

DISCUSSION PAPER SERIES

DP11566

ORDERED CONSUMER SEARCH

Mark Armstrong

INDUSTRIAL ORGANIZATION



ORDERED CONSUMER SEARCH

Mark Armstrong

Discussion Paper DP11566

Published 09 October 2016

Submitted 09 October 2016

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Mark Armstrong

ORDERED CONSUMER SEARCH

Abstract

The paper discusses situations in which consumers search through their options in a deliberate order, in contrast to more familiar models with random search. Topics include: the existence of ordered search equilibria with symmetric sellers (all consumers first inspect the seller they anticipate sets the lowest price, and a seller which is inspected first by consumers will set the lowest price); the use of price and non-price advertising to direct search; the impact of consumers starting a new search at their previous supplier; and the incentive a seller can have to raise its own search cost. I also show how ordered search can be reformulated as a simpler discrete choice problem without search frictions or dynamic decision making.

JEL Classification: D21, D43, D83, L11, L15, M37

Keywords: consumer search, sequential search, ordered search, directed search, discrete choice, oligopoly, advertising, obfuscation

Mark Armstrong - mark.armstrong@economics.ox.ac.uk
Department of Economics, University of Oxford and CEPR

Acknowledgements

Parts of the paper reflect joint work with John Vickers and Jidong Zhou, and I am grateful to them for further discussion. I am also grateful to Laura Doval, Renaud Foucart, Tobias Gamp, José Luis Moraga, Alexsei Parakhonyak, Vaiva Petrikaite, Régis Renault, Andrew Rhodes, David Sappington, Sandro Shelegia and Chris Wilson for many helpful comments.

Ordered Consumer Search

Mark Armstrong*

September 2016

Abstract

The paper discusses situations in which consumers search through their options in a deliberate order, in contrast to more familiar models with random search. Topics include: the existence of ordered search equilibria with symmetric sellers (all consumers first inspect the seller they anticipate sets the lowest price, and a seller which is inspected first by consumers will set the lowest price); the use of price and non-price advertising to direct search; the impact of consumers starting a new search at their previous supplier; and the incentive a seller can have to raise its own search cost. I also show how ordered search can be reformulated as a simpler discrete choice problem without search frictions or dynamic decision making.

Keywords: Consumer search, sequential search, ordered search, directed search, discrete choice, oligopoly, advertising, obfuscation.

1 Introduction

Consider a consumer who wishes to purchase one product from the several variants which are available. In some cases she might know exactly what she wants in advance, and no prior market investigation is needed. In other situations, she might be so ill-informed that she stumbles randomly from one option to another, incurring a search cost each time, until she discovers something “good enough”. Between these extremes, though, a consumer prefers a particular order of search (which may differ across consumers). For instance, in an online book market De los Santos, Hortacsu and Wildenbeest (2012, Table 6) report: of those consumers who searched for a book only once, 69% inspected *Amazon*; of those

*Department of Economics and All Souls College, University of Oxford. This paper was presented as the Marshall Lecture at the 2016 Annual Conference of the European Economic Association. Parts of the paper reflect joint work with John Vickers and Jidong Zhou, and I am grateful to them for further discussion. I am also grateful to Laura Doval, Renaud Foucart, Tobias Gamp, José Luis Moraga, Alexsei Parakhonyak, Vaiva Petrikaite, Régis Renault, Andrew Rhodes, David Sappington, Sandro Shelegia and Chris Wilson for many helpful comments.

who searched two sellers, 57% inspected *Amazon* first, and *Amazon* had a very large share of the second inspections of those consumers who first inspected other sellers.

In this paper I say that a seller is “prominent” when it is inspected early in a consumer’s chosen order of search. The classical theory of optimal sequential search outlined in the next section shows that a product is more likely to be prominent for a consumer if that product: (i) has a distribution for match utility which is better (in the sense of first-order stochastic dominance) or riskier; (ii) is expected (or observed) to have a lower price, or (iii) is less costly to inspect. These three factors are discussed in more detail in the following.

Nelson (1970, page 312) writes that consumer search need not “be conducted at random. Prior to sampling, a consumer can obtain information from relatives and friends, consumer magazines, or even from advertising”. Likewise, a consumer might know in advance some component of her tastes—someone looking for a house might care about its location, say—but needs a closer inspection to discover other product attributes. Such a consumer is likely to first inspect those products which have the most favourable “known” attributes. Alternatively, the consumer might consult an expert intermediary to recommend a search order. For instance, enquiring about hotels in a city on specified dates from an online travel agent may generate a list of available options ordered according to the travel agent’s ranking algorithm, and if the algorithm is any good the consumer will do well to search in the suggested order.¹ A seller’s willingness-to-pay to be recommended to consumers can also be informative about the suitability of its product. For instance, when consumers use a search engine it might be that the most relevant seller for them is the seller that bids the most to be displayed first in the sponsored search results, in which case consumers should inspect sellers in the order they appear in the sponsored results.

Turning to factor (ii), if consumers view sellers as otherwise symmetric but expect that one seller offers a lower price, all consumers will optimally inspect that seller first. Conversely, when a seller is prominent for most consumers, that seller will often have an incentive to set a lower price than its rivals. (A seller placed further back knows that the few consumers it encounters have been disappointed in the offers received so far, and so

¹Ursu (2015) studies such a travel agent empirically. An important feature of her dataset is that the travel agent randomized its recommendations to some consumers, which allows her to study the effect of rank on click rates and purchase decisions in a clean way. In the random order treatment, consumers clicked on links with decreasing frequency further down the page (since presumably they believe the ranking has some content) but their purchase probability contingent on clicking does not depend on rank. However, when the travel agent’s true ranking was displayed, the purchase probability did depend strongly on rank, suggesting that the ranking algorithm was indeed useful to consumers as a guide to search.

it can afford to set a high price.) Because of this, if other consumers use a rule of thumb for choosing which seller to inspect first—for example, they first inspect the seller which runs the most visible advertising campaign—then an individual consumer should do the same. This kind of self-fulfilling prophecy means that ordered search, which involves price dispersion and a skewed pattern of sales and profits across sellers, can be an equilibrium even in symmetric environments.

In other situations sellers advertise their prices to consumers in advance, for instance on the internet, and consumers do not need to search to discover prices. In many such situations, though, consumers still need to investigate other product attributes. When products are otherwise symmetric consumers will choose to inspect products in order of increasing price, and more generally advertised prices can be used to attract searchers (which is not possible when prices are discovered after search, when a seller's price can be used only to retain consumers who choose to inspect that seller). In sum, when prices are hidden a prominent seller will often set a low price, while when prices are advertised, the seller with a low price often becomes prominent.

Some products have lower search costs than others (factor (iii) above). For instance, geography or shop layout influences a consumer's search order, and she might choose to inspect the nearest option first (which might be different for different consumers). In a physical store it is easier to inspect products displayed at eye level or on the ground floor, regulation might require certain products to be on the "top shelf", while unhealthy snacks aimed at children might profitably be placed at a lower height. Alternatively, some products might have their characteristics described in a transparent way, while others may be harder to evaluate without the consumer spending time on inspection. Judicious design of store layout and product display might mean that a multiproduct seller can *force* the consumer to consider its products in a particular order. Consumers might find it less costly to inspect a new product from a supplier from which they have purchased before than from a new supplier, perhaps because they have contact details readily available.

Another way to reduce search costs is to first visit a location (either physical or online) known to have a concentration of varieties of the relevant product. A cluster with several independent suppliers not only allows a better chance of finding a good product, but competition between suppliers within the cluster might lead to low prices. Similarly, a consumer might choose to first visit a "big box" store which stocks more varieties of the

product in question. However, because pricing is coordinated within the store, it is likely that the store will set a higher price than its smaller rivals, and consumers choosing their search order have to trade off the one-stop shopping benefits of greater variety with the higher price they will have to pay there.

The rest of this paper is organized as follows. Section 2 describes the principles governing optimal sequential search in fairly general terms, and shows how this search problem can be reformulated as a simpler discrete choice problem without search frictions and dynamic decision making. Section 3 uses this theory to describe outcomes in a model where sellers choose prices for their products. In simple settings where each consumer views the sellers as symmetric *ex ante*, this model often exhibits multiple equilibria: the consumer search order depends on which sellers choose lower prices, and the prices that sellers choose depend on where they are in the search order. Random consumer search is one equilibrium (and is often the focus of existing oligopoly models), although it is often unstable. However, if the demand system is “smoothed”, by giving individual consumers heterogeneous preferences over which sellers to inspect first, the market might have a single equilibrium. Extensions to this basic model are presented in section 4, which illustrate several of the reasons for ordered search described above. These include discussions of how sellers choose prices when they anticipate that consumers start a new search process with their previous supplier, how it might be profitable for a seller to deliberately increase its inspection costs, and the use of both price and non-price advertising as a guide for search. Section 5 suggests some promising options for further research. The relevant literature is discussed as I present these various aspects of ordered search, while more technical arguments are left to the appendix.

2 Opening the box

Consider a consumer who wishes to select one product from several variants which are available. One way to model this decision problem is to suppose that the consumer knows in advance her idiosyncratic match utility for each product i , say v_i , and knows in advance each product’s price, p_i , and chooses the option with the highest net surplus $v_i - p_i$ provided this is positive. This is the “discrete choice” problem. This paper studies another scenario, where before purchase the consumer needs to incur a cost s_i to discover product i ’s characteristics, v_i and p_i . (In section 4.4 I also study a scenario between these two

extremes, where consumers know each product’s price in advance but need to discover the associated match utility.) The kinds of products where consumers have idiosyncratic tastes, and which they usually wish to inspect in some way before buying (even if they know the price in advance), include cameras, cars, clothing, furniture, hotels, novels, perfume, pets and phones.²

Before studying in the next section how equilibrium prices are determined, we first describe the risk-neutral consumer’s optimal search strategy for an exogenous set of options. Weitzman (1979) provides the key to understanding optimal sequential search through a finite number of mutually exclusive options (“boxes”) with uncertain payoffs. The consumer’s payoff from option i is a random variable v_i , where her payoffs are independently distributed across options with CDF $F_i(v_i)$ for option i . Discovering the realization of v_i inside box i involves the non-refundable inspection cost s_i , and the consumer cannot select a box without first opening it. There is free recall, so that the consumer can costlessly return to claim the payoff from a box opened earlier. The consumer can select the outside option, which has deterministic payoff zero, at any point, in which case we can suppose that the random variable v_i is non-negative. The consumer wishes to consume at most one of the options, and aims to maximize the expected value of the selected option net of total search costs. To do this she decides both the order in which to inspect options and the rule for when to terminate search (in which case she consumes the best option opened so far). Weitzman refers to this as “Pandora’s problem”, and so the consumer is female in this paper.³

For now, suppose that $\mu_i \equiv \mathbb{E}_i v_i > s_i$ for each i , for otherwise it is never optimal to open box i and this option can be eliminated from her choice problem. (Here, \mathbb{E}_i denotes taking expectations with respect to the distribution for v_i .) Define the “reservation price” of box i to be the unique price r_i which satisfies

$$\mathbb{E}_i \max\{v_i - r_i, 0\} = s_i . \tag{1}$$

(Since $\mu_i > s_i$, this reservation price is positive.) Here, r_i is the highest price such that the consumer is willing to incur the sunk cost s_i for the right to purchase the product at

²As discussed by De los Santos, Hortacsu, and Wildenbeest (2012), even an apparently homogeneous product such as a specific book might have different match utilities from different online retailers (reflecting whether they currently stock the book, for example), which require inspection to discover.

³This is not particularly apt terminology, as in the myth there was just one box. Weitzman’s framework is more general than that presented here, and allows for time discounting as well as inspection costs.

that price once she has discovered her match utility. Another interpretation of r_i is that it is the threshold for her current payoff which determines when it is worth opening this box: the expected benefit of inspecting box i given that the consumer already has secured a potential payoff $x \geq 0$ to which she can freely return is

$$\mathbb{E}_i \max\{v_i, x\} - s_i - x = \mathbb{E}_i \max\{v_i - x, 0\} - s_i , \quad (2)$$

which is positive if and only if $x < r_i$.

Weitzman shows that an optimal search strategy in this context—“Pandora’s rule”—is as follows.⁴

Selection rule: If another box is to be opened, it should be the unopened box with the highest reservation price;

Stopping rule: Terminate search whenever the maximum payoff discovered so far exceeds the reservation price of all unopened boxes (which is zero if no box remains unopened), and consume the option with this maximum payoff.

