly random search behavior. If instead, one seller S becomes more salient than its competitors such that both consumers and sellers should expect that an arbitrarily small but positive mass of consumers is more likely to sample S first, S will optimally set a lower price than its competitors, as shown by Armstrong and Zhou (2011) or Zhou (2011). If consumers are free to choose at what seller to begin searching while all sellers' products are ex ante identical, this creates an incentive for all consumers to begin searching at seller S. Consequently, the seller sampled first is not a random choice and search becomes partially ordered. Similarly, consumer beliefs about which seller to sample afterwards can tip easily and so can beliefs about all sellers' position in the search process. Then, random search becomes perfectly ordered as the same argument applies to the second seller etc.

The market tipping property of random search prevails because sellers who are visited first have an incentive to price lower than sellers who are sampled later in the search process. In the current framework instead, this is not true if $g(x)$ is weakly increasing, suggesting that random search would be stable. Suppose that a positive mass of consumers would not search randomly but begin with a particular seller M such that M has a larger share of consumers with shorter search histories. As $g(x)$ is increasing and high types are less likely to be among consumers with longer search histories, seller M has an incentive to raise its price above the uniform price charged by others.[19] Hence, the more salient seller M sets the highest price if $g'(x) > 0$. This prevents the search market from tipping as due to its higher price, consumers would avoid searching for seller M first, including those that were expected to sample it first on purpose. Consequently, consumers prefer to search randomly over coordinating on more salient sellers.

[18] See the proof of lemma 11

[19] By not raising the price by too much, consumer's stopping rule would remain unaffected so that the stopping rule remains unaffected at least for small price deviations.

8 Conclusion

Most of users' activities across the internet are tracked by third parties. Accessing these browsing data is particularly attractive to sellers if the average search behavior varies across different consumer types. Then, tracking a consumer's search path enables sellers to learn about a consumer's willingness to pay and gives rise to search path-dependent price discrimination. This paper presents a rich and tractable framework integrating consumer heterogeneity and tracking into a model of sequential consumers search to address the major implications of tracking for market outcomes.

First, I show that, in the unique equilibrium, tracking implies search history-dependent pricing. Specifically, prices increase in the order of search if the difference between consumer types is the peculiarity and nicheness of their taste. Since niche consumers are more likely to search longer, demand from consumers with long search histories is less elastic. Consequently, sellers set higher prices conditional on observing longer search histories.

Second, I compare the market outcome under search with tracking with the unique equilibrium under search without tracking to evaluate its welfare consequences. Because initial prices are lower while later prices are higher than the price under search without tracking, the surplus of niche consumers decreases while the surplus of mass consumers increases due to tracking. Besides, overall welfare effects depend how tracking affects consumers' search persistence, which, in turn, depends on the level of search costs. For a wide range of intermediate search costs, consumers sample more sellers in the absence of tracking because the average no-tracking price makes continuing search with a long search history more attractive. This may cause welfare losses due to forgone matches, particularly at the cost of sellers who would have extracted most of the matching surplus. However, tracking increases welfare if search costs are very low such that search persistence remains unaffected and if search costs are very high. In the latter case, this is because low prices conditional on short search histories ensure that consumers have an incentive to begin searching, thus keeping the market active. In the absence of tracking instead, the no-tracking price is often too high, thus leading to a market breakdown for the same level of search costs. Perhaps surprisingly, consumers may always be better off from tracking whereas sellers make less profits at least for some search costs.

Third, I investigate whether tracking prevails endogenously when consumers can dynamically opt out from tracking. I find that, in the unique equilibrium, consumers always prefer to disclose their search history as the price conditional on hiding it always exceeds the price at least some consumers could obtain after disclosing it. Therefore, the entire search history is disclosed in equilibrium. Even though the equilibrium is unique, the full tracking prediction is not cast in stone. As discussed in the previous section, partial tracking, where consumers disclose their search history only up to some threshold, obtains if a positive mass of consumers always disables tracking by default. Besides, the endogenous tracking result is interesting not only because it explains why many internet users do not prevent tracking, but also because it has important implications for policy makers. That is, if tracking is surplus-increasing, no intervention is necessary because tracking prevails even though it entails increasing prices.

Often, sellers refrain from personalized pricing because of consumers' prejudices against price discrimination or legal uncertainty. Then, improved targeting constitutes an alternative practice to capitalize on traceable search histories. That is, a seller possessing multiple products of the same category might be able to use information conveyed through search histories to offer more suitable products to individual consumers. Following Johnson and Myatt (2006), a seller's product choice could be integrated in this paper's framework by allowing sellers to rotate the matching probability function. Again, it would be interesting to examine whether the profit-maximizing design induces inefficiently low search persistence or whether it can be welfare increasing as well.

Finally, Turow et al. (2009) find that the majority of consumers oppose personalized pricing, thus confirming the anecdotal evidence that significant prejudices against tracking still prevail. In the light of the fact that tracking often has desirable consequences for consumer surplus and welfare, the question of where this negative view comes from deserves more attention. If consumers obtain additional utility from anonymity, avoidance of tracking and personalized pricing may, of course, be welfare maximizing despite the foregone matching surplus. If, however, the preference for anonymity is based on false beliefs about the consequences of tracking, consumers may be harmed from being misinformed.

9 Appendix

9.1 Proofs of Section 2.3

Proof of Lemma 2

Proof. In **Part I**, I show that if consumers follow the stopping rule as specified in lemma 1, there is a unique sequence of increasing prices that maximize profits, denoted by $\{p_k^+\}_{k=1,...,N}$. Under the equilibrium stopping rule $\mathcal{R}^*$, consumers always buy if $v_{ik} > p_k$. Taking all other forthcoming sellers' expected prices as given, a price by seller k can induce a different stopping rule for at least some types $x \in X$ if $p_k > p_{k+1}^e + \delta, \ \delta > 0$ such that those types find continuing to search worthwhile despite $v_{ik} > p_k$. In **Part II**, I show that any price $p_k > p_k^+$ inducing an alternative stopping rule cannot be profitable if it is not profitable under $\mathcal{R}^*$.

<u>**Part I**</u>

Existence. Define $q_k(p) := -\frac{D_k(p)}{D_k'(p)}$ such that the FOC writes: $p = q_k(p) \ \forall \ k = 1, ..., N$. Next, write

$$q_k(p) = \frac{\int_p^1 g(x) f_k(x) \mathrm{d}x}{g(p) f_k(p)}$$

and observe that $q_k(p)$ is continuous $\forall p > 0$ and satisifes: $\lim_{p \to \underline{v}} q_k(p) > 0$ and $\lim_{p \to 1} q_k(p) = 0$. Hence, there always exists a $p \in \mathbb{R}$ that solves $p = q_k(p)$.

