

Firm pricing with consumer search*

Simon P. ANDERSON[†] and Régis RENAULT[‡]

December 21, 2016

Prepared for the Handbook of Game Theory and Industrial Organization, edited by Luis Corchon and Marco Marini, Edward Elgar Publishing, 2017.

ABSTRACT Economists could not properly capture the impact of internet on markets without a proper theory of consumer search. As a result, this theory has been rediscovered and developed further since the early 2000s. It can address such critical questions as the impact of reduced search cost on prices, variety, and product choice as well as advertising practices (such as search advertising). This theoretical development has also fed into a rich empirical literature exploiting the wealth of data that is now available regarding both consumers' and firms' online activity. The goal of this chapter is to present the basic concepts underpinning the theory of imperfectly competitive markets with consumer search. We stress that appropriate theoretical frameworks should involve sufficient heterogeneity among agents on both sides of the market. We also explain why the analysis of ordered search constitutes an essential ingredient for modeling recent search environments.

Keywords: Reservation rule, Diamond paradox, sequential/simultaneous search, random/ordered search, price dispersion, product differentiation, internet, two-sided market.

JEL Classification: D42 D82 D83 L11 L15 M37 M38

*Acknowledgements: We thank Mark Armstrong, José Luis Moraga-González, Andrew Rhodes, Vaiva Petrikaitė and Jidong Zhou for invaluable comments. We also thank Emily Cook for admirable research assistance. The first author thanks the NSF for financial support; the second author thanks Labex MME-DII for financial support.

[†]CEPR and Department of Economics, University of Virginia, sa9w@virginia.edu

[‡]Université de Cergy-Pontoise, Thema. regis.renault@u-cergy.fr

1 Introduction

Consumer search in the modern economy seems more prevalent than ever before, with the advent of internet shopping opportunities. The lower per search cost enabled through the internet, concurrent with low transactions costs for shipping final purchases, has enabled consumers to access a huge variety of options. Heretofore the effective choice set was limited by the local market, and even shopping in that market was constrained by quite high costs of getting informed as to what products were offered at what prices. Search costs have always been important, but, since few people actually actively searched much, it was perhaps not apparent that search costs have a major constraining role in stifling economic activity. Ironically, because search costs were less visible, economists often ignored them and addressed seemingly more pressing questions about pricing, product variety, and market power, while treating the set of available products as effectively given and known.

Now that search is visibly at the fore of the modern economy, economists are developing (or re-developing) models of market interaction in which search costs play a central role. Search costs are a friction in markets, and their reduction facilitates transactions that are likely to involve both better matching (of consumer to product) and lower prices (because there is more effective competition). One key reason why there is now so much more accessible choice is that in earlier times high search costs effectively deterred access to available options: search costs silently curtailed the set of goods and services brought to market. Nonetheless, significant search frictions remain.

To set the stage from the perspective of economic theory of oligopoly pricing, we start by describing how costly search impacts economic outcomes in the simplest scenario, which we can then embellish to address more subtle features. The first result is both striking and quite extreme.

Assume that firms set prices for a homogeneous good, and produce at a common constant marginal cost. Absent search costs, this market interaction gives us the Bertrand Paradox (itself a striking and extreme result): *price is at marginal cost* as long as at least two firms compete. Now introduce a small search cost, s , for the consumer to find out the price set by a firm, but assume she observes a first price quote for free. Suppose that consumers search sequentially among options, and close the model with the rational expectations stipulation that they are correct in equilibrium about the price(s) they expect. To make it simple, suppose that the consumers all have the same unit demand curve, with maximal valuation r for the product. Then, the outcome – the so-called “Diamond Paradox” after its first proponent (Diamond, 1971) – has all firms pricing at the *monopoly price*, r . Roughly, one can think of prices being “ratcheted up” (by the amount of the search cost) for any putative lowest price below the monopoly level, r . (In the text, we develop and extend the arguments and contexts that give rise to the Paradox.) Thus even a small search cost radically tips the Bertrand Paradox to the opposite extreme.¹ There is a second leg to the Paradox: the

¹Some classic ways of softening the Bertrand Paradox (see e.g. Tirole, 1988) include introduc-

context is a search model, and yet no one does actually search again in equilibrium. This is logical because prices are rationally expected to be the same for all firms. And it is because there is no search that the monopoly price sustains.²