For instance, suppose there are three boxes with respective reservation prices r_i and realized payoffs v_i given by

$$(r_1, v_1) = (5, 2) ; (r_2, v_2) = (10, 4) ; (r_3, v_3) = (3, 7) . \quad (3)$$

Then the consumer (who of course does not know the realized payoff v_i until she inspects that box) should first inspect box 2 as that has the highest reservation price, should go on to inspect box 1 (since that box has reservation price above her current payoff $v_2 = 4$), then come back to consume the payoff in box 2 without inspecting box 3 (since v_2 is above both v_1 and r_3). Unless all boxes are identical *ex ante*, it can be optimal to return to consume an earlier-opened box before all boxes are opened.⁵

The reservation price in (1) depends only on the properties of that option, i.e., s_i and F_i . The reservation price for a box is not the same as that box’s stand-alone surplus,

⁴The optimality of this rule depends on a number of factors. When the consumer cannot freely return to an earlier opened box, it may not be optimal to inspect boxes in order of their reservation values. See Salop (1973) for an investigation of optimal order of search when there is no recall (which is a common assumption in the job search literature). Olszewski and Weber (2015) discuss how the rule needs to be modified when the agent gains utility from all opened boxes, while Doval (2014) considers the situation where the agent can consume the contents of a box without inspecting it first (and without incurring the search cost).

⁵De los Santos, Hortacsu and Wildenbeest (2012, Table 3) show that among the multi-seller searchers around 40% went back to buy from an earlier option, usually without exhausting all options first.

$\mu_i - s_i$. If the consumer could only choose one box to open, she would choose the box with the highest value of $\mu_i - s_i$, which need not be the box with the highest r_i .⁶ As is intuitive, the reservation price in (1) is decreasing in s_i and increasing in $[1 - F_i(\cdot)]$. In addition, since it depends on the right-tail of the distribution, all else equal it increases with the “riskiness” of the option.

Pandora’s rule can conveniently be re-expressed as a static discrete choice problem without search frictions or dynamic choice.⁷ Specifically, Pandora’s rule is equivalent to the choice rule whereby the consumer selects the box with the highest index

$$w_i \equiv \min\{r_i, v_i\}, \quad (4)$$

where v_i is the realized payoff inside box i . (This is the case in (3) above, when box 2 was ultimately selected.) To see this I show that box j is not chosen under Pandora’s rule when $w_j < w_i$. If $r_j < r_i$, then box i will be inspected before j , and j could then be chosen only if it is inspected, which requires $v_i < r_j$, and then only if $v_i < v_j$, which taken together contradict the assumption $w_j < w_i$. If instead $r_j > r_i$, then box j is inspected before i . But then $w_j < w_i$ implies $v_j < \min\{r_i, v_i\}$, which implies that box j is not consumed before first inspecting i , which then reveals a superior payoff. In either case, the inequality $w_j < w_i$ implies that box j is not chosen. Since some box is eventually chosen under Pandora’s rule, we deduce it is the box with the highest w_i . In sum, while the box-specific index r_i determines which box the consumer opens first, the box-specific index w_i determines which box is ultimately selected.⁸ As is implicitly shown in Theorem 1 in Kleinberg, Waggoner, and Weyl (2016), which itself builds on Weber (1992), this reformulation of Pandora’s rule allows for an elegant proof of Weitzman’s result which is presented in Appendix A below.

When a population of consumers choose their options it will often be the case that consumers differ in their reservation prices r_i for box i , and hence differ in their optimal order of search.⁹ Consumers might differ in their cost of inspecting a given box (e.g., due to their different geographic locations) or in their prior distribution for a box’s match

⁶Chade and Smith (2006) analyze which boxes to inspect in non-sequential search.

⁷This discussion develops the analysis in Armstrong and Vickers (2015, pages 303-4), where we showed how a search problem with free recall of earlier options can be recast as a discrete choice problem. (This reformulation is not possible without free recall of earlier options.)

⁸Since $\mathbb{E}_i w_i = \mathbb{E}_i [v_i - \max\{v_i - r_i, 0\}] = \mu_i - s_i$, the expected value of the parameter w_i is the stand-alone surplus available from box i .

⁹The data on online search behaviour for books, presented in the Introduction, shows that consumers differ in their choice of which seller to inspect first.

utility. An individual consumer is characterized by her list of reservation prices (r_1, r_2, \dots) and her list of realized payoffs (v_1, v_2, \dots) which via (4) generate the list (w_1, w_2, \dots) . This heterogeneous population of consumers selecting an option via optimal sequential search can equivalently be modelled as engaging in a discrete choice problem, where the type- (w_1, w_2, \dots) consumer selects the option with the highest w_i . The joint distribution of (w_1, w_2, \dots) in the consumer population then determines the aggregate demand for each option.

3 Ordered search with strategic sellers

We now put strategic sellers inside these boxes. Suppose there are a finite number of sellers, labelled $i = 1, 2, \dots$, which each supply a single variant of a product, where seller i has constant marginal cost of production c_i . Consumers want at most one product and have idiosyncratic match utilities for the product from seller i , denoted v_i , where v_i is not observed by the seller. A particular consumer incurs search cost s_i to inspect seller i , anticipates that seller i 's match utility comes from the CDF $F_i(v_i)$ and believes that her match utilities are independently distributed across sellers. Seller i chooses price \tilde{p}_i , and a consumer who buys from i obtains payoff $v_i - \tilde{p}_i$ (excluding her search costs). The consumer discovers seller i 's price \tilde{p}_i and the corresponding match utility v_i only after paying the search cost s_i . I assume that the consumer can freely return to a previously inspected product and there is no danger of a popular product being sold out.

Since the consumer's decision to inspect a seller, and the order in which she inspects sellers, depends on equilibrium, not actual, prices, write p_i for the equilibrium price from seller i (while \tilde{p}_i is that seller's actual, possibly non-equilibrium, price). In equilibrium we will require that $\tilde{p}_i = p_i$. Given a list of equilibrium prices (p_1, p_2, \dots) , the consumer's optimal search order is described by Pandora's rule. To understand what this means we calculate the reservation price of the lottery inside box i , which has anticipated payoff $\max\{v_i - p_i, 0\}$. As discussed in section 2, if $p_i > r_i$, where r_i in (1) is the reservation price for the match utility v_i , it is not worthwhile for this consumer ever to inspect seller i . If $p_i < r_i$, though, the reservation price of the box with random payoff $\max\{v_i - p_i, 0\}$ and search cost s_i is positive and equal to $r_i - p_i$.¹⁰ Therefore, according to Pandora's rule this consumer should first inspect the seller with the highest $r_i - p_i$, if this is positive, and

¹⁰The reservation price of this box is the x which satisfies $\mathbb{E}_i \max\{v_i - p_i - x, 0\} = s_i$, so that $x + p_i = r_i$.

keep searching until her maximum sampled payoff $v_k - \tilde{p}_k$ (where \tilde{p}_k is seller k 's actual price) is above all the $r_j - p_j$ for uninspected products. In general, consumers will differ in their reservation prices, and Figure 1 depicts the optimal search order with two sellers in terms of (r_1, r_2) and equilibrium prices (p_1, p_2) . In particular, a consumer's decision about which seller to inspect first is akin to a discrete choice problem where a consumer values option i at r_i and chooses the option with the highest payoff $r_i - p_i$ (or, as shown in the shaded region, the outside option zero if that is superior to engaging in search). Equilibrium occurs when (i) consumers choose their order of search optimally given the prices they anticipate sellers choose, and (ii) each seller chooses its price to maximize its profit given the consumer search order and the prices chosen by rival sellers, and this price coincides with the price anticipated by consumers.

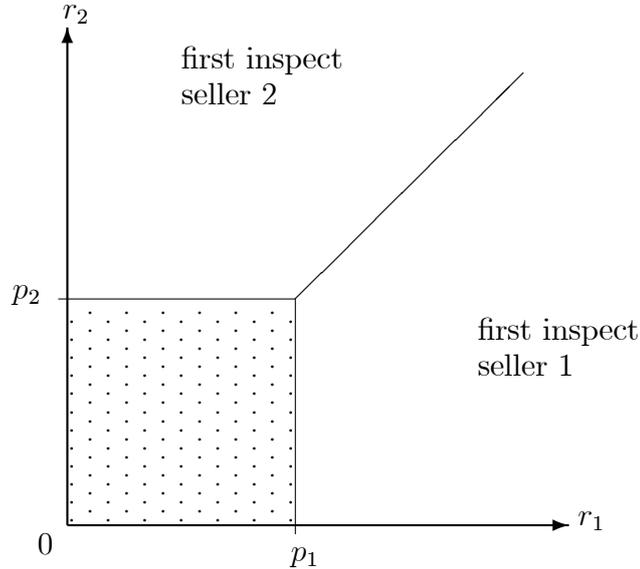


Figure 1: The pattern of search with two sellers

Equivalently, with equilibrium prices (p_1, p_2, \dots) and actual prices $(\tilde{p}_1, \tilde{p}_2, \dots)$ the discrete choice reformulation in section 2 shows that seller i 's demand is the fraction of consumers for whom the index

$$\min\{r_i - p_i, v_i - \tilde{p}_i\} \quad (5)$$

is positive and higher than the corresponding index from all rival sellers. Equilibrium in this market occurs when Bertrand equilibrium in actual prices $(\tilde{p}_1, \tilde{p}_2, \dots)$ given anticipated prices (p_1, p_2, \dots) coincides with these anticipated prices. In equilibrium (when $\tilde{p}_i = p_i$) expressions (4) and (5) imply that seller i 's demand consists of those consumers for whom

$w_i - p_i$ is positive and greater than $w_j - p_j$ for all other sellers. In the case of duopoly this pattern of demand in terms of (w_1, w_2) and equilibrium prices (p_1, p_2) is shown on Figure 2.

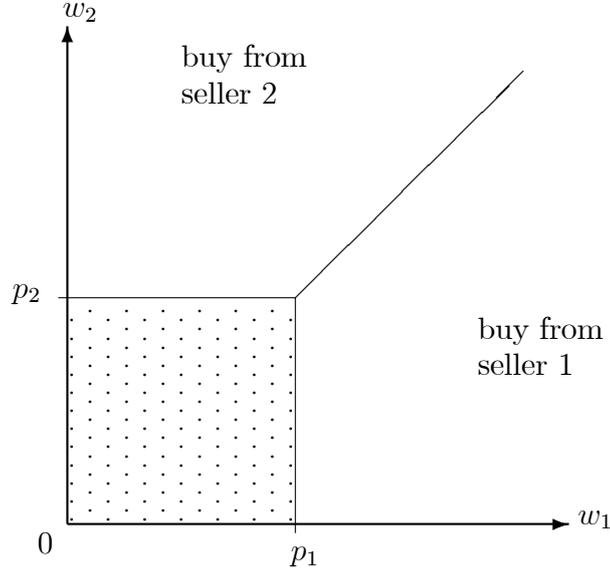


Figure 2: The pattern of demand with two sellers

From (5), a seller competes against its rivals (and the outside option) on two margins. If a consumer's preferences satisfy $r_i > v_i$, then for local deviations $\tilde{p}_i \approx p_i$ it is the size of $v_i - \tilde{p}_i$ which determines whether this consumer will buy from the seller, and the seller can affect this likelihood via its choice of price \tilde{p}_i . Otherwise, though, it is the size of $r_i - p_i$ which determines its demand, and this portion of demand does not depend on the seller's actual price \tilde{p}_i . If search costs become negligible for all sellers and consumers, then the case $r_i > v_i$ applies and this oligopoly model converges to an oligopoly model where consumers have complete information about match utilities (v_1, v_2, \dots) and actual prices $(\tilde{p}_1, \tilde{p}_2, \dots)$. Alternatively, if all consumers accurately know in advance their match utility v_i from seller i and have positive search costs s_i for this seller, it follows that $r_i < v_i$. Since this seller's demand is then perfectly inelastic with respect to its actual price \tilde{p}_i in the neighborhood of p_i , its profit cannot be maximized at $\tilde{p}_i = p_i$ and there is no equilibrium in which this seller has positive demand.¹¹ This is Diamond (1971)'s famous paradox.