Uniqueness. Define $\theta_k(x) := g(x) f_k(x) \ \forall \ k \geq 1$. Then, expected demand at seller k conditional on observing a history $h = k - 1$ writes:

$$D_k(p) = \int_p^1 \theta_k(x) \mathrm{d}x = \Theta_k(1) - \Theta_k(p),$$

where $\Theta_k(\cdot)$ is the antiderivative of $\theta_k(\cdot)$. Thus, I can rewrite $q_k(p)$ as:

$$q_k(p) = \frac{\Theta_k(1) - \Theta_k(p)}{-\theta_k(p)}.$$

Notice that that $\theta_k(p)$ is log-concave for all k since first, $\theta_1(x) = f_1(x) g(x) = f(x) g(x)$ is log-concave as multiplication preserves log-concavity and because second, log-concavity of $\theta_k(p)$ implies log-concavity of $\theta_t(p) \ \forall \ t \geq k$. The second statement follows from:

$$\theta_{k+1} = g(x) f_{k+1}(x) = \frac{g(x)\big(1 - g(x)\big) f_k(x)}{\int_{\underline{v}}^1 \big(1 - g(t)\big) f_k(t) \mathrm{d}t} = \frac{1 - g(x)}{C_k} \theta_k(x),$$

where C_k is a constant. The argument in footnote 1 shows that log-concavity of $g(x)$ implies log-concavity of $1 - g(x) \ \forall x \in \{x : g(x) \leq 1\}$. Again, multiplication of two positive and log-concave functions preserves log-concavity. Hence, log-concavity of $\theta_k(x)$ implies log-concavity of $\theta_{k+1}(x)$. Consequently, $\theta_k(p)$ is log-concave $\forall \ k \leq K^*$.

Further, log-concavity of $\theta_k(p)$ implies log-concavity of the anti-derivative $\Theta_k(p)$. Define $\Delta_k(p) = \big(\Theta_k(1) - \Theta_k(p)\big)/O_k$ where $O_k := \int_{\underline{v}}^1 g(x) f_k(x) \mathrm{d}x$ is a normalization to make $\Delta_k(p)$ a probability measure. Then, $\Delta_k(p)$ is log-concave over its positive domain[20]. Rewriting $q_k(p)$ yields:

$$q_k(p) = \frac{\Delta(p)}{-\Delta'(p)} \frac{O_k}{O_k}$$

As $\Delta_k(p)$ is a log-concave probability distribution, $q_k(p)$ is decreasing (see Bagnoli and Bergstrom (2005) and, thus, has a unique fix point.

The unique sequence of prices is increasing in k. Consider again $q_k(p)$ for arbitrary k:

$$q_k(p) \leq q_{k+1}(p) \textit{ iff } \frac{\int_p^1 g(x) f_k(x) \mathrm{d}x}{g(p) f_k(p)} \leq \frac{\int_p^1 \frac{g(x)\big(1-g(x)\big) f_k(x)}{\int_{\underline{v}}^1 \big(1-g(t)\big) f_k(t) \mathrm{d}t} \mathrm{d}x}{\frac{1-g(p)}{\int_{\underline{v}}^1 \big(1-g(t)\big) f_k(t) \mathrm{d}t} g(p) f(p)} \tag{11}$$

$$\Leftrightarrow \int_p^1 g(x)\big(1-g(p)\big) f_k(x) \mathrm{d}x \leq \int_p^1 g(x)\big(1-g(x)\big) f_k(x) \mathrm{d}x \tag{12}$$

$$\Leftrightarrow \int_p^1 g(x)\big(g(p)-g(x)\big) f_k(x) \mathrm{d}x \geq 0 \textit{ iff } g'(x) < 0 \tag{13}$$

Because inequality (13) always holds for $g'(x) < 0$, $q_k(p) \leq q_{k+1}(p)$ holds as well. The following proof goes by contradiction. Hence, assume that $p^*_{k+1} < p^*_k$ for at least some k. Then, 13 implies that:

$$p^*_{k+1} = q_{k+1}(p^*_{k+1}) > q_k(p^*_{k+1})$$

As $q_k()$ is decreasing, it follows from the assumption of $p^*_{k+1} < p^*_k$ that:

$$q_k(p^*_{k+1}) > q_k(p^*_k)$$

But then uniqueness of $p = q_k(p)$ implies $p^*_{k+1} = q_{k+1}(p^*_{k+1}) > q_k(p^*_k) = p^*_k$, contradicting the assumption.

<u>Part II</u>

Now consider seller k setting a price $p'_k > p^+_k$ such that continuing to search is worthwhile for at least some consumers, i.e. $\exists x \in X$ s.th. $V_{k+1}(x_i) > 0$ even if $v_{ik} > p'_k$. If for some type x, $V_{k+1}(x_i) > 0$ even if $v_{ik} > p'_k$, then $x \in \hat{X}_k$, with

$$\hat{X}_k := \{x \in X : g(x)\big(p'_k - p_{k+1}\big) > s\}. \tag{14}$$

[20] Subtracting $\Theta_k(1)$ preserves log-concavity. Further, multiplication with a log-concave function (i.e. (-1)) preserves log-concavity over the composed function's positive domain and O_k is just a constant.

This is because $p^e_{k+2} > p^e_{k+1}$, and hence the surplus from sampling any seller beyond the next seller cannot be positive if the surplus from sampling only the next seller is not positive. By construction, $\exists! \; \bar{x} \in \hat{X}_k$ (with $\bar{x} > p'_k$) s.th. $x \in \hat{X} \Rightarrow x > \bar{x}$. If $x_i > \bar{x}$, consumer i might return with some probability and buy from k. Denote that probability by ρ. It is sufficient to note that $\rho \leq 1 \; \forall x_i > \bar{x}$ with strict inequality for a positive mass of consumers(If $\rho = 1$ always, there would be no reason to continue search in the first place. Hence, it would not correspond to prices that lead to an alternative stopping rule for at least some consumers.) Due to $p'_k > p^e_{k+1} + \delta$, seller k's demand $\hat{D}_k(p)$ under any alternative stopping rule writes

$$\hat{D}_k(p) = \int_p^{\bar{x}} \rho \cdot g(x) f_k(x) \mathrm{d}x + \int_{\bar{x}}^{\bar{v}} g(x) f_k(x) \mathrm{d}x \tag{15}$$

$$< \int_p^{\bar{x}} g(x) f_k(x) \mathrm{d}x + \int_{\bar{x}}^{\bar{v}} g(x) f_k(x) \mathrm{d}x \qquad = D_k(p). \tag{16}$$

where $D_k(p)$ denotes the demand if consumers always followed $\mathcal{R}^*$. As shown in Part I, the problem $\max_p D_k(p)p$ has a unique solution at $p = p^+_k$ and thus:

$$p^+_k \cdot D_k(p^+_k) > p'_k \cdot D_k(p'_k) \; \forall \; p'_k \neq p^+_k$$

Therefore, it follows from inequality (16) that:

$$p^+_k \cdot D_k(p^+_k) > p'_k \cdot \hat{D}_k(p'_k)$$

□

Proof of Lemma 3

Proof. Define $j^* = \sup\{k : p^e_k < p^+_k, k \leq N\}$. By construction, it holds that $p^e_{j^*} < p^e_k \; \forall \; k \geq j^*$, implying that for $p_{j^*} = p^e_{j^*}$, a consumer buys from j^* if $v_{ij^*} \geq p^e_{j^*}$. Notice also that it does not matter whether a consumer with history $h = j^* - 1$ has encountered other matches previously. Since $p^e_{j^*} < p^e_k \; \forall \; k > j^*$, the fact that sampling j^* must have been worthwhile to a consumer implies that $p^e_{j^*} < p_k \; \forall \; k \in \{j < j^* : v_{ij} > 0\}$. Thus, all consumers arriving at seller j^* will buy if $v_{ij^*} > p_{j^*}$ and $p_{j^*} \leq p^e_{j^*} + \delta$, where $\delta > 0$ is determined such that for prices $p_{j^*} > p^e_{j^*} + \delta$, some consumers are induced to continue searching or return to a previous seller despite $v_{ij^*} > p_{j^*}$. Hence, for any price $p_{j^*} \leq p^e_{j^*} + \delta$, seller j^* has full monopoly power. Thus, if the price belongs to this interval, it solves

$$\max_p \Big(pD_{j^*}(p) \Big)$$

Consequently, there always exists a profitable deviation from $p_{j^*} = p^e_{j^*} < p^e_{j^*} + \delta$ if $\partial_p p D_{j^*}(p)|_{p=p^e_{j^*}} > 0$. Let $\tilde{p} = \min_{k<j^*}\{p_k : v_{ik} > p_k\}$ and denote by $\hat{X} := \big\{ x \in X :$

$g(x)\big(\min\{x,\tilde{p}\}-p^e_{j^*}\big) > s\big\}$ the set of consumers whose expected surplus from sampling j^* despite available matches is positive due to potential price savings. If potential savings in the price are sufficiently low, $\hat{X} \subseteq \emptyset$. Then, consumers follow $\mathcal{R}^*$ and seller j^*'s demand is fully captured by the expression given for $D_{j^*}(p)$. From lemma 2, we know that

$$p^e_{j^*} < p^+_{j^*} = \frac{D_{j^*}(p^+_{j^*})}{-D'_{j^*}(p^+_{j^*})} < \frac{D_{j^*}(p^e_{j^*})}{-D'_{j^*}(p^e_{j^*})} \tag{17}$$

where the second inequality holds as $D_{j^*}(p)$ is log-concave under $\mathcal{R}^*$ such that the RHS of the FOC is decreasing in p. Hence, $\partial_p pD_{j^*}(p)|_{p=p^e_{j^*}} > 0$, rendering $p'_{j^*} > p^e_{j^*}$ a profitable deviation.

Next, consider the case where some types' stopping behavior does change such that $\hat{X} \not\subseteq \emptyset$ and denote by $\hat{D}_{j^*}(p)$ the resulting expected demand at seller j^*. As before, denote by $f_k(x)$ the PDF of arriving consumers if they follow $\mathcal{R}^*$ such that $\hat{X} \subseteq \emptyset$ and but by $\hat{f}_k(x)$ the PDF of types if $\hat{X} \not\subseteq \emptyset$. By construction, $\exists\ \underline{x} \in \mathbb{R}$ s.th. $x \in \hat{X} \Rightarrow x \geq \underline{x}$. For $f(x)$ and $\hat{f}(x)$, this implies

$$f_{j^*}\big(x,\mathcal{R}'\big) \geq f_{j^*}\big(x,\mathcal{R}^*\big) \forall\ x \geq \underline{x},$$

with strict inequality for all $x \in \hat{X}$. By continuity of $g(x)$, $\int_{x\in\hat{X}} \mathrm{d}x > 0$ and since $\hat{f}_{j^*}(x)$ is a PDF, it follows that $\hat{f}_{j^*}(x) = z \cdot f_{j^*}(x)\ \forall\ x < \underline{x}$ where $z < 1$ is a normalization. I can thus conclude about the RHS of (17) that:

$$\begin{aligned}
\frac{D_{j^*}(p)}{-D'_{j^*}(p)} &= \frac{\int_p^{\underline{x}} g(x)f_{j^*}(x)\mathrm{d}x + \int_{\underline{x}}^{\bar{v}} g(x)f_{j^*}(x)\mathrm{d}x}{-g(p)f_{j^*}(x)} \\
&< \frac{\int_p^{\underline{x}} g(x)f_{j^*}(x)\mathrm{d}x + \int_{\underline{x}}^{\bar{v}} (1/z)g(x)\hat{f}_{j^*}(x)\mathrm{d}x}{-g(p)f_{j^*}(x)} \\
&= \frac{z\int_p^{\underline{x}} g(x)f_{j^*}(x)\mathrm{d}x + \int_{\underline{x}}^{\bar{v}} g(x)\hat{f}_{j^*}(x)\mathrm{d}x}{-z\cdot g(p)f_{j^*}(x)} \\
&= \frac{\hat{D}_{j^*}(p)}{-\hat{D}'_{j^*}(p)}\ \forall\ p < p^e_{j^*} + \delta
\end{aligned}$$

Combining the above inequality with (17) yields:

$$p^e_{j^*} < \frac{\hat{D}_{j^*}(p^e_{j^*})}{-\hat{D}'_{j^*}(p^e_{j^*})}$$

and thus $\partial_p p\hat{D}_{j^*}(p)|_{p=p^e_{j^*}} > 0$. Consequently, seller j^* would always deviate to a higher price $p > p^e_{j^*}$ if $p^e_{j^*} < p^+_{j^*}$ and hence, $p^e_{j^*} \geq p^+_{j^*}$.

As $p^e_{j^*} < p^+_{j^*}$ is eliminated, another seller $j^* = \sup\{k : p^e_k < p^+_k, k \leq N\}$ might exist. However, by going backwards and using the same argument as above, any $p^e_k < p^+_k$ can be ruled out until $j^* = 1$. $\square$

Proof of Proposition 1

Proof. Begin with the first seller a consumer samples, i.e. $k = 1$. By Lemma 3, $p_k^e \geq p_k^+ \;\forall\; k$ and thus $p_1^+ < p_k^e \;\forall\; k > 1$. Define $\hat{X}_1$ as in (14). Then $\exists!\delta > 0$ s.th. $\hat{X}_1 \subseteq \emptyset$ *iff* $p_1 \leq p_1^+ + \delta$. Note that $\hat{X}_1 \subseteq \emptyset$ induces consumers to adopt the search behavior $\mathcal{R}^*$ and by Part I of lemma 2, $p_1^+ = \text{argmax } D_1(p)p$ where $D_1(p)$ is a hypothetical demand function for $p_1 > p_1^+ + \delta$ that would prevail if continuing to search despite a match were ruled out by assumption (and consumers followed $\mathcal{R}^*$).