To illustrate, consider an example with two sellers and costless production. Each consumer has match utility v_i for product $i = 1, 2$ which is uniformly distributed on the

¹¹More generally, the same issue arises when match utility for product i takes the binary form in which match utility takes some specified positive value with a specified probability and is zero otherwise. Here, whenever the match utility is positive it is greater than the reservation price.

interval $[0, 1]$ and has reservation price r_i for this product which is also (independently) uniformly distributed on $[0, 1]$. (From (1), the search cost corresponding to r_i is given by $s_i = \frac{1}{2}(1 - r_i)^2$ which lies in the interval $[0, \frac{1}{2}]$.) Then w_i in (4) is the minimum of two independent uniform variables, and so has support $[0, 1]$ and density $2(1 - w_i)$. Among those situations where both sellers are active, this example has a unique equilibrium and in this equilibrium each seller chooses the price $p \approx 0.49$.¹² (Details for this example are presented in Appendix B below.) As in Figure 1, if a consumer searches at all she will first inspect the seller for which she has the higher r_i . Each consumer searches in a deterministic order, but that order differs across consumers.¹³

If this example is modified so that search costs are zero—in which case $r_i \equiv 1$ and w_i in (5) is uniformly distributed on $[0, 1]$ —the symmetric equilibrium price is $p = \sqrt{2} - 1 \approx 0.41$, which is below the price with search frictions. It is intuitive that greater search frictions will tend to increase equilibrium prices, as in this example. Consider a particular seller in the market. If the inspection costs for this seller rise, it will tend to encounter fewer but “more desperate” consumers who have not found a good option from other sellers, and this will typically give it the opportunity to raise its price. Likewise, if the inspection cost for rival seller j increases, this shifts the distribution for r_j downwards, and hence shifts w_j downwards as well, and again this tends to give the seller an incentive to raise its price. In sum, if inspection costs rise, either for a single seller or across the market, this is likely to raise each seller’s equilibrium price. Thus, we expect there will be positive own- and cross-cost passthrough of inspection costs to prices, as well as the more familiar positive passthrough of production costs.

Multiple equilibria: The “double uniform” example above is a demand system which is smooth, in the sense that small changes in anticipated prices p_i do not lead to discrete changes in demands. In other situations—which include those commonly studied in the

¹²As usual, there are also less interesting equilibria in which consumers anticipate that seller i chooses such a high price that it is not worthwhile to inspect this seller, and then this seller might as well set this very high price. In this example, for instance, there is also an equilibrium where seller 1 sets price $p_1 = 1$ and no one inspects it, while seller 2 sets the monopoly price $p_2 = \frac{1}{2}$ and a consumer inspects it if $r_2 \geq \frac{1}{2}$ (and then buys if $v_2 \geq \frac{1}{2}$).

¹³Siekman (2016) analyzes a related model, where a consumer’s match utility takes the form (using adjusted notation) $v_i = \lambda a_i + (1 - \lambda)b_i$, where a_i is a component which the consumer knows in advance, b_i is a component for which the consumer needs to incur inspection costs to discover, and λ is a parameter which reflects how well-informed consumers are *ex ante*. In a symmetric market where consumers anticipate that all sellers choose the same price, a consumer will first inspect the seller for which she has the highest a_i . (A duopoly example where a and b are uniformly distributed is examined in detail.)

literature—the demand system is not smooth. Specifically, consider the situation in which each consumer considers sellers to be symmetric *ex ante*, so that (in the duopoly case) reservation prices on Figure 1 lie on the 45° line. Here, when one seller is expected to offer a lower price *all* consumers who search will choose to inspect it first. There is a strong possibility of multiple equilibria in such a market: the consumer search order depends sensitively on anticipated prices, while a seller’s price usually depends on where it is placed in the search order.

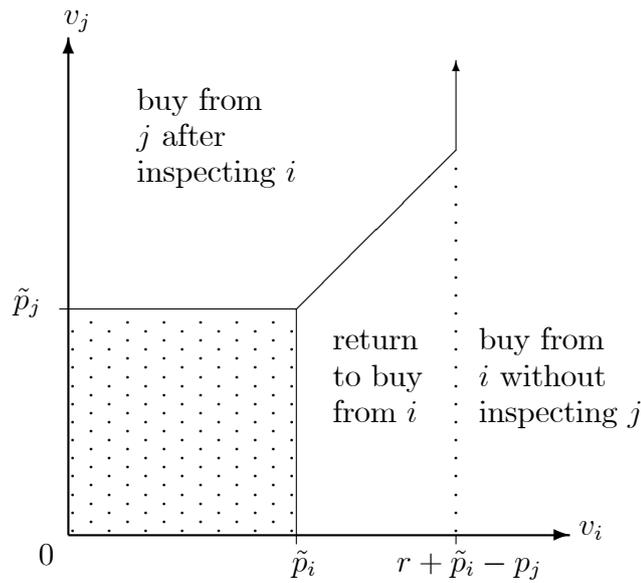


Figure 3: The pattern of demand with two sellers when $p_i < p_j$

To discuss this point in more detail, suppose each consumer has the same CDF $F(v)$ for match utility and the same inspection cost s for each seller, and hence has the same reservation price r for each seller’s product. Suppose also that each seller has the same production cost c . This is the framework analyzed in the influential models of Wolinsky (1986) and Anderson and Renault (1999) under the assumption that consumers search randomly through sellers. In contrast to these earlier papers, suppose that all consumers search through sellers in the same order.¹⁴ If the hazard rate for match utility, $f(v)/(1 - F(v))$, is strictly increasing in v , then more prominent sellers (i.e., sellers closer to the start of this search order) have more elastic demand than those sellers placed further back. For this reason, more prominent sellers typically set lower prices, which in turn rationalizes the assumed consumer search order. Intuitively, a seller inspected earlier in a consumer’s

¹⁴The following discussion is based on the analysis (for the uniform distribution) in Armstrong, Vickers, and Zhou (2009) and Zhou (2011).

search order knows that a prospective consumer is likely to have a superior outside option relative to the situation where a seller is inspected later—a later seller only encounters a consumer if that consumer was disappointed by her options so far—and with an increasing hazard rate, a seller who knows a consumer has a better outside option will choose to set a lower price. A more detailed argument for why a prominent seller faces more elastic demand is presented in Appendix C below.

To see how ordered search can be an equilibrium in a symmetric environment, suppose there are two symmetric sellers and look for an equilibrium where $p_i < p_j$ so that seller i is inspected first by all consumers. Here, the pattern of demand for the two sellers is shown in Figure 3.¹⁵ Appendix D below calculates equilibrium prices for various search costs when production is costless and match utility is uniformly distributed on the interval $[0, 1]$, and the left-hand graph on Figure 4 shows how the prominent firm chooses a lower price in equilibrium. (In this example, the maximum search cost s which allows consumer search in equilibrium is $s = \frac{1}{8}$. The prices of the two sellers are equal at the two extremes where search frictions are absent and where $s = \frac{1}{8}$ when both sellers set the monopoly price.) Thus there are two equilibria with ordered search, one where all consumers inspect seller 1 first and one where they inspect seller 2 first. There is also a third, symmetric, equilibrium, where exactly half the consumers first inspect each seller and where the two sellers set the same price. However, this symmetric equilibrium—which is the focus of the analysis in Wolinsky (1986) and Anderson and Renault (1999)—is *unstable*: if slightly more consumers first inspect one seller when anticipated prices are equal there is no equilibrium with equal prices. Thus, this is a classic “tipping” market, and we expect one low-price seller will be inspected first by all consumers even though sellers are symmetric *ex ante*.¹⁶ Consumers may well be worse off in an equilibrium with ordered search where they all inspect one seller first compared to the equilibrium with random search.¹⁷ Intuitively, faced with the increasing price path which goes with ordered search, consumers cease their search too

¹⁵This pattern is generated by observing that a consumer buys from seller i if $\min\{r - p_i, v_i - \tilde{p}_i\}$ is positive and greater than $\min\{r - p_j, v_j - \tilde{p}_j\}$.

¹⁶In situations where the hazard rate is decreasing, a seller which is first inspected by more consumers sets a higher price than its rival, and the unique and stable equilibrium has the two sellers setting the same price and half the consumers first inspect each seller. (In the knife-edge case of an exponential distribution for match utility, where the hazard rate is constant, a seller’s price does not depend on where it is in the search order, and no network effects are present.)

¹⁷Armstrong, Vickers, and Zhou (2009) and Zhou (2011) show this to be so in the case with a uniform distribution for match utilities. Zhou shows in this example that *all* prices can be higher with ordered search relative to the price with random search.

early and competition between sellers is weakened.

When consumers search in the same order the prominent seller has larger sales for two reasons: even with equal prices its demand would be larger because it is inspected first (its extra demand consists of the “north-east” region on Figure 3), while its lower price amplifies this effect. The result is that the distribution of sales across sellers is more skewed than it would be in a market with random search or in a market without search frictions. In this example, sales are equal for the two sellers when search frictions are absent, but the prominent seller sells up to twice as much as its rival when search costs are larger (see the right-hand graph on Figure 4).

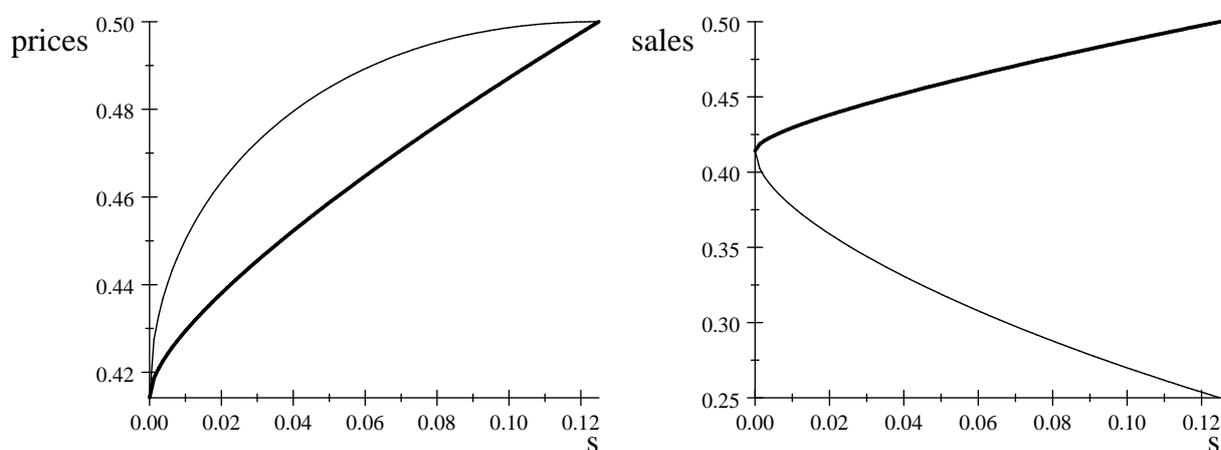


Figure 4: Prices and sales of the prominent seller (bold curves) and non-prominent seller

Likewise, the prominent seller obtains greater profit than its rival. (The prominent seller *could* choose the equilibrium price of its rival, in which case it has greater demand and more profit, but in general is even better off with another price.) In this example, the equilibrium profits of the two sellers are obtained by multiplying the respective curves on Figure 4, and are plotted on Figure 5. The impact on profit of an increase in search frictions will often differ for the two sellers.¹⁸ Profit for the prominent seller will rise with s since both its price and its demand do. The impact on the non-prominent seller’s profit, though, depends on two opposing forces—its price rises, but its demand is likely to fall—and the result is that its profit can be non-monotonic in the search cost as shown on the

¹⁸See also the discussion and Figure 2 in Zhou (2011). By contrast, with random search (and no outside option), Anderson and Renault (1999, Proposition 1) show that the symmetric equilibrium profit for each seller increases with the search cost, provided the hazard rate is increasing.

figure (which focusses on small search costs when the non-monotonicity is most apparent). When there are no search frictions both sellers choose the same price and obtain the same profit. As s increases both sellers' profits initially rise, but for larger s the non-prominent seller's profit falls as search frictions increase.

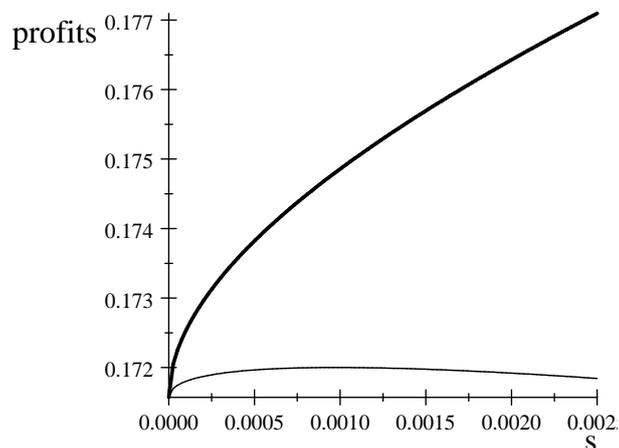


Figure 5: Profits of the prominent seller (bold) and non-prominent seller

This situation with ordered consumer search can usefully be contrasted with classical models of directed search in labour markets. For instance, Montgomery (1991) and Burdett, Shi, and Wright (2001) study a market with two firms, each with one vacancy, which advertise their wages to two potential workers. If both workers apply to the same firm, that firm chooses one at random and other worker is not able to apply to the second firm. Given a pair of wage offers there are usually three equilibria for the workers: two pure strategy equilibria in which one worker applies to one firm and the other worker to the other firm, and a symmetric mixed strategy equilibrium where each worker applies to a firm with the same probability (where workers are more likely to apply to the firm offering the higher wage). The authors argue that the most plausible of these equilibria is the symmetric mixed strategy equilibrium, since it is hard for workers to coordinate their applications to distinct firms, and firms then follow a symmetric pure strategy for wage choice. By contrast, in consumer markets coordination requires agents to follow the *same* order (e.g., if they all anticipate a lower price from one seller), and this seems easier to achieve than coordination on different orders.