It remains to show that $p_1 = p_1^+$ yields larger profits than any price $p_1 > p_1^+ + \delta$, inducing a stopping rule different from $\mathcal{R}^*$ for at least some types. I use the notation of the previous proof and denote the demand that arises under an alternative stopping rule by $\hat{D}(p)$. By Part II of lemma 2, $\hat{X}_1 \not\subseteq \emptyset$ implies that $p_1' D_1(p) > p_1' \hat{D}_1(p_1') \;\forall\; p_1' > p_1^+ + \delta$ and thus

$$p_1^+ \cdot D_1(p_1^+) > p_1^+ + \delta \tag{18}$$

Hence, profits are maximized globally at $p_1^+ = \text{argmax } D_1(p) \cdot p$, irrespective of whether other prices could induce an alternative stopping rule.

Next, consider seller $k = 2$. Since $p_1 = p_1^+ < p_k^e \;\forall\; k > 1$, any consumer sampling $k = 2$ satisfies $v_{i1} = 0$. Hence, the distribution of arriving consumers $f_2(x)$ is equal to (1) as derived under the consumers' stopping rule $\mathcal{R}^*$. Also, $v_{i1} = 0$ implies that consumers never return to the previous seller and thus always buy if both $v_{i2} > p_2$ and $\hat{X}_2 \subseteq \emptyset$. Since $p_2 \leq p_2^+ + \delta \Rightarrow \hat{X}_2 \subseteq \emptyset$, the previous argument for seller $k = 1$ now applies to $k = 2$. By induction, this implies $p_k = p_k^+ \;\forall k \leq N$. $\square$

Proof of Proposition 2

Proof. The FOC from the maximization problem (7) in the main text writes:

$$0 = \sum_{i=1}^{N} \phi_i \big(pD_i'(p) + D_i(p)\big) = p \sum_{i=1}^{N} \phi_i D_i'(p) + \sum_{i=1}^{N} \phi_i D_i(p)$$

Define $\tilde{q}(K^*, p) = \frac{\sum_{i=1}^{K^*} \phi_i D_i(p)}{-\sum_{i=1}^{K^*} \phi_i D_i'(p)}$. Then price p maximizing (7) must solve $\tilde{q}(K^*, p) = p$, where ϕ_k is defined by (6). Rewriting $\tilde{q}$ yields:

$$\tilde{q}(p) = \frac{\sum_{i=1}^{K^*} \phi_i \Theta_i(1) - \sum_{i=1}^{K^*} \phi_i \Theta_i(p)}{-\sum_{i=1}^{K^*} \phi_i \theta_i(p)}$$

As $\Theta_i(p)$ is log-concave $\forall\; i$, $\sum_{i=1}^{K^*} \phi_i \Theta_i(p)$ is log-concave. By the same argument made in proving proposition 1, this implies that $\bar{\Delta}(p) = \sum_{i=1}^{K^*} \phi_i \Theta_i(1) - \sum_{i=1}^{K^*} \phi_i \Theta_i(p)$ is log-concave. This permits to write $\tilde{q}(p) = \frac{\bar{\Delta}(p)}{-\bar{\Delta}'(p)}$. By Bagnoli and Bergstrom (2005), log-concavity of $\bar{\Delta}(p)$ is sufficient that $\tilde{q}(p)$ is decreasing in p. Further, it holds that $\lim_{p\to\underline{v}}\tilde{q}(p) > 0$ and $\lim_{p\to 1}\tilde{q}(p) = 0$. Hence, the price p solving the FOC $p = \tilde{q}(p)$ exists and is unique.[21]

□

Proof of Lemma 4

Proof. The proof is based on the following algebraic property:

$$\frac{x_1}{y_1} < \frac{x_2}{y_2} \Rightarrow \frac{x_1}{y_1} < \frac{\phi_1 x_1 + \phi_2 x_2}{\phi_1 y_1 + \phi_2 y_2} < \frac{x_2}{y_2} \quad \textit{for} \;\; x, y, \phi > 0 \tag{19}$$

Suppose $\phi_1, \phi_2 > 0$. The LHS follows from the following algebra:

$$\begin{aligned}\frac{x_1}{y_1} < \frac{x_2}{y_2} \;\; &\Leftrightarrow \phi_1 x_1 y_2 < \phi_1 x_2 y_1 \\ &\Leftrightarrow \phi_1 x_1 y_2 + \phi_2 x_1 y_1 < \phi_1 x_2 y_1 + \phi_2 x_1 y_1 \\ &\Leftrightarrow x_1(\phi_1 y_1 + \phi_2 y_2) < y_1(\phi_1 x_1 + \phi_2 x_2) \\ &\Leftrightarrow \frac{x_1}{y_1} < \frac{\phi_1 x_1 + \phi_2 x_2}{\phi_1 y_1 + \phi_2 y_2}.\end{aligned}$$

The proof of the RHS of inequality (19) is a tautology. Iterating over inequality (19) implies that $\frac{\sum_{i=1}^{K^*}\phi_i D_i(p)}{-\sum_{i=1}^{K^*}\phi_i D_i'(p)} < \frac{D_{K^*}(p)}{-D'_{K^*}(p)}$. Besides, the proof of lemma 2 tells us that $\frac{D_{K^*}(p)}{-D'_{K^*}(p)} < \frac{D_{K^*+1}(p)}{-D'_{K^*+1}(p)}$. Writing $x_1 = \sum_{i=1}^{K^*}\phi_i D_i(p)$, $y_1 = -\sum_{i=1}^{K^*}\phi_i D_i'(p)$, $x_2 = D_{K^*+1}(p)$. and $y_2 = -D'_{K^*+1}(p)$, it follows from inequality (19) that:

$$\frac{\sum_{i=1}^{K^*}\phi_i D_i(p)}{-\sum_{i=1}^{K^*}\phi_i D_i'(p)} < \frac{\sum_{i=1}^{K^*}\phi_i D_i(p) + \phi_{K^*+1}D_{K^*+1}(p)}{-\sum_{i=1}^{K^*}\phi_i D_i'(p) + \phi_{K^*+1}D'_{K^*+1}(p)}.$$

Hence, $\tilde{q}(p, K^*) < \tilde{q}(p, K^*+1)$ and more generally, $\tilde{q}(p, K^*) < \tilde{q}(p, K_2^*)$ if $K_2^* > K_1^*$. Since profit maximization implies $\tilde{q}(p, K_1^*, \hat{h}) = p(K_1^*)$ and $\tilde{q}(p, K_2^*, \hat{h}) = p(K_2^*)$, it follows that $p(K_2^*) > p(K_1^*)$. □

Proof of Proposition 3

The statement follows immediately from the fact that $p(K^*)$ is increasing in K^* (lemma 4) if K^* is weakly decreasing in s. while $p(K^*)$ is increasing in K^*. The argument below, including the subsequent lemmata 13, 14 and 15, shows that K^* is weakly decreasing in s.

[21] The existence and unique proofs are basically identical to those used in proving lemma 2.

For the simplest case, suppose that $V_k\big(K^*, p\big(K^*(s)\big), s\big) > 0 \; \forall \; k \leq N$.[22] Then, $K^* = N$ is the unique equilibrium. Otherwise however, whether an equilibrium with active search exists, depends crucially on the ordering of the elements of $\{\hat{s}_k(K^*)\}_{k=1,..,N}$ for every K^*, where

$$\hat{s}_k(K^*) \in \{s : V_k\big(K^*, p\big(K^*(s)\big), s\big) = 0\}$$

is the threshold level of search costs specifying that for a given price $p(K^*)$ and search persistence K^*, sampling seller k is worthwhile if and only if $s \leq \hat{s}_k$.

It is instructive to begin with the case where consumers become more pessimistic while searching, meaning that $V_k\big(K^*, p\big(K^*(s)\big), s\big) < V_{k+1}\big(K^*, p\big(K^*(s)\big), s\big) \; \forall \; k < K^*$ and for all K^*. This implies that $\hat{s}_k(K^*) > \hat{s}_{k+1}(K^*) \; \forall \; k < K^*$ such that search persistence decreases smoothly as there are no "jumps" in K greater than one.

Lemma 13. *There is a sequence of intervals $\{(\hat{s}_{K+1}(K), \hat{s}_K(K)]\}_{K=1,..,N}$ separated by closed neighborhoods such that consumers follow a stopping rule in pure strategies and sample up to*

$$K^* = K \text{ sellers for } s \in (\hat{s}_{K+1}(K), \hat{s}_K(K)]$$

Lemma 14. *There is a sequence of intervals $\{(\hat{s}_K(K), \hat{s}_K(K-1)\}_{K=1,..,N}$ disjoint from the intervals characterized in lemma 13 and separated by closed neighborhoods such that consumers follow a stopping rule in pure strategies up to seller $K^* = K-1$ and randomize over sampling and not sampling seller K with some probability $m(K) \in (0,1)$.*

Proof. Since $\hat{s}_k(K^*) > \hat{s}_{k+1}(K^*) \; \forall \; k < K^*$, $K^* < N$ only if$s > \hat{s}_N(N)$. Further, because continuation values are decreasing, $\hat{s}_{K^*}(K^*) > \hat{s}_{K^*+1}(K^*+1) \; \forall \; K^* < N$. Consequently, $K^* = N-1$ if $\hat{s}_{N-1} < s < \hat{s}_N - \epsilon, \; \epsilon > 0$.

At $s = \hat{s}_N(N)$, the the direct effect of a decrease in K^* from N to $N-1$ is zero, since the benefit of sampling seller N is exactly offset by s. However, the indirect effort through the price is strictly larger than zero. By proposition 4, it reduces the price and thus raises the continuation value such that $V_N(N-1, p(N-1), \hat{s}_N(N)) > 0$, implying that consumers would actually sample N sellers.

This inconsistency arises for all search costs in the range of $\hat{s}_N(N) \leq s < \hat{s}_N(N-1) < \hat{s}_{N-1}(N-1)$, where $\hat{s}_N(N-1)$ is determined by the general rule $\hat{s}_K(K-1) \in \{s : V_K(K, p(K-1), s) = 0\}$.

For $s \in \big[\hat{s}_N(N), \hat{s}_N(N-1)\big)$, consumers thus choose a mixed stopping rule, sampling seller N only with some probability $m(N) \in (0,1)$. From substituting $\phi'_N = \phi_N \cdot m(N)$ for ϕ_N into equation (7) and looking at the resulting FOC, it follows that the uniform optimal price $p(m(N))$ is decreasing in $m(N)$ and always satisfies $p(m) \in \big(p(N-1), p(N)\big) \; \forall \; m \in (0,1)$.

[22] I switched back to the notion of "seller k" here since I believe it makes the analysis more tractable. Notice that this is equivalent to referring to the continuation value of a consumer with $h = k-1$.

Since $V_N\big(N, p(N), s\big) < 0 < V_N\big(N, p(N-1), s\big)$ for $s \in \big[\hat{s}_N(N), \hat{s}_N(N-1)\big)$, there thus always exists an $m(N) \in (0,1)$ such that $V_N\big(N, p(m), s\big) = 0$ and $V_k\big(N, p(m), s\big) > 0 \ \forall\ k < N$.

Applying a randomized stopping rule prior to any seller $k < N$ cannot be an equilibrium strategy. Since $V_k\big(N, p(N), s\big) > 0 \ \forall\ k < N$, it follows that $V_k\big(K^*, p(m(k)), s\big) > 0$ for any $K^* < N$ and $m(k) < 1$ for some seller k by proposition 4. However, randomization at k requires $V_k = 0$, thus leading to a contradiction.

By the same argument, consumers use a only pure stopping strategies if $s \in (\hat{s}_N(N-1), \hat{s}_{N-1}(N-1)]$ and (conditional on $v_{ik} = 0 \ \forall\ k < N-1$) randomize over sampling an not sampling seller $N-1$ for $s \in (\hat{s}_{N-1}(N-1), \hat{s}_{N-1}(N-2)]$. If the threshold levels are ordered, this patterns repeats until $K^* = 1$. □

In general, the sequence of threshold levels $\{\hat{s}_k(K^*)\}_{k=1,..,N}$ need not be decreasing for every K^*. For every possible search persistence K', define $\underline{\hat{s}}(K') \in \min_k \big\{\hat{s}_k(K')\big\}_{k=1,..,K'}$. If $\hat{s}_{K'}(K') = \underline{\hat{s}}(K')$, the stopping rule is given by lemma 4 and 5 for all s such that $K^*(s) \geq K'$ with $K^*(s) \in \mathbb{K}$ as defined in (5). However, if there exists a $j < K'$ with $\hat{s}_j(K') = \underline{\hat{s}}(K^*)$, consumers do not begin to randomize over sampling seller K' if $s > \hat{s}_{K'}(K')$ as specified in lemma 5. Instead, already for $\hat{s}_{K'}(K')s > s_j(K')$, they randomize over sampling seller j and then continue sampling sellers up to K'. Formally, denote the upper bound on search persistence in the latter case by $\underline{K}^*(s) \in \max_K \big\{K : \hat{s}_K(K) > \underline{\hat{s}}(K)\big\}$ and denote by $j \in \{k : \hat{s}_k(\underline{K}^*) = \underline{\hat{s}}(\underline{K}^*\}$ the seller with the lowest threshold given $\underline{K}^*$. Lemma 15 summarizes the general randomized stopping rule:

Lemma 15. *Consumers follow a stopping rule in pure strategies up to seller $j-1$ and randomize over continuing to sample seller j with some probability $m(j) \in (0,1)$ for $s \in (\hat{s}_j(\underline{K}^*), \hat{s}_j(j-1)]$. Consumers who sample j continue sampling sellers up to $k = \underline{K}^*$ if $v_{ik} = 0 \ \forall\ k < \underline{K}^*$.*

Proof. Consumers with history $h = \underline{K}^*$ would find sampling seller $\underline{K}^*$ worthwhile if $s \leq \hat{s}_{\underline{K}^*}(\underline{K}^*)$. However, consumers do not "reach" seller $\underline{K}^*$ when the price is $p(\underline{K}^*)$ since $\hat{s}_j(\underline{K}^*) < \hat{s}_{\underline{K}^*}(\underline{K}^*)$ by construction. That is, they would stop sampling at seller $j < \underline{K}^*$.