Monopolistic competition: The existence of multiple equilibria can make it hard to do comparative statics, such as whether a higher-quality firm (where the match valuation

distribution comes from a better CDF) sets a higher price or is inspected first or whether a firm with a higher inspection cost is inspected later. For this reason a smooth demand system, where different consumers prefer to search in different orders, might work better (as well as often being more plausible). However, such a demand system can be cumbersome to work with beyond specific examples or without resorting to numerical methods.

One convenient way to simplify the framework is to study monopolistic competition with *many* symmetric sellers (Wolinsky, 1986), when a stable equilibrium with symmetric prices exists in a broad class of cases.¹⁹ When a consumer expects all sellers to offer the same price $p < r$, a consumer will search until she finds a product with $v \geq r$ and will never return to a seller inspected earlier.²⁰ Thus a seller has no “return demand” (by which we mean consumers who come back to purchase after investigating other sellers) which was the source of the incentive for prominent firms to set lower prices, and a seller sets the same price regardless of its place in the search order. The result is that all sellers set the same price and consumers do not care how other consumers choose to search. The symmetric equilibrium price, p say, when consumers are also symmetric is derived as follows. Consider a seller who meets a consumer. If it chooses price \tilde{p} the consumer will buy from it if $v - \tilde{p} \geq r - p$, and so its profit from this consumer is $(\tilde{p} - c) \times [1 - F(\tilde{p} + r - p)]$. In equilibrium, this must be maximized at $\tilde{p} = p$, which yields the unique first-order condition

$$p = c + \frac{1 - F(r)}{f(r)} . \quad (6)$$

The equilibrium markup and industry profit in this market, $(1 - F(r))/f(r)$, depends on the shape of the CDF $F(v)$ and the magnitude of search frictions. (How sales and profits are divided across the sellers is not pinned down in this framework.) Consumers have an incentive to participate in this market provided that p in (6) is below r . When the hazard rate $f/(1 - F)$ increases, the equilibrium price in (6) decreases with r and hence increases with the search cost s . In such cases, a reduction in search frictions yields a double benefit to consumers: their average match utility is higher and the price they pay is lower. Although this model of monopolistic competition does not necessarily involve

¹⁹Anderson and Renault (1999) show that a symmetric equilibrium with monopolistic competition exists provided that the hazard rate is increasing.

²⁰Anderson and Renault (2015) discuss another way to obtain this simplifying feature. They suppose that the distribution for match utility from a given seller has an atom at zero combined with a continuous distribution with a support well away from zero, and they find equilibria with ordered search where each seller sets price so that any consumer it encounters buys immediately when she has non-zero match utility.

ordered search, it is useful starting point for some of the applications and extensions to this basic framework presented in the next section.

4 Applications and extensions

4.1 Obfuscation

One issue of interest is whether it might ever be profitable for a firm deliberately to raise its own inspection cost—that is, to “obfuscate”—in order to make its rival the prominent seller. For instance, in the UK some well-known insurance companies currently advertise that their products do *not* appear on price-comparison websites. To discuss this possibility, suppose the initial situation is that there are two symmetric sellers and no search frictions (so the market is a duopoly version of Perloff and Salop, 1985). Then consumers will investigate both sellers’ offers and buy from the seller with the higher $v_i - p_i$ (if this is positive). In regular cases the equilibrium will be symmetric, and firms obtain equal profit.

If one firm artificially introduces an inspection cost, $s > 0$, this will induce all consumers to inspect the rival first since they have nothing to lose by doing so. The new equilibrium prices will, given an increasing hazard rate, involve the prominent rival choosing the lower price, which reinforces the incentive to inspect this firm first. However, this lower price will typically still be higher than the equilibrium price without search frictions, and this could compensate the obfuscating seller for its disadvantaged position. As discussed in section 3, the non-prominent firm has an incentive to raise its price since the consumers it encounters are not satisfied with their offer from the prominent seller, and because prices are strategic complements this induces the prominent firm to raise its own price too. For instance, consider the example depicted in Figure 5 above. When one firm introduces a small inspection cost, the equilibrium profits are as shown on this figure, and we see that a small inspection cost boosts the obfuscating firm’s profit a little (although the rival’s profit is boosted more).

Wilson (2010) analyzes this question using a different duopoly model with a homogeneous product. There are two kinds of consumers: those who can see both prices without cost (even with obfuscation), and those who incur obfuscation costs if they are artificially introduced. In this market with a homogeneous product, without obfuscation there is Bertrand competition and zero profit. Wilson shows that it is always in a firm’s interest to obfuscate, with the result that firms choose their prices according to an asymmetric

mixed strategy, costly searchers inspect the transparent firm first, and both firms make positive profit. One difference between Wilson’s model and the one presented here is that in my model the obfuscating firm sets a higher price, while in Wilson’s model that firm (on average) sets a lower price.²¹

The example on Figure 5 was rather delicate, and a seller had only a small incentive to obfuscate. More striking and robust results are seen in the context of a multiproduct monopolist considering how to price and present its products. (We maintain the assumption that the consumer wishes to buy at most one product from the seller.) For simplicity, consider a situation where the seller has costless production and supplies two symmetric products, 1 and 2, where each consumer’s match utility for product i is an independent draw from the CDF $F(v)$. Suppose that unless the seller deliberately obfuscates, a consumer observes both prices and both match utilities from the start, and chooses the product with the higher surplus $v_i - p_i$ (if this is positive). In regular cases, the seller will then choose the same price p for both products, which as shown on Figure 4a below is chosen to maximize

$$p[1 - F^2(p)] . \tag{7}$$

However, the seller can always do better than this by making one product, say product 2, costly to inspect, while leaving product 1 costless to inspect (which is therefore inspected first by all consumers). Indeed, we can show that the seller can then obtain as profit the maximum value of

$$p_1[1 - F(p_1)] + F(p_1)p_2[1 - F(p_2)] . \tag{8}$$

Here, (8) represents the profit obtained if the seller could first offer product 1 to the consumer, at price p_1 , and the consumer myopically chooses whether or not to buy this product myopically, without considering the subsequent option to buy product 2. Since the profit in (8) coincides with the “frictionless” profit (7) with uniform prices $p_1 = p_2 = p$, it is clear that the maximum profit in (8) is above the maximum profit without obfuscation.

Maximizing (8) involves choosing $p_2 = p^M$, where p^M is the monopoly price for a single product, i.e., which maximizes $p[1 - F(p)]$. Suppose the seller makes the consumer incur

²¹Ellison and Wolitzky (2012) also study a model with a homogeneous product but where consumers do not observe a firm’s chosen inspection cost in advance, and so a firm cannot use obfuscation to influence search order. They assume that search costs are convex, in the sense that the more time a consumer spends obtaining one seller’s offer, the less inclined she is to investigate other sellers. Armstrong and Zhou (2016) study another way to create search frictions artificially, which is when a seller chooses to make it hard for a consumer to return to buy from it after inspecting a rival seller, for instance by offering the consumer a discount if she chooses to buy immediately without further search.

the search cost s for the second product so that a consumer is just willing to inspect this product when priced at p^M if she has no other option, so that s satisfies

$$s = \mathbb{E} \max\{v - p^M, 0\}. \quad (9)$$

Suppose also that the seller chooses p_1 to maximize (8) given $p_2 = p^M$, which implies $p_1 > p^M = p_2$. Although consumers cannot observe p_2 in advance, it is an equilibrium for consumers to anticipate the price $p_2 = p^M$ and for the seller to maximize its profit by choosing this price. (Regardless of the price p_1 , if consumers anticipate $p_2 = p^M$, and have to incur the search cost s in (9) to find the corresponding match utility, Pandora's rule states they will only choose to inspect this product if their match utility from the first product, v_1 , is below p_1 . Therefore, no consumer will ever return to buy the first product if they inspect the second, and so the seller chooses p_2 maximize its profit as if it only sold this single product, i.e., it chooses $p_2 = p^M$.) Figure 6 shows how the pattern of demand is altered by the introduction of an inspection cost for the second product.

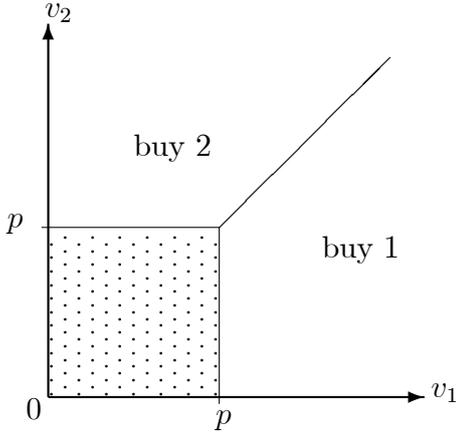


Figure 6a: frictionless selling

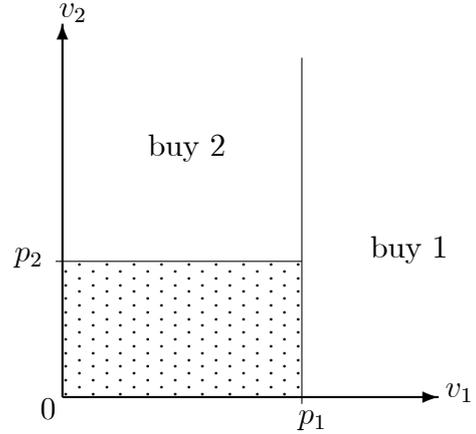


Figure 6b: product 2 obfuscated

In sum, it is a profitable and credible strategy for the seller to choose the two prices p_1 and p_2 which maximize (8) and to introduce the artificial inspection cost s defined in (9) for product 2, and this strategy generates higher profit than the situation without search frictions.²² This policy involves an expensive product being prominently displayed, perhaps at eye level, while a cheaper product is deliberately displayed inconveniently so

²²A similar argument applies with more than two products. For instance, with three products the seller can choose its prices to maximize $p_1[1 - F(p_1)] + F(p_1)[p_2[1 - F(p_2)] + F(p_2)p_3[1 - F(p_3)]]$, which entails $p_1 > p_2 > p_3 = p^M$. Here product 1 is expensive but presented without obfuscation, product 2 has medium obfuscation and a medium price, while product 3 is cheap but hard to inspect.

that the consumer has to look “high and low” for a good deal. Somewhat akin to the model in Deneckere and McAfee (1996), where a seller deliberately reduces the quality of one product variant in order to facilitate price discrimination, here the seller chooses to “damage” its retail environment. In regular cases, the prominent product’s price p_1 will rise while the obscure product’s price p_2 falls, relative to the situation without search frictions.²³ In such cases, consumers are harmed by the obfuscation policy: the search cost (9) eliminates all consumer surplus from product 2, while the price for product 1 is increased.

I presented here the simplest version of the model in Petrikaite (2016), who analyzes cases with more products, with asymmetric products, and with competition between multiproduct sellers. For instance, she shows that when products differ in quality, the seller will choose to obfuscate the lower quality product.²⁴ Gamp (2016) studies a related model where all prices are observed by consumers from the start, and the seller can choose to make it harder to inspect a product’s match utility. This framework is conceptually less complicated, in that we need not consider what the equilibrium hidden price is for the obscure product. Intuition for why obfuscation is profitable in this framework with advertised prices is clear-cut. Suppose there are two asymmetric products and in the situation without artificial search costs the most profitable strategy involves a higher markup over cost for product 1. If the seller keeps these prices unchanged but introduces a modest search cost for the less profitable product 2, this does not affect total demand for the two products but diverts some consumers to the more profitable product which boosts profit. Gamp shows that the seller typically wishes to guide consumers towards “niche” products first, where some consumers are willing to pay high prices, while more “mass market” products are obfuscated and used as insurance by the firm in case a consumer does not like its niche products. He also shows by example that obfuscation can increase total welfare: the

²³Starting from the uniform price which maximizes $p[1 - F^2(p)]$, one can check that profit in (8) is locally increasing in p_1 and decreasing in p_2 . Note in particular that with coordinated pricing the prominent product is likely to have a higher price than more obscure products, which is the reverse pattern predicted in section 3 when products are supplied by separate sellers. See Zhou (2009) for further discussion of this comparison.