For $j > 1$, the issue is resolved with a unique mixed strategy. Randomizing over the decision to sample seller j also affects demand at all seller $k' > j$. While by construction $V_k\big(\underline{K}^*, p\big(m(j), \underline{K}^*\big), s\big) > 0 \ \forall\ j < k \leq \underline{K}^*$ is unfeasible since consumers are not indifferent, the mass of consumers is reduced by the same fraction $1 - m(j)$ for all sellers $k > j$.

As in lemma 14, there exists a $m(j) \in (0,1)$ such that $V_j\big(\underline{K}^*, p(m(j)), s\big) = 0$ for $s > \hat{s}_j(\underline{K}^*$ and where $p(m)$ maximizes sellers profits. Notice that for $m(j) \to 0$, $K^* = j-1$ effectively. Hence, the threshold search level for which no mixed strategy $m(j) \in (0,1)$ can yield $V_j\big(\underline{K}^*, p(m(j)), s\big) = 0$ is given by $\hat{s}_j(j-1)$ since$V_j\big(j-1, p(j-1), \hat{s}_j(j-1)\big) = 0$

by construction. For $s < \hat{s}_j(j-1)$, consumers follow again a pure stopping strategy with $K^* = j-1$ by lemma 13.

The mixed strategy equilibrium is unique. To see why, note that generally, $V_{j'} \neq V_j$ for $j \neq j'$ and that randomizing over the decision to sample any seller j' requires that $V_{j'} = 0$ in equilibrium. (Note that I drop some arguments of V here).

Suppose that randomizing over sampling $j' > j$ was an equilibrium and that $m(j') < 1$. Then, a price ensuring $V_{j'}\big(p(m(j'))\big) = 0$ implies $V_j(p(m(j')) < 0$ since by construction of j, $V_j(p) < V_k(p)\ \forall\ k \leq \underline{K}^*$ for any price p.[23] Hence, consumers would neither sample seller j nor any $j' > j$, which is a contradiction to $m(k) \in (0,1)$. Next, suppose that consumers randomize at some $k' < j$. Then, $V_{k'}\big(p(m(k'))\big) = 0$ implies $V_j\big(p(m(k'))\big) < 0$. However, then consumers sample at most $k' < j$ sellers, and it follows from proposition 4 that $p(m(k')) < p(m(j))\ \forall\ m(j) \in [0,1]$. But $V_j\big(p(m(k'))\big) > 0$, which is a contradiction. □

To conclude, consumers' search persistence is always weakly decreasing in the level of search costs s, though it may involve mixing over sampling additional sellers or discontinuous jumps if search costs are above a certain threshold. Hence, proposition 3 obtains.

Proof of Proposition 4

Proof. Note that $p(m(j))$ is increasing in $m(j)$ only if $j > 1$. Too see why, consider the FOC for a randomized stopping rule $m(j)$:

$$p = \frac{\sum_{i=1}^{j-1} \phi_i D_i(p) + m(j) \sum_{i=j}^{\underline{K}^*} \phi_i D_i(p)}{-\Big(\sum_{i=1}^{j-1} \phi_i D_i'(p) + m(j) \sum_{i=j}^{\underline{K}^*} \phi_i D_i'(p) \Big)} \tag{20}$$

If $j = 1$, $m(j)$ cancels out from the RHS of (20). Thus, randomizing with $m(j) < 1$ has no effect on on the equilibrium price. Consequently, no mixed strategy $m(j)$ for $j = 1$ exists that renders $V_1 = 0$. Moreover, as shown in the proof of lemma 13 and 14, there also exists no mixed strategy for $m(k) \in (0,1)$ for $k > j = 1$. □

[23]Seller j is defined in the main text.

9.2 Proofs of Section 2.4

Proof of Proposition 5

Proof. Using inequality (19), it follows immediately that:

$$\frac{D_{1^*}(p)}{-D'_{1^*}(p)} < \frac{\sum_{i=1}^{K^*} \phi_i D_i(p)}{-\sum_{i=1}^{K^*} \phi_i D'_i(p)} < \frac{D_{K^*}(p)}{-D'_{K^*}(p)}$$

Since by lemma 2, prices with and without tracking are unique, the stated result obtains. □

Proof of Lemma 5

Proof. The proof is provided in the main text. Suppose that consumers followed the stopping rule $\mathcal{R}^*$, irrespective of the sellers price, thereby eliminating price competition entirely as in the unique equilibrium with tracking. In principle, sellers can choose any price, including the price they would choose if tracking were not available. However, even when assuming that consumers follow $\mathcal{R}^*$, sellers choose different prices to maximize profits by proposition 5. Hence, they cannot obtain less profits if K^* under search with tracking is at least as high as K^* under search without tracking.

□

Proof of Proposition 6

Proof. The result is an immediate consequence of lemma 4. If given the same level of search costs, the market is active under tracking while it is inactive under no tracking, strictly more matches are realized under tracking, thus leading to both higher consumer surplus and higher profits. □

Formulas used in Surplus Computations

Consumer surplus is captured by the continuation value prior to sampling the first sellers: Substituting $k = 1$ in equation (4) und using the specification for $f()$ and $g()$ as provided in the main text, I obtain

$$V_1 = \sum_{k=0}^{K^*} \left(\int_{p_k}^{1} (1-x)(x-p_k)\big(x\big)^k \mathrm{d}x - s \int_0^1 \big(x\big)^k \mathrm{d}x \right).$$

Under search with tracking, a seller in position k (observing $h = k-1$) maximizes the following profit function:

$$\pi_k = p_k \cdot \left(\int_{p_k}^{1} (1-x)(x)^{k-1} \frac{1}{\int_0^1 (x)^{k-1} \mathrm{d}t} \mathrm{d}x \right). \tag{21}$$

To derive the uniform price without tracking conditional on search persistence K^*, the above profit functions from equation must be weighted by ϕ_k. I do not normalize conditional on K^* because the sellers FOC would remain unchanged and because I look for industry profit, not for seller profit per consumer.

$$\Pi(K^*, p) = p \cdot \left(\sum_{k=1}^{K^*} \phi_k \cdot \int_{p}^{1} (1-x)(x)^{k-1} \frac{1}{\int_0^1 (x)^{k-1} \mathrm{d}t} \mathrm{d}x \right) \tag{22}$$

where the probability ϕ_k is given in its general form in (6) and now writes $\phi_k = \int_0^1 (x)^{k-1} \mathrm{d}x$. Hence, overall profits without tracking take a very simple form:

$$\Pi(K^*, p) = p \cdot \left(\sum_{k=1}^{K^*} \cdot \int_{p}^{1} (1-x)(x)^{k-1} \mathrm{d}x \right) \tag{23}$$

The formulas above are implemented in a *Mathematica* code to derive optimal prices for every possible K^*. Using expressions for V_k, the equilibrium K^* is computed for every level of search costs. The code can be obtained from the corresponding author upon request.