²⁴Salop (1977) studied a related model much earlier. In his model a monopolist supplies a homogeneous product and offers consumers repeated draws from a price distribution chosen by the seller, where each draw incurs an exogenous search cost which varies across consumers. (All consumers obtain the first price quote for free.) When consumer valuations for its product are positively correlated with their search costs, Salop shows that the monopolist can extract more profit by choosing its price stochastically, so that consumers who choose to search more will obtain a lower price, than by offering a uniform price.

benefits of the low price for the obscure product outweigh the extra search costs consumers must incur.

4.2 Repeat business

In markets with ordered consumer search, tiny asymmetries between sellers can translate into major differences in their sales and profits. There is a significant advantage to a seller being placed early in the consumer search order, simply because it meets more consumers than its rivals placed further back. Since consumers in near-symmetric situations are near-indifferent about which seller to inspect first, it does not take much inducement to favour one seller in the search order, which then causes that seller to enjoy a discrete jump in its sales and profit. For instance, consider the case of monopolistic competition, where the equilibrium price from each seller is (6) regardless of the search order. With random search each seller obtains negligible sales and profit, while if one seller manages to be placed first in the search order—perhaps because it is slightly easier to inspect or it has a slightly superior CDF for match utility—its profit jumps to

$$(p - c)[1 - F(r)] = \frac{[1 - F(r)]^2}{f(r)}. \quad (10)$$

One way to introduce small asymmetries between sellers is in a framework where sellers supply various products over time and consumers wish to purchase each of these products. (This discussion is taken from Armstrong and Zhou, 2011, section 3.) It is plausible that a consumer who previously purchased from one seller might first inspect the same seller when she searches for a second product, even if there is no correlation in her match utilities for the various products from the same seller.²⁵ For instance, she may have the contact details of this seller to hand, in which case the seller’s inspection cost is a little lower than that of its rivals. Here, the supplier of one product to a consumer becomes prominent for that consumer when she starts the search process for the next product.

In more detail, suppose there are two product categories, 1 and 2, which consumers buy sequentially—for instance, a bank account first and then a mortgage—both of which are jointly supplied by many sellers in monopolistic competition. For product category $i = 1, 2$, the reservation price is r_i , the CDF for the match utility is $F_i(v)$, the production

²⁵If a consumer’s match utilities for a given seller’s products were positively correlated, then presumably a consumer has a strict incentive to start her next search at the seller she purchased from previously. Such a model is considerably more complicated to analyze, however.

cost is c_i , while the factor sellers use to discount profits from the second product when they supply the first product is δ . As in expression (10), when a supplier sells product 1 to a consumer, that consumer will go on to generate expected profit from the second product equal to $[1 - F_2(r_2)]^2/f_2(r_2)$. (A seller sets the same price for product 2 regardless of whether or not a consumer purchased product 1 from it.) Anticipating this profit the equilibrium price for the first product will be

$$p_1^* = c_1 + \frac{1 - F_1(r_1)}{f_1(r_1)} - \delta \frac{[1 - F_2(r_2)]^2}{f_2(r_2)} .$$

Thus, the promise of profit from the second product induces firms to lower the price for the first product, perhaps to below its cost. When search frictions are larger for the second product, this will usually lead firms to offer a lower price for the first product. This outcome is reminiscent of markets with switching costs. When switching costs are small, though, they have little impact on the outcome, while here a tiny “default bias” leads to large effects. Moreover, since the default bias has no impact on the product-2 price but causes the product-1 price to fall, this form of default bias will benefit consumers.

Garcia and Shelegia (2015) present a related model, where a consumer is inclined to start her search process at a seller where she has observed someone *else* make a purchase. Again, because making a sale makes the seller prominent in a consumer’s mind, albeit a different consumer’s mind, sellers have an incentive to price low in equilibrium.

4.3 Non-price advertising

There are situations in which consumers are more likely to first inspect those sellers which spend the most on their marketing efforts. This behaviour might be due to psychological factors, if consumers most easily recall sellers from which they have seen adverts. Alternatively, as discussed in more detail in the remainder of this section, rational consumers could use advertising intensity to guide their search towards products which are likely to be cheaper or more suitable. We will discuss three broad forms of non-price advertising: (a) a traditional advertising campaign using television, print media, and so on; (b) payments for prominent positions or special endorsements within a retail outlet, and (c) auctions for prominent positions in sponsored search.

Consider first situation (a), where consumers are assumed to be able to accurately observe relative spending on ad campaigns and use this to coordinate on which seller to inspect first, anticipating that the price will be lower from the seller which spends the most

on advertising. As discussed in section 3, assuming the hazard-rate condition holds, when a fraction of consumers use the rule of thumb of first inspecting that seller which advertises the most, that seller will choose a lower price and it is optimal for an individual consumer to mimic that search order. Consider a two-stage framework where two symmetric risk-neutral sellers first choose their advertising intensities and then choose their prices. Let π_H be a seller's equilibrium profit (excluding advertising costs) in the second stage if it is prominent and let $\pi_L < \pi_H$ be the corresponding profit for the non-prominent seller. The first stage, in which sellers compete to become prominent by advertising the most, is then a symmetric all-pay auction with complete information. It is clear that no pure strategy equilibrium for advertising can exist, since spending a little more than your rival on advertising generates a discrete jump in profit. However, a symmetric mixed strategy equilibrium for advertising exists. If $H(a)$ denotes the equilibrium probability that a seller spends less than a on advertising in the first stage, a seller's expected profit when it spends a on advertising is

$$H(a)\pi_H + [1 - H(a)]\pi_L - a . \quad (11)$$

(With probability $H(a)$ it wins the contest and enjoys high profit π_H , and otherwise it becomes the less prominent seller.) Since a seller can obtain profit π_L by not advertising at all, in equilibrium profit in (11) must be identically equal to π_L over the range of advertising used, so that $H(a) = \frac{a}{\pi_H - \pi_L}$ for $0 \leq a \leq \pi_H - \pi_L$. Thus in equilibrium each seller chooses its advertising according to a uniform distribution on the interval $[0, \pi_H - \pi_L]$.²⁶ The seller which advertises more will offer the lower price. Each seller spends $\frac{1}{2}(\pi_H - \pi_L)$ on advertising on average, and since we expect that the benefit of prominence, $\pi_H - \pi_L$, is higher when the search cost s is higher (for instance, see Figure 6), this model suggests that there is more advertising in markets with higher search frictions. Competition to advertise the most acts to dissipate profit so that sellers earn average profit equal to the low level associated with not being prominent, π_L , which as shown on Figure 6 can be non-monotonic in the search cost s .

This analysis is similar to Bagwell and Ramey (1994), who analyze a model where sellers offer a homogeneous product and consumers choose from where to buy before they observe

²⁶Of course, here and elsewhere in this section another equilibrium exists: consumers do not respond to advertising and choose to inspect sellers in random order, in which case sellers are not prepared to invest in advertising. However, this random search equilibrium is fragile, in the sense that if some fraction of consumers do use advertising to guide their search order, it pays all consumers to do the same.

the seller's choice of price. Because sellers in their model have increasing returns to scale, when more consumers turn up to buy from a seller that seller offers a lower price. Thus, consumers benefit from coordinating on a seller, and a fraction of consumers rationally go to the seller which advertises the most heavily as the means with which to do this. By contrast, the model above assumes no economies of scale but instead uses the feature that sellers who attract more first visits have more elastic demand. Bagwell and Ramey (1994) also report empirical studies (for eye-glasses, liquor, and prescription drugs) which show how prices were lower and market structure was more concentrated in markets where advertising was permitted, even when prices could not be advertised.

Haan and Moraga-Gonzalez (2011) analyze another related model, where a seller's share of first-inspections is continuous in its choice of advertising intensity instead of the "winner take all" formulation discussed above. They focus on a symmetric equilibrium where sellers advertise with the same intensity, obtain the same share of initial searches, and so charge the same price. As with the simpler model presented above, they find that net profits can be non-monotonic in the search cost while advertising intensity typically increases with the search cost. Their model involves behavioural rather than fully rational consumers: if one seller advertises slightly more heavily than its rival and this induces more consumers to inspect it first, that seller will charge a lower price and hence *all* consumers should rationally visit that seller first.

Consider next a situation where consumers use marketing efforts to find more a suitable product, rather than the cheapest supplier of otherwise symmetric products. To study this situation, suppose that competing sellers are invited by a retailer to pay for a prominent position or endorsement within its store. For instance, a bookstore might charge publishers a fee for a book to be that retailer's "book of the week". Armstrong, Vickers and Zhou (2009, section 3) present a model of monopolistic competition where sellers differ in the quality of the product they supply (in the sense that their distributions for match utilities are ordered in terms of first-order stochastic dominance). In more detail, a seller with quality parameter q has a product which provides a match utility with CDF $F_q(v)$. A higher q improves the distribution, in the sense of first-order stochastic dominance, so that $F_q(v)$ decreases with q . There are many sellers, and the distribution of q in the population of sellers has CDF $G(q)$. For simplicity, suppose that each seller has costless production, and let the equilibrium price of the type- q seller be denoted p_q . The assumption of monopolistic

competition implies a seller's price is not altered if it is made prominent in the consumer search order.

Initially, a consumer must search randomly through the sellers without being able to identify a seller's quality before inspection, where there is a search cost s per seller. In this situation with random search, let V denote a consumer's optimal expected payoff from random search in this market. Since with many sellers a consumer will not return to a seller already inspected, the type- q seller will choose its price to maximize $p[1 - F_q(p + V)]$. Clearly, a higher- q seller will obtain greater expected profit from a consumer it encounters than a seller with lower q , and so the seller who is willing to pay the most for prominence has the highest possible q . (Its demand is greater for any given price due to stochastic dominance.) Whether a higher- q seller sets a higher price, and whether a consumer obtains a higher payoff from inspecting a higher- q seller first, depends on the demand curve $1 - F_q(v)$ and how it depends on q . In those cases where a consumer wishes to inspect the highest- q seller first, the ability of sellers to pay for a prominent position guides consumers towards better, and better value, products, and this improves consumer surplus and total profit relative to the situation with random search.

To illustrate, consider the case where $1 - F_q(v) = q[1 - F(v)]$, where $0 \leq q \leq 1$ and $F(v)$ is a reference CDF.²⁷ Here, a type- q seller provides a product which with probability $1 - q$ is worthless to the consumer and with probability q the product has match utility with CDF $F(v)$. This formulation has the property that the equilibrium price from a type- q seller does not depend on q , and each seller choose the price p which maximizes $p[1 - F(p + V)]$. Because price does not depend on q , it is clear that a consumer prefers to inspect a seller with higher q first, since that makes it more likely she will find a suitable product at her first attempt.²⁸

The third and final scenario in this section involves sponsored search. Perhaps the most well-known paper on this topic which allows for optimal search by consumers is Athey and

²⁷Alternatively, Armstrong, Vickers, and Zhou (2009) examine the case where both v and q are uniformly distributed.

²⁸Eliasz and Spiegel (2011) analyze the related situation (with the same specification for $F_q(v)$) where an intermediary *only* presents a population of high- q sellers to consumers, in contrast to the situation in the text where the pool of sellers is unchanged but one seller is selected out of the pool for prominent display. The fact that the pool presented by the intermediary contains only relatively good sellers affects a consumer's search payoff V and the equilibrium price. While consumers would prefer that the intermediary only presents the very highest q sellers to them, Eliasz and Spiegel (2011) find that the intermediary often has an incentive to "dilute" the pool with some lower- q sellers, in order to relax competition and boost the revenue they can extract from sellers.

Ellison (2011).²⁹ As above, sellers differ in how likely consumers are to find their products suitable, although here prices are not modelled explicitly. Specifically, a type- q seller offers a product which each consumer has probability q of finding suitable, in which case the consumer obtains payoff 1 from the product, and with remaining probability $1 - q$ she obtains payoff zero from the product, while a seller obtains payoff 1 each time its product is selected. Here, sellers differ only in their probability of being suitable, q , which is private information. Clearly, a consumer will select the first product which is found to be suitable, and to minimize her search costs she would like to inspect sellers in order of decreasing q if she could identify that order. Provided that better sellers do choose to pay more, it is optimal for consumers to inspect sellers—that is, to click on their links—in the order they appear on the results page.

In broad terms, sponsored search auctions allocate sellers to prominent positions, require sellers to pay the search engine fee each time a consumer clicks on their link, and use a generalized second-price auction format. With this format, higher- q sellers are indeed willing to bid more in the auction, with the result that consumers rationally sift through their options in the order they appear on the search engine’s results page. In this setting, the observation that consumers click more often on results placed higher up the sponsored results page is driven by the information content of the ordering, rather than “inert” consumers who mechanically follow the suggested order. As Athey and Ellison put it (page 1215): “a search engine that presents sponsored links should be thought of as an information intermediary that contributes to welfare by providing information (in the form of an ordered list) that allows consumers to search more efficiently”.³⁰

The models presented in this section provide the best case for non-price advertising to be used to guide search: the seller which pays the most to achieve prominence is the seller which consumers would like to inspect first. In other situations, this coincidence of interests need not hold perfectly or at all, and advertising expenditure is then a less reliable guide for consumers. For example, in the model of sponsored search, sellers might differ both in their likelihood of being suitable and in their expected profit from a click.

²⁹See also Chen and He (2011) for a related model.

³⁰The position auction model just discussed assumes that position is the only way that sellers can communicate with consumers. In reality, the displayed link also contains a small portion of text (or sometimes a small photograph) which can also guide search. Jeziorski and Segal (2015) document empirically how consumers click on the displayed links, and find that a large portion of consumers miss out some links as they move down the page, and many click on higher links after clicking on lower links.

In such a setting, the seller which is willing to bid the most per-click is not always the seller which the most consumers will find suitable. Because of this, search engines usually use measures of relevance as well as willingness-to-pay per click when they determine the order of sponsored links.³¹

4.4 Price advertising

The final extension to the basic model supposes that consumers observe all prices at the start of their search process, and given these prices they choose the order to inspect sellers to discover the corresponding match utility.³² Unlike the model in section 3 where prices were hidden until inspection, when prices are advertised they can be used to influence a consumer's search order. In addition, because prices are known in advance, the network effects discussed in section 3 do not arise. In the model in section 3 a seller which is prominent often chooses to set a lower price, while in the current scenario with price advertising this causality is reversed and a seller can become prominent by virtue of advertising a lower price.

Consider first the monopolistic competition framework above, where the equilibrium price in the absence of price advertising was given in (6). When firms advertise their price, the only equilibrium involves all sellers choosing price equal to marginal cost.³³ Setting price equal to cost is an equilibrium because when all its rivals do so, a seller cannot do better with a higher price since no consumers will ever inspect it. However, if all rivals advertised a price $p > c$, a seller could boost its profit by advertising a slightly lower price which attracts all consumers to inspect it first. Thus, even though there may be significant search frictions and horizontal product differentiation, the ability to use prices to become prominent in the search order drives price down to cost.

This framework is more interesting in oligopoly, when, in contrast to the situation in

³¹See Athey and Ellison (2011, section VI) for further discussion. See also Gomes (2014) for a model in which an intermediary selects one seller from a pool of heterogeneous sellers to display to consumers. The revenue-maximizing mechanism for the intermediary, when it can only obtain revenue from sellers, is a "scoring auction" which balances a seller's willingness-to-pay for display with the consumer's valuation for clicking on the seller's link.

³²We continue to assume that consumers cannot purchase from a seller without incurring the search cost. Even in cases where it is possible to buy without incurring the search cost and without observing the realized match utility, it will be optimal to inspect products before purchase if the search cost is small enough.

³³If sellers can *choose* whether or not to advertise their price (and can advertise costlessly), the following discussion remains valid if consumers anticipate that a seller which does not reveal its price in advance has in fact set a high price.

which prices are hidden, equilibrium prices often *decrease* with search frictions. Consider first the (apparently) simple situation where consumers view sellers as symmetric *ex ante*. In this case, consumers will all choose to first inspect the firm with the lowest advertised price and so slightly undercutting a rival’s price causes a discrete jump in profit.³⁴ In such situations, prices will be chosen according to mixed strategies in equilibrium. However, unlike more familiar models of random pricing, such as Varian (1980), here the prices offered by higher-price sellers continue to affect the demand of the “winning” seller, since some consumers will buy from more expensive sellers if those products have a better match utility. This combination of product differentiation and mixed strategies is usually hard to solve explicitly.³⁵

The framework is made more tractable, as well as often more plausible, if consumer demand is smoothed by supposing that consumers differ *ex ante* in which seller they wish to inspect first, given a set of advertised prices.³⁶ Given advertised prices (p_1, \dots, p_n) , expression (5) implies that a consumer chooses to buy from the seller with the highest value of $\min\{r_i - p_i, v_i - p_i\} = w_i - p_i$, provided this is positive, where w_i is defined in (4). This demand system in the case of duopoly is as depicted above in Figure 2, except that (p_1, p_2) on the figure represent any advertised prices not necessarily equilibrium prices. The demand for each seller’s product in terms of prices can then be calculated given the distribution for (w_1, \dots, w_n) in the consumer population. One can use this demand system to derive the equilibrium prices as with any discrete choice Bertrand oligopoly. In other words, the convenient Bertrand model, which assumes consumers are fully informed from start, can also be used to study this far more complex market where consumers uncover market information in an optimal sequential manner. To do so requires that the relevant consumer valuations are (w_1, w_2, \dots) rather than the (higher) raw valuations (v_1, v_2, \dots) .

³⁴Baye, Gatti, Kattuman, and Morgan (2009) document empirically that a seller on a price-comparison website has demand which jumps substantially when it reduces its price to become the lowest-price seller on the website.

³⁵Armstrong and Zhou (2011, section 2) analyze one version which can be solved, where a consumer’s match utility from one seller is negatively correlated with her value for the rival’s product. This implies that a consumer knows her payoff at the second seller once she inspects the first, and so there is no “return demand”, and as usual this simplifies the demand functions. In this stylized framework we find that equilibrium prices decrease with the size of search frictions.

³⁶Choi, Dai, and Kim (2016), Haan, Moraga-Gonzalez, and Petrikaite (2015) and Shen (2014) study models of price advertising in situations of this form. In each paper, the authors suppose that a consumer’s match utility is the sum of two independent components, and the consumer knows one component from the start which gives rise to *ex ante* brand preferences. (The paper by Siekman (2016), discussed earlier in the context of hidden prices, also models match utility in this manner.) Choi, Dai, and Kim (2016) also use a discrete choice re-formulation of the search model to obtain their results.

One can use this approach to obtain familiar comparative statics results about the impact on price of changes in production costs or the number of sellers. In addition, equilibrium prices in a market with advertised prices are likely to be lower than the corresponding market with hidden prices. Intuitively, a seller’s demand is more elastic with respect to changes in its price with advertised prices than when its price is only discovered after the consumer pays the search cost. This can be seen formally from expression (5), where a seller’s loss of demand when it increases its price \tilde{p}_i is smaller than if both p_i and \tilde{p}_i were increased simultaneously (as is the case with advertised prices).

As mentioned earlier, more surprising comparative statics in this market with advertised prices are obtained with respect to changes in search costs. Since the advantage of being prominent increases with search frictions, with advertised prices it is plausible that an increase in search frictions intensifies competition to be prominent, with the result that equilibrium prices fall. Again, this can be understood using the discrete choice perspective, where an increase in search costs reduces a consumer’s list of reservation prices r_i and hence causes her adjusted valuations w_i to fall stochastically. It is natural that, in many cases, a fall in the distribution for valuations in an oligopoly discrete choice model will lead to lower equilibrium prices.

To illustrate these observations, consider the “double uniform” example from section 3, where variation in r_i stems from heterogeneous inspection costs s_i rather than from heterogeneous distributions for match utility F_i . Using Figure 2 one can calculate the symmetric equilibrium advertised price which in this example is about 0.31. In the same example with hidden prices, we saw that the symmetric equilibrium price was about 0.49 which, as expected, is above the equilibrium advertised price. If the example is modified so that search costs are identically zero, in which case w_i is uniformly distributed on $[0, 1]$, the equilibrium price (with or without price advertising) would be $\sqrt{2} - 1 \approx 0.41$, which is higher than in the situation with search frictions and advertised prices.³⁷

5 Searching questions

This article has discussed a range of situations in which consumers plausibly search through their options in a deliberate order. We saw how sellers which were first in line often had an

³⁷Although the price is lower when search frictions are present, the price reduction in this example is not sufficient to outweigh the direct harm to consumers caused by costly search, and consumer surplus is higher when there are no search costs despite the higher price.

incentive to set lower prices than those further back, and this gave an individual consumer an incentive to mimic the search order followed by other consumers. Likewise, consumers wish to click first on the advertiser which supplies the most relevant product, and advertisers with the most relevant products often had the greatest incentive to pay for a prominent position. These kinds of self-fulfilling prophecies can make ordered search a stable equilibrium.³⁸ Because the seller which is inspected first typically makes greater profit than its rivals, sellers have an incentive to move to the front of the queue for consumers. Ways they might have to do this include advertising intensively, advertising the lowest price, selling other products to the pool of consumers in the past, forming a cluster with other sellers or selling multiple varieties under one roof.³⁹

There are two broad themes in this paper. One is more methodological, which is to show how the complex scenario with optimal sequential search can be reformulated as a static discrete choice problem without search frictions. An implication of this is that oligopoly search models often “work better” when the demand system is smoothed by giving consumers heterogeneous preferences over which seller to inspect first. (This can be done by making consumers heterogeneous in their brand preferences or heterogeneous with respect to their seller-specific search costs.) By contrast, when each consumer views sellers as symmetric *ex ante*, small changes in anticipated or advertised prices lead to a discrete jump in a seller’s demand. When prices are hidden, this often gives rise to the multiple equilibria discussed above, while when prices are advertised this feature leads to the use of mixed pricing strategies in equilibrium.

A second theme is how changes in search costs affect outcomes. A quite robust result

³⁸Similar effects can arise in the labour market. Mailath, Samuelson, and Shaked (2000) present a model in which employers search for workers, and *ex ante* identical workers decide whether to invest in skills before they enter the labour market. Workers are exogenously labelled with one of two colours, which for a given skill level have no effect on their productivity, and colour is observed by employers before search. In this framework asymmetric equilibria exist in which employers concentrate their search efforts on workers of one colour, while workers of the other colour are less likely to invest in skills.

³⁹Space constraints prevented detailed discussion of the important issues of clustering and multi-variety selling. Fischer and Harrington (1996) present a model with differentiated products, one cluster location and many “peripheral” sellers, where consumers who choose their search order based on rational expectations of prices chosen by sellers in the cluster and by peripheral sellers and their own idiosyncratic search costs. Moraga-Gonzalez and Petrikaite (2013) discuss whether consumers choose to inspect a multi-variety seller (which sets a relatively high price) before single-variety sellers. Among other results, they show that a merger of single-variety sellers can be profitable for merging firms but can reduce the profit of non-merging firms since they are pushed further back in the consumer search order. This contrasts with a more standard analysis of Bertrand price competition, where a merger typically raises the profits of non-merging firms.

was that a multiproduct seller had an incentive to “damage” its retail environment: instead of permitting consumers to see product characteristics transparently, it was more profitable to deliberately make it costly to inspect some products. When separate sellers each supply a single product and prices are hidden, a rise in one seller’s inspection cost typically causes that seller and its rivals to raise their prices. A higher inspection cost means that many of the consumers the seller encounters are unsatisfied with other options and so the seller can afford to set a high price, while a rival also usually has an incentive to raise its price since it knows consumers have less desirable alternative options. Thus, we expect there will be positive own- and cross-cost passthrough of inspection costs, as well as the more familiar positive passthrough of production costs. Because of this, a seller may have an incentive to raise its own inspection cost artificially as a way to relax competition. Although an industry-wide increase in search frictions will typically raise equilibrium prices when those prices are hidden, the opposite is likely to be the case when prices are advertised: higher search costs make a consumer more likely to buy at the first seller they inspect, and this intensifies competition to become the seller consumers wish to investigate first.

Several questions remain for further study. Pandora’s problem is a special case of a wider class of bandit problems which can be solved by means of the so-called Gittins index. (See Gittins, Glazebrook, and Weber (2011) for an account of this theory.) It would be interesting to investigate which other problems in this class, if any, can be reformulated as static discrete choice problems.

Turning to consumer markets, a worthwhile extension would be to situations where a seller has some freedom to choose the distribution for its match utility—i.e., its product design—as well as its price. For example, a seller might be able to choose between a “niche” or a “mass market” design, where the former is associated with a riskier distribution for match utility. Bar-Isaac, Caruana, and Cunat (2012) discuss this issue in the context of monopolistic competition and random consumer search, focussing on the impact of lower search costs on the choice of product design. It would be interesting to investigate this in the context of ordered search. Does a prominent seller have different incentives to choose its design from those sellers further back? While Pandora’s rule suggests that consumers would like to inspect niche (i.e., riskier) products first, it is less clear that a prominent seller has an incentive to offer a niche product design.