9.3 Proofs of Section 2.5

Proof of Lemma 6

Proof. **(a) Existence and uniqueness** If consumers follow $\mathcal{R}^*$, seller k's demand writes:

$$D_k(p) = \int_{p}^{\bar{v}} g(x) \frac{\big(1-g(x)\big)^{k-1} f(x)}{\int_{\underline{v}}^{1} \big(1-g(t)\big)^{k-1} f(t) \mathrm{d}t} \mathrm{d}x$$

Profit maximization with respect to the above demand function yields first-order conditions which are equivalent to those shown in the proof of the existence of an increasing price sequence in proposition 1. Hence, a solution to

$$\max_p D_k(p)p$$

exists. Further, since the FOC $\frac{D_k(p)}{-D_k'(p)}$ is decreasing, it is also unique. Moreover, the same reasoning as in Part II regarding the possibility of setting to a price that changes the stopping rule applies. Consequently, uniqueness is preserved when accounting for deviations from $\mathcal{R}^*$.

(b) Decreasing price sequence

If $g'(x) > 0$, the inequality in (13) is reversed. Hence,

$$q_j(p) > q_{j+1}(p) \; \forall \; p \in [\underline{v}, \bar{v}].$$

Since the solution to $q_k(p) = p$ is unique $\forall \; j$, it follows that $p_k < p_{k-1} \; \forall \; k \leq N$ in equilibrium. $\square$

Proof of Lemma 7

Proof. By lemma 6, $\{p_k^*\}_{k=1,\ldots,N}$ is optimal if consumers follow $\mathcal{R}^*$ always. As in any PBE, expectations are correct, it suffices to show that given sellers' optimal prices if $\hat{\Delta} > g'(x) \; \forall \; x \in X$, applying $\mathcal{R}^*$ is optimal for consumers.

Using previous notation $\hat{X}_k = \{x \in X : g(x)\big(p_k' - p_{k+1}\big) > s\}$, recall that

$$\hat{X} \subseteq \emptyset \;\Leftrightarrow\; \mathcal{R}^* \textit{ is optimal.}$$

Hence, all types apply $\mathcal{R}^*$ if

$$g(x)\big(p_k^e - p_{k+1}^e\big) \leq s \; \forall \; k < N, \; x \in \; X$$

As $g'(x) > 0$, the type $x_i = \bar{v}$ has the highest probability of encountering a match. Hence, $\hat{X} \subseteq \emptyset$ if $g(\bar{v})\big(p_k^e - p_{k+1}^e\big) \leq s \; \forall \; k$. Equation (13) shows that the difference in prices is a function of the slope of $g(x)$. In particular, $|p_k^e - p_{k+1}^e\,| \to 0 \; \textit{if} \; g'(x) \to 0 \; \forall \; x$. Hence, it is possible to find a $g(x)$ sufficiently flat such that $\hat{X} \subseteq \emptyset$ for every $s > 0$. $\square$

Proof of Lemma 8

Proof. Replace p_k^+ by p_k^* in the proof of lemma 3. The result follows immediately. The reason why the same argument as in the proof of lemma 3 applies is given in the main text. $\square$

Proof of Proposition 6

Proof. Replace p_k^+ by p_k^* in the proof of Proposition 2. The result follows immediately. The reason why the same argument as in the proof of Proposition 2 applies is given in the main text. $\square$

9.4 Proofs of Section 2.6

Proof of Lemma 11

Proof. **(a) Existence and Uniqueness:** Random consumer search from the perspective of sellers for all histories $h \geq \hat{h}$ implies that a seller at position $k > \hat{h}$ computes the probability that his position is k conditional on $K^* \geq k > \hat{h}$. This probability is denoted by $\phi_k(\hat{h})$ with:

$$\phi_k(\hat{h}) = \phi_k / \sum_{i=\hat{h}+1}^{K^*} \phi_i \;\forall\; k > \hat{h} \tag{24}$$

where ϕ_i is defined as in (6). Notice that as $\phi_k < \phi_{k+1}$ by construction, it also holds that $\phi_k(\hat{h}) < \phi_{k+1}(\hat{h})$. Conditional on not observing the search history h, the optimal price solves the following FOC:

$$\frac{\sum_{i=\hat{h}+1}^{K^*} \phi_i(\hat{h})\Theta_i(1) - \sum_{i=\hat{h}+1}^{K^*} \phi_i(\hat{h})\Theta_i(p)}{-\sum_{i=\hat{h}+1}^{K^*} \phi_i(\hat{h})\theta_i(p)} = p.$$

Analogously to the previous arguments, it follows that the optimal price exists and is unique for every $\hat{h}$.

(b) Comparison between $p(\hat{h})$ and $p_{\hat{h}+1}$: $p(\hat{h})$ is uniquely defined by:

$$p(\hat{h}) = \frac{\sum_{i=\hat{h}+1}^{K^*} \phi_i D_i\big(p(\hat{h})\big)}{-\sum_{i=\hat{h}+1}^{K^*} \phi_i D_i'\big(p(\hat{h})\big)}.$$

Using $D_{\hat{h}+1}(p)/-D'_{\hat{h}+1}(p) < ... < D_{K^*}(p)/-D'_{K^*}(p)$ and applying inequality (19) again implies:

$$\frac{\sum_{i=\hat{h}+1}^{K^*} \phi_i D_i(p)}{-\sum_{i=\hat{h}+1}^{K^*} \phi_i D_i'(p)} > \frac{D_{\hat{h}+1}(p)}{-D'_{\hat{h}+1}(p)} \;\forall\; p \in X$$

As $p_{\hat{h}+1} = \frac{D_{\hat{h}+1}(p_{\hat{h}+1})}{-D'_{\hat{h}+1}(p_{\hat{h}+1})}$, it follows by the same argument as in proposition 1 that $p(\hat{h}) > p_{\hat{h}+1} \;\forall\; 1 \leq \hat{h} < K^* - 1$. If $\hat{h} = K^* - 1$, the FOC determining $p(\hat{h})$ reduces to $p = \frac{D_{K^*}(p)}{-D'_{K^*}(p)}$ and is thus identical to the FOC for seller K^* if the history $\hat{h}$ is disclosed. Hence, $p(\hat{h}) = p_{K^*}$ for $\hat{h} = K^* - 1$. □