It would be interesting to study how knowledge of which product is the current “bestseller”

can be used to guide optimal search. As with the classical herding models of Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), suppose that consumers buy in sequence. Unlike these herding models, though, consumers have no private signal about which product is better, but merely observe which product has sold the most to date. Here, it seems plausible that, given that previous consumers have started their search process with the then-current bestseller, a consumer should optimally start her search with the current bestseller.⁴⁰ In this scenario, it seems possible that the first consumer buys an inferior product and all subsequent consumers follow suit, with the result that an inferior popular product becomes the best-seller in the long run. This would not be possible in an alternative situation where all consumers searched randomly.

The discussion of optimal pricing and store layout in section 4.1 introduces what seems to be a rich seam for further investigation, namely, how a multiproduct retailer should choose its prices and its “product placement” to maximize profit. Related issues arise also in a number of public policy discussions, and this framework might be a fruitful way to add to those discussions. For instance, is the optimal way to discourage consumption of unhealthy products to tax them, to place them away from eye level, or to use a mixture of both policies? In addition, it would be interesting to investigate the *optimal* way to sell within some natural class of contracts and layout plans. If feasible, for instance, the multiproduct seller in section 4.1 could *charge* the consumer the cost s in (9) to inspect the product rather than make the consumer pay this cost as a search friction from which the seller obtains no direct gain.

This paper focussed on situations in which consumers could freely choose their order of search. In other environments the search order is exogenously imposed. A driver on a motorway looking for fuel encounters service stations in order, and consumers have to decide now whether to buy the current model of phone or wait to see if a better or cheaper model is released next year. When consumers differ in their per-seller search cost, a likely outcome is that equilibrium prices *fall* as consumers move through the exogenous search order, in contrast to the situations discussed in section 3. Consumers with high search costs buy quickly, leaving later sellers to face a pool of consumers who are more inclined to shop around. For instance, currency exchange in an airport’s arrivals hall might be more expensive than outside the airport, to exploit those travellers reluctant to search for

⁴⁰See Sorensen (2007) and Cabral and Natividad (2016) for evidence that becoming a best-seller boosts subsequent demand.

a better deal.⁴¹

Relatedly, this paper has focussed on situations in which consumers are rational, and optimally choose the order in which they consider options. However, consumers sometimes search through options in the order they are presented, even when there is no obvious information content or differential inspection costs to being placed in certain positions. For instance, random position on the ballot paper can affect vote share in elections. Ho and Imai (2008) estimate that first-listed candidates in primary or non-partisan elections for US state or federal offices gain about two percentage points. They suggest voter behaviour can be modelled as a search problem, where voters work their way down the list of candidates on the ballot until they find one which meets the required quality threshold. Coauthors of an article are often ordered alphabetically and papers in the bibliography are listed in order of first author's name. This may mean that articles with one author early in the alphabet garner more citations and that scholars early in the alphabet are asked to act as referee more frequently.⁴² Unlike people, firms can choose their name to affect alphabetic ordering. McDevitt (2014) describes how 21% of Chicago plumbing firms have listed names starting with 'A' or a number, and also that these firms attract a disproportionate number of consumer complaints. He interprets this as being consistent with a market with expert and uninformed consumers, where the latter are assumed to search for a supplier in alphabetic order (even though this is not the optimal way to search given that low-quality plumbers apparently choose to be listed first).

While preparing this paper, my search for useful articles was far from random. I was guided by keyword searches combined with citation counts from search engines and by the bibliographies in papers I had read previously, I investigated authors who worked in the wider area and whose work I already admired, I inspected journals I expected to contain relevant papers, I solicited recommendations from colleagues, and was biased towards sources which were more easily inspected (journal articles and working papers rather than book chapters, say, or "economics" rather than "marketing" papers). Doubtless I have missed interesting and relevant work. One disadvantage of ordered search, relative to the ancient process of browsing more randomly through library shelves, is that serendipitous discovery becomes less likely.⁴³

⁴¹For further discussion see Arbatskaya (2007) and Armstrong, Vickers and Zhou (2009, section 4).

⁴²See Huang (2015) and Richardson (2008), respectively.

⁴³See Foster and Ford (2003) for further discussion.

Technical Appendix

A. Proof of the optimality of Pandora's rule:

For a given search procedure, let \mathbb{A}_i be the indicator function for the consumer selecting box i (so $\mathbb{A}_i = 1$ if the consumer selects this box and otherwise $\mathbb{A}_i = 0$), and let \mathbb{I}_i be the indicator function for inspecting box i . Since the consumer can only select a single box, we have $\sum_i \mathbb{A}_i \leq 1$, and since she must inspect a box if she selects it we have $\mathbb{A}_i \leq \mathbb{I}_i$. The consumer's expected payoff from using the search procedure is

$$\begin{aligned} \mathbb{E}(\sum_i \mathbb{A}_i v_i - \sum_i \mathbb{I}_i s_i) &= \mathbb{E}(\sum_i \mathbb{A}_i v_i - \sum_i \mathbb{I}_i \max\{v_i - r_i, 0\}) \\ &\leq \mathbb{E}(\sum_i \mathbb{A}_i [v_i - \max\{v_i - r_i, 0\}]) \\ &= \mathbb{E}(\sum_i \mathbb{A}_i w_i) \\ &\leq \mathbb{E}(\max\{w_i\}) . \end{aligned}$$

Here, the first equality holds using the definition of r_i together with the fact that \mathbb{I}_i and v_i are independent variables, while the first inequality holds since $\mathbb{A}_i \leq \mathbb{I}_i$. However, the search procedure determined by Pandora's rule has equality in these two inequalities. For the first, note that if the consumer inspects but does not select box i , i.e., if $\mathbb{A}_i < \mathbb{I}_i$, then it must be that $v_i \leq r_i$, and for the second we have already shown that the procedure selects the box with the highest w_i . We deduce that Pandora's rule generates the highest expected payoff for the consumer.

B. Duopoly example where r and v are each uniformly distributed on $[0,1]$:

Without loss of generality, look for an equilibrium in which equilibrium prices are $p_1 \leq p_2$ and let \tilde{p}_i denote seller $i = 1, 2$'s actual price. In the region $0 \leq \tilde{p}_1 \leq p_1 \leq p_2 \leq 1$, one can calculate that seller 1's demand as a function of \tilde{p}_1 , which from (5) is the fraction of consumers for whom

$$\min\{r_1 - p_1, v_1 - \tilde{p}_1\} \geq \max\{\min\{r_2 - p_2, v_2 - p_2\}, 0\} ,$$

is

$$\begin{aligned} &(1 - p_1)(1 - \tilde{p}_1)(1 - (1 - p_2)^2) + \int_{p_2}^1 (2(1 - w)(1 - p_1 - w + p_2)(1 - \tilde{p}_1 - w + p_2))dw \\ &= \frac{4}{3}p_2 - \frac{2}{3}p_1 - \frac{2}{3}\tilde{p}_1 + \tilde{p}_1 p_2^2 - \frac{1}{3}\tilde{p}_1 p_2^3 + p_1 p_2^2 - \frac{1}{3}p_1 p_2^3 - p_2^2 + \frac{1}{6}p_2^4 + \tilde{p}_1 p_1 - \tilde{p}_1 p_2 - p_1 p_2 + \frac{1}{2} , \end{aligned}$$

which is linear in \tilde{p}_1 . Since seller 1 chooses \tilde{p}_1 to maximize \tilde{p}_1 times this demand, and this optimal \tilde{p}_1 must equal p_1 , we obtain the first-order condition for p_1 given p_2 given by

$$2p_1^2 - p_1p_2^3 + 3p_1p_2^2 - 3p_1p_2 - 2p_1 + \frac{1}{6}p_2^4 - p_2^2 + \frac{4}{3}p_2 + \frac{1}{2} = 0 . \quad (12)$$

Similarly, in the region $0 \leq p_1 \leq p_2 \leq \tilde{p}_2 \leq 1$, one can calculate that seller 2's demand is

$$(1 - p_2)(1 - \tilde{p}_2)(1 - (1 - p_1)^2) + \int_{p_1}^{1 - \tilde{p}_2 + p_1} (2(1 - w)(1 - p_2 - w + p_1)(1 - \tilde{p}_2 - w + p_1))dw ,$$

and so the first-order condition for p_2 given p_1 is

$$3p_1^2p_2 - 2p_1^2p_2^2 - p_1^2 + \frac{5}{3}p_1p_2^3 - 3p_1p_2 + \frac{4}{3}p_1 - \frac{1}{2}p_2^4 + 2p_2^2 - 2p_2 + \frac{1}{2} = 0 . \quad (13)$$

A candidate equilibrium consists of a solution to the pair of equations (12)–(13).

A symmetric equilibrium, where $p_1 = p_2 = p$ say, must satisfy the equation

$$(1 - p)(3 - p - 13p^2 + 5p^3) = 0 ,$$

which has a unique root in the interval $(0, 1)$ equal approximately to $p \approx 0.49$. One can check that when $p_1 = p_2$ is equal to this root, it is indeed the optimal strategy for seller 1 to make the same choice $\tilde{p}_1 = p$, and so this constitutes an equilibrium.⁴⁴ One can manipulate (12)–(13) to obtain p_1 as an explicit function of p_2 , and then substitute this p_1 back into one of (12)–(13). Doing so reveals there exists no asymmetric equilibrium with $0 < p_1 < p_2 < 1$.

C. Details of argument that more prominent sellers face more elastic demand:

Suppose all consumers search through the n sellers in the same order. Suppose hypothetically that sellers are expected to charge the same price p (where $p < r$ so that consumers search at all), and consider each seller's elasticity of demand with respect to a small change in its actual price. By Pandora's rule, when consumers expect the same price $p < r$ from all sellers they will buy from the first seller which offers payoff $v - \tilde{p}$ above $r - p$ (where \tilde{p} is a seller's actual price), and if no payoff meets this threshold they will buy from the seller with the highest payoff provided this is positive. Using the terminology in Armstrong, Vickers, and Zhou (2009), demand from consumers who buy immediately (i.e., if $v - \tilde{p} \geq r - p$) is a seller's "fresh demand", while demand from those consumers who

⁴⁴This procedure only allows for seller 1 to choose a lower price $\tilde{p}_1 < p$, but one can check that an upward deviation in its price also reduces its profit.

buy from the seller only after exhausting all options is its “return demand”.⁴⁵ The seller which is m^{th} in the search order has fresh demand in terms of its actual price \tilde{p} equal to

$$q_F(\tilde{p}) = F^{m-1}(r)[1 - F(r + \tilde{p} - p)] . \quad (14)$$

(This is because a consumer only reaches it if she did not find a match v above r from the previous $m - 1$ sellers, and the consumer will then buy immediately if $v - \tilde{p} \geq r - p$.) With n sellers in all, this seller’s return demand is

$$q_R(\tilde{p}) = \int_{\tilde{p}}^{r+\tilde{p}-p} F^{n-1}(v + p - \tilde{p})f(v)dv = \int_p^r F^{n-1}(\tilde{v})f(\tilde{v} + \tilde{p} - p)d\tilde{v} \quad (15)$$

where the first equality follows since a firm sells to a return consumer if $v - \tilde{p}$ is below the threshold $r - p$ (for otherwise she would have purchased immediately) and above all the other firms’ offers and the outside option of zero, and the second follows after changing variables from v to $\tilde{v} = v + p - \tilde{p}$. Thus, fresh demand is proportional to $1 - F(r + \tilde{p} - p)$, scaled down geometrically by $F(r)$ as the seller is placed further back in the search order, while return demand does not depend on the seller’s position in the search order. When $f(v)/(1 - F(v))$ increases with v , any seller’s fresh demand is more elastic than its return demand (evaluated at price $\tilde{p} = p$). To see this, note that

$$\begin{aligned} -q'_R(p) &= - \int_p^r F^{n-1}(\tilde{v})f'(\tilde{v})d\tilde{v} \leq \int_p^r F^{n-1}(\tilde{v})f(\tilde{v})\frac{f'(\tilde{v})}{1 - F(\tilde{v})}d\tilde{v} \\ &\leq \frac{f(r)}{1 - F(r)} \int_p^r F^{n-1}(\tilde{v})f(\tilde{v})d\tilde{v} = -\frac{q'_F(p)}{q_F(p)}q_R(p) , \end{aligned}$$

which establishes the claim.⁴⁶ Since a more prominent seller (i.e., a seller with smaller m in (14)) has a greater volume of fresh demand relative to its return demand, it follows that this seller’s total demand is more elastic than that of its rivals further back.

D. Duopoly example where r is constant and v is uniformly distributed on $[0,1]$:

Look for an equilibrium in which equilibrium prices satisfy $p_1 < p_2 < r$, where $r = 1 - \sqrt{2s}$, so that all consumers inspect seller 1 first and if they do not buy immediately

⁴⁵For the final seller in the search order, all of whose demand is really immediate, for consistency divide its demand into an “immediate” portion with $v - \tilde{p} \geq r - p$ and a “return” portion with $v - \tilde{p} \leq r - p$.

⁴⁶Here, the first inequality follows from the assumption that $f/(1 - F)$ is increasing, the second inequality also follows from $f/(1 - F)$ being increasing and the fact that $\tilde{v} \leq r$, while the final equality follows from the definitions of q_F and q_R .

from seller 1 they go on to inspect seller 2. From Figure 3, one can calculate that seller 1's demand as a function of its actual price \tilde{p}_1 given the equilibrium price p_2 from its rival is

$$\underbrace{\frac{1}{2}(r^2 - p_2^2)}_{\text{return demand}} + \underbrace{1 - (r - p_2 + \tilde{p}_1)}_{\text{immediate demand}}$$

which is linear in its price \tilde{p}_1 . With costless production, seller 1 chooses \tilde{p}_1 to maximize \tilde{p}_1 times this demand, and this optimal \tilde{p}_1 must equal p_1 , and so we obtain the first-order condition for p_1 given p_2 given by

$$\frac{1}{2}(r^2 - p_2^2) + 1 + p_2 - r = 2p_1 . \quad (16)$$

Likewise, in the same region $p_1 < p_2 < r$ seller 2's demand in terms of its actual price \tilde{p}_2 and equilibrium prices (p_1, p_2) is

$$(1 - \tilde{p}_2)(r + p_1 - p_2) - \frac{1}{2}(r - p_2)^2$$

which again is linear in price \tilde{p}_2 . Since seller 2 chooses \tilde{p}_2 to maximize \tilde{p}_2 times this demand, and this optimal \tilde{p}_2 must equal p_2 , the first-order condition for p_2 given p_1 is

$$(1 - 2p_2)(r + p_1 - p_2) = \frac{1}{2}(r - p_2)^2 . \quad (17)$$

The equilibrium prices can then be explicitly (but messily) solved for each reservation price $\frac{1}{2} \leq r \leq 1$ from the pair of conditions (16)-(17), and hence in terms of the underlying search cost s , and are depicted in the left-hand graph in Figure 4 in the text. (When $r < \frac{1}{2}$, the search cost is so high that there is no equilibrium where consumers participate in the market.) These prices increase with the search cost s . The prices coincide when there are no search frictions ($s = 0$) or when the search cost is so high that a consumer is just willing to inspect a seller which sets the monopoly price $p^M = \frac{1}{2}$ (when $s = \frac{1}{8}$). Otherwise, though, an equilibrium exists where all consumers first inspect seller 1 and that seller sets a strictly lower price. (Of course, a similar equilibrium also exists where all consumers first inspect seller 2.) The equilibrium sales of the two sellers is shown in the right-hand graph in Figure 4, where the prominent seller has the greater sales. Here, the prominent seller's equilibrium demand increases with s , while the non-prominent seller sells less when search frictions are greater.

When does an equilibrium exist in which sellers set the same price, say p , and consumers are indifferent about their order of search? Suppose a fraction σ of consumers first inspect

seller 1 when anticipated prices are equal. By considering Figure 3 above, one can check that seller 1’s demand when it chooses actual price \tilde{p} is

$$\underbrace{\sigma \left(\frac{1}{2}(r^2 - p^2) + 1 - (r - p + \tilde{p}) \right)}_{\text{first inspect seller 1}} + \underbrace{(1 - \sigma) \left((1 - \tilde{p})r - \frac{1}{2}(r - p)^2 \right)}_{\text{first inspect seller 2}},$$

and the first-order condition for this seller to choose $\tilde{p} = p$ is

$$\sigma \left(\frac{1}{2}(r^2 - p^2) + 1 - (r + p) \right) + (1 - \sigma) \left((1 - 2p)r - \frac{1}{2}(r - p)^2 \right) = 0 .$$

The corresponding expression for seller 2 to be willing to choose $\tilde{p} = p$ is the same except σ and $1 - \sigma$ are permuted. As such, the only situation in which both sellers are willing to choose the same price is when $\sigma = \frac{1}{2}$, in which case the equilibrium price satisfies $(1+r)p+p^2 = 1$. Thus, if the “tie-breaking” rule is such that more than half the consumers first inspect one seller when anticipated prices are equal, the only equilibria in this market involves ordered search where one seller sets a strictly lower price than its rival.

References

- ANDERSON, S., AND R. RENAULT (1999): “Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model,” *Rand Journal of Economics*, 30(4), 719–735.
- (2015): “Search Direction,” mimeo.
- ARBATSKAYA, M. (2007): “Ordered Search,” *Rand Journal of Economics*, 38(1), 119–126.
- ARMSTRONG, M., AND J. VICKERS (2015): “Which Demand Systems can be Generated by Discrete Choice?,” *Journal of Economic Theory*, 158, 293–307.
- ARMSTRONG, M., J. VICKERS, AND J. ZHOU (2009): “Prominence and Consumer Search,” *Rand Journal of Economics*, 40(2), 209–233.
- ARMSTRONG, M., AND J. ZHOU (2011): “Paying for Prominence,” *Economic Journal*, 121(556), F368–395.
- (2016): “Search Deterrence,” *Review of Economic Studies*, 83(1), 26–57.
- ATHEY, S., AND G. ELLISON (2011): “Position Auctions with Consumer Search,” *Quarterly Journal of Economics*, 126(3), 1213–1270.

- BAGWELL, K., AND G. RAMEY (1994): “Coordination Economies, Advertising, and Search Behavior in Retail Markets,” *American Economic Review*, 84(3), 498–517.
- BANERJEE, A. (1992): “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107(3), 797–817.
- BAR-ISAAC, H., G. CARUANA, AND V. CUNAT (2012): “Search, Design, and Market Structure,” *American Economic Review*, 102(2), 1140–1160.
- BAYE, M., R. GATTI, P. KATTUMAN, AND J. MORGAN (2009): “Clicks, Discontinuities, and Firm Demand Online,” *Journal of Economics and Management Strategy*, 18(4), 935–975.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 100, 992–1026.
- BURDETT, K., S. SHI, AND R. WRIGHT (2001): “Pricing and Matching with Frictions,” *Journal of Political Economy*, 109(5), 1060–1086.
- CABRAL, L., AND G. NATIVIDAD (2016): “Box-Office Demand: The Importance of Being No. 1,” *Journal of Industrial Economics*, 64(2), 277–294.
- CHADE, H., AND L. SMITH (2006): “Simultaneous Search,” *Econometrica*, 74(5), 1293–1307.
- CHEN, Y., AND C. HE (2011): “Paid Placement: Advertising and Search on the Internet,” *Economic Journal*, 121(556), F309–F328.
- CHOI, M., Y. DAI, AND K. KIM (2016): “Consumer Search and Price Competition,” mimeo, University of Iowa.
- DE LOS SANTOS, B., A. HORTACSU, AND M. WILDENBEEST (2012): “Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior,” *American Economic Review*, 102(6), 2955–2980.
- DENECKERE, R., AND P. MCAFEE (1996): “Damaged Goods,” *Journal of Economics and Management Science*, 5(2), 149–174.

- DIAMOND, P. (1971): “A Model of Price Adjustment,” *Journal of Economic Theory*, 3(2), 156–168.
- DOVAL, L. (2014): “Whether or Not to Open Pandora’s Box,” mimeo.
- ELIAZ, K., AND R. SPIEGLER (2011): “A Simple Model of Search Engine Pricing,” *Economic Journal*, 121(556), F329–F339.
- ELLISON, G., AND A. WOLITZKY (2012): “A Search Cost Model of Obfuscation,” *Rand Journal of Economics*, 43(3), 417–441.
- FISCHER, J., AND J. HARRINGTON (1996): “Product Variety and Firm Agglomeration,” *Rand Journal of Economics*, 27(2), 281–309.
- FOSTER, A., AND N. FORD (2003): “Serendipity and Information Seeking: An Empirical Study,” *Journal of Documentation*, 59(3), 321–340.
- GAMP, T. (2016): “Guided Search,” mimeo, University of Bonn.
- GARCIA, D., AND S. SHELEGIA (2015): “Consumer Search with Observational Learning,” mimeo.
- GITTINS, J., K. GLAZEBROOK, AND R. WEBER (2011): *Multi-Armed Bandit Allocation Indices*. John Wiley and Sons, UK.
- GOMES, R. (2014): “Optimal Auction Design in Two-Sided Markets,” *Rand Journal of Economics*, 45(2), 248–272.
- HAAN, M., AND J.-L. MORAGA-GONZALEZ (2011): “Advertising for Attention in a Consumer Search Model,” *Economic Journal*, 121, 552–579.
- HAAN, M., J.-L. MORAGA-GONZALEZ, AND V. PETRIKAITE (2015): “Price and Match-Value Advertising with Directed Consumer Search,” mimeo.
- HO, D., AND K. IMAI (2008): “Estimating Causal Effects of Ballot Order from a Randomized Natural Experiment: The California Alphabet Lottery, 1978-2002,” *Public Opinion Quarterly*, 72(2), 216–240.
- HUANG, W. (2015): “Do ABCs get more Citations than XYZs?,” *Economic Inquiry*, 53(1), 773–789.

- JEZIORSKI, P., AND I. SEGAL (2015): “What Makes Them Click: Empirical Analysis of Consumer Demand for Search Advertising,” *American Economic Journal: Microeconomics*, 7(3), 24–53.
- KLEINBERG, R., B. WAGGONER, AND G. WEYL (2016): “Descending Price Coordinates Approximately Efficient Search,” mimeo.
- MAILATH, G., L. SAMUELSON, AND A. SHAKED (2000): “Endogenous Inequality in Integrated Labor Markets with Two-sided Search,” *American Economic Review*, 90(1), 46–72.
- MCDEVITT, R. (2014): “‘A’ Business by Any Other Name: Firm Name Choice as a Signal of Firm Quality,” *Journal of Political Economy*, 122(4), 909–944.
- MONTGOMERY, J. (1991): “Equilibrium Wage Dispersion and Interindustry Wage Differentials,” *Quarterly Journal of Economics*, 106(1), 163–179.
- MORAGA-GONZALEZ, J., AND V. PETRIKAITE (2013): “Search Costs, Demand-Side Economies and the Incentives to Merge Under Bertrand Competition,” *Rand Journal of Economics*, 44(3).
- OLSZEWSKI, W., AND R. WEBER (2015): “A More General Pandora Rule?,” *Journal of Economic Theory*, 160, 429–437.
- PETRIKAITE, V. (2016): “Consumer Obfuscation by the Multiproduct Firm,” mimeo.
- RICHARDSON, M. (2008): “Alphabetic Bias in the Selection of Reviewers for the American Journal of Roentgenology,” *American Journal of Roentgenology*, 191, W213–W216.
- SALOP, S. (1973): “Systematic Job Search and Unemployment,” *Review of Economic Studies*, 40(2), 191–201.
- (1977): “The Noisy Monopolist: Imperfect Information, Price Dispersion and Price Discrimination,” *Review of Economic Studies*, 44(3), 393–406.
- SHEN, J. (2014): “Ex Ante Preference in a Consumer Search Market,” mimeo.
- SIEKMAN, W. (2016): “Directed Consumer Search,” mimeo, University of Groningen.

- SORENSEN, A. (2007): “Bestseller Lists and Product Variety,” *Journal of Industrial Economics*, 55(4), 715–738.
- URSU, R. (2015): “The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions,” mimeo.
- VARIAN, H. (1980): “A Model of Sales,” *American Economic Review*, 70(4), 651–659.
- WEBER, R. (1992): “On the Gittens Index for Multiarmed Bandits,” *Annals of Applied Probability*, 2(4), 1024–1033.
- WEITZMAN, M. (1979): “Optimal Search for the Best Alternative,” *Econometrica*, 47(3), 641–654.
- WILSON, C. (2010): “Ordered Search and Equilibrium Obfuscation,” *International Journal of Industrial Organization*, 28(5), 496–506.
- WOLINSKY, A. (1986): “True Monopolistic Competition as a Result of Imperfect Information,” *Quarterly Journal of Economics*, 101(3), 493–511.
- ZHOU, J. (2009): “Prominence in Search Markets: Competitive Pricing and Central Pricing,” mimeo.
- (2011): “Ordered Search in Differentiated Markets,” *International Journal of Industrial Organization*, 29(2), 253–262.