Proof of Proposition 10

Proof. Suppose that $\hat{h} < K^*$. In equilibrium, the strategy $\hat{h}$ must be optimal, i.e. $p^e(NT) < p^e_{\hat{h}+1}$. In any symmetric equilibrium, sellers' beliefs must satisfy $\mu(h) =$

$0\ \forall\ h < \hat{h}$ and $\mu(h) = \phi_{h+1}\ \forall h \geq \hat{h}$, where ϕ_{h+1} equals a seller's probability of being in position $h+1$ in a consumer's search process as defined in (6). Then by Lemma 7, the seller's optimal price conditional on $\hat{h}$ satisfies $p(NT) = p(\hat{h}) \geq p_{\hat{h}+1}\ \forall\ \hat{h}$. If $\hat{h} = K^*$, $p(\hat{h}) \geq p_{\hat{h}+1}$ does not affect a consumer's choice because she does not sample seller $\hat{h}+1$. However if $\hat{h} < K^*$, the optimal no tracking price $p(NT)$ contradicts the expectation of $p^e(NT) < p^e_{\hat{h}+1}$, which is necessary to sustain the equilibrium strategy $\hat{h} < K^*$. Besides, a price $p(NT) > p_{K^*}$, despite being an action that is never chosen in equilibrium, cannot be part of the seller's equilibrium strategy. If a seller observes the off-equilibrium choice NT, the maximum possible history can be $h = K^* - 1$ as otherwise the consumer would have ended search. Consequently, setting $p'(NT) = p_{K^*} < p(NT)$ constitutes a profitable deviation. Notice also that unique full disclosure outcome is robust to assuming the opposite tie-breaking rule in favor of no disclosure if $p^e(NT) = p_{\hat{h}+1}$. Changing the tie-breaking rule in that way allows for an equilibrium with $\hat{h} = K^* - 1$. However, consumers with a history $h = \hat{h}$ are the only consumers choosing no disclosure. Hence, the choice of $d = NT$ perfectly reveals the type, sellers set $p(NT) = p_{\hat{h}+1} = p_{K^*}$ and the outcome is equivalent to the unique equilibrium under the alternative tie-breaking rule.

$\square$

Proof of Lemma 12 and Proposition 11

Reversing the inequality $h \geq \hat{h}$ in the proofs of Lemma 11 and Proposition 10 (to $h \leq \check{h}$) yields the results.

References

Acquisti, A. and Varian, H. R. (2005). Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381.

Aguiar, M. and Hurst, E. (2007). Life-cycle prices and production. *American Economic Review*, 97(5):1533–1559.

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

Anderson, S. P. and Renault, R. (2015). Search direction. Technical report, mimeo.

Arbatskaya, M. (2007). Ordered search. *The RAND Journal of Economics*, 38(1):119–126.

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

Armstrong, M., Vickers, J., and Zhou, J. (2009). Prominence and consumer search. *The RAND Journal of Economics*, 40(2):209–233.

Armstrong, M. and Zhou, J. (2010). Conditioning prices on search behaviour. Technical report, University Library of Munich, Germany.

Armstrong, M. and Zhou, J. (2011). Paying for prominence. *The Economic Journal*, 121(556).

Armstrong, M. and Zhou, J. (2015). Search deterrence. *The Review of Economic Studies*, 83(1):26–57.

Bagnoli, M. and Bergstrom, T. (2005). Log-concave probability and its applications. *Economic theory*, 26(2):445–469.

Bar-Isaac, H., Caruana, G., and Cuñat, V. (2012). Search, design, and market structure. *The American Economic Review*, 102(2):1140–1160.

Belleflamme, P. and Peitz, M. (2015). *Industrial organization: markets and strategies.* Cambridge University Press.

Belleflamme, P. and Vergote, W. (2016). Monopoly price discrimination and privacy: The hidden cost of hiding. *Economics Letters*, 149:141–144.

Borgesius, F. Z. and Poort, J. (2017). Online price discrimination and eu data privacy law. *Journal of Consumer Policy*, 40(3):347–366.

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

Chen, Y. and Zhang, Z. J. (2009). Dynamic targeted pricing with strategic consumers. *International Journal of Industrial Organization*, 27(1):43–50.

Conitzer, V., Taylor, C. R., and Wagman, L. (2012). Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science*, 31(2):277–292.

De Cornière, A. (2016). Search advertising. *American Economic Journal: Microeconomics*, 8(3):156–88.

Diamond, P. A. (1971). A model of price adjustment. *Journal of economic theory*, 3(2):156–168.

Ellison, G. and Ellison, S. F. (2014). Match quality, search, and the internet market for used books1.

Fudenberg, D. and Tirole, J. (2000). Customer poaching and brand switching. *RAND Journal of Economics*, pages 634–657.

Haan, M. A. and Moraga-González, J. L. (2011). Advertising for attention in a consumer search model. *The Economic Journal*, 121(552):552–579.

Hart, O. D. and Tirole, J. (1988). Contract renegotiation and coasian dynamics. *The Review of Economic Studies*, 55(4):509–540.

Johnson, J. P. and Myatt, D. P. (2006). On the simple economics of advertising, marketing, and product design. *The American Economic Review*, 96(3):756–784.

Mikians, J., Gyarmati, L., Erramilli, V., and Laoutaris, N. (2012). Detecting price and search discrimination on the internet. In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*, pages 79–84. ACM.

Milgrom, P. and Roberts, J. (1986). Relying on the information of interested parties. *The RAND Journal of Economics*, pages 18–32.

Montes, R., Wilfried, S.-Z., and Valletti, T. (2017). The value of personal information in markets with endogenous privacy.

Moraga-González, J. L. and Petrikaitė, V. (2013). Search costs, demand-side economies, and the incentives to merge under bertrand competition. *The RAND Journal of Economics*, 44(3):391–424.

Moraga-González, J. L., Sándor, Z., and Wildenbeest, M. R. (2017). Prices and heterogeneous search costs. *The RAND Journal of Economics*, 48(1):125–146.

Taylor, C. and Wagman, L. (2014). Consumer privacy in oligopolistic markets: Winners, losers, and welfare. *International Journal of Industrial Organization*, 34:80–84.

Taylor, C. R. (2004). Consumer privacy and the market for customer information. *RAND Journal of Economics*, pages 631–650.

Turow, J., King, J., Hoofnagle, C. J., Bleakley, A., and Hennessy, M. (2009). Americans reject tailored advertising and three activities that enable it.

Villas-Boas, J. M. (1999). Dynamic competition with customer recognition. *The Rand Journal of Economics*, pages 604–631.

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

Yang, H. (2013). Targeted search and the long tail effect. *The RAND Journal of Economics*, 44(4):733–756.

Zhou, J. (2011). Ordered search in differentiated markets. *International Journal of Industrial Organization*, 29(2):253–262.