To be sure, economists were thinking about search before Diamond’s market equilibrium solution. Stigler (1961) was intrigued by the price dispersion (and the implicit repealing of the Law of One Price) that he saw in markets from anthracite coal to Chevrolets (see also Lach, 2002, for a study documenting price dispersion in Israeli supermarkets for instant coffee, frozen chickens, flour, and refrigerators; Ellison and Fisher Ellison, 2014, for on-line books, etc.). This led Stigler to begin formulating consumer search theory (we develop this theme below by using the formal model of Burdett and Judd, 1983, as our jumping-off point for the formal analysis with equilibrium pricing.) The theory of consumer search—with the consumer facing a set of options with exogenous utility distributions—was developed quite briskly. It culminated in the beautiful characterization of the consumer problem in Weitzman (1979) whereby consumers search remaining options in order of their reservation utility scores until they reach a stopping point at which remaining options have scores below their current holding. This rule, and some of the work that led up to it, is discussed in more detail below.

Somewhat ironically though, even though Stigler felt that consumer search frictions should be at the heart of observed price dispersion, the Diamond result suggested otherwise – at least under the assumption that consumers search sequentially across options (we discuss simultaneous search with reference to Stigler in the next section). There are of course other reasons for price dispersion, such as heterogeneity in consumer tastes across products, and cost differences across firms, but search costs are not one of them – at least in the Diamond set-up. This observation leads us to consider what could generate price dispersion and equilibrium search when search is indeed sequential (contrast Stigler on this point), and products are homogeneous to consumers (i.e., absent product differentiation). As we argue below, pure production cost heterogeneity is insufficient (Reinganum, 1979): while this generates equilibrium price dispersion (with a pooling atom at the price that just deters further search) it generates no search. Demand heterogeneity alone is also unlikely to give dispersion, and in many instances pure search cost heterogeneity alone does not

ing product differentiation, repeated interaction, and decreasing returns to scale. We will discuss below how product differentiation relaxes the Diamond Paradox. The relevant repeated interaction in the search context would be long term relationships between buyers and sellers with repeat purchases, where prices would be kept low in order to ensure ongoing business in the future. Allowing for decreasing returns (increasing marginal costs) with consumers still searching randomly would imply that “Diamond” equilibrium prices decrease with the number of firms in the market, as the “monopoly” market is spread in more pieces. Conversely, locally increasing returns (not so severe as to violate second order conditions) cause prices to *rise* with entry.

²This argument, as we shall see below, is why product differentiation breaks the Paradox: people search for better product matches, and thus firms are brought into competition with each other.

either (Rob, 1985, Stahl, 1996). However, there is a notable exception, which arises when some consumers have zero search costs. Indeed, Stahl (1989) has provided an enduring work-horse model of pricing under consumer search that generates price dispersion. In this model, products are homogeneous but a fraction of consumers have strictly positive search costs while the rest are assumed to observe all prices (and they can be construed as a mass of consumers with zero search costs).³ The equilibrium is in mixed strategies, and constructing it (via an expository device suggested by Janssen, Moraga-González, and Wildenbeest, 2005) builds off the ingredient models of Varian’s (1980) model of sales coupled to the constraint that the consumer reservation price for the high-cost types just deters them from searching (engaging one striking consequence of Weitzman’s search rule – that if a myopic search is not optimal, then it is optimal not to search at all).⁴

Allowing for differentiated products is an appealing way to break the Diamond Paradox. Wolinsky (1986) first explored this avenue, and Anderson and Renault (1999) develop further the economics of the approach. The idea is that search is not just for price but also for product suitability with a consumer’s tastes, and different products suit better different consumers (“horizontal differentiation”). In this case a consumer may well consider searching even if she encounters the same price she expects elsewhere (or even if she encounters a lower one), because she may seek a better match to her desired product specification. The situation is most tractably modeled as a differentiated product discrete choice model with i.i.d. valuations (idiosyncratic taste matches) across consumers. Under symmetric demands, there is a symmetric equilibrium price, so this approach does not per se deliver price dispersion.⁵ It does nonetheless deliver consumer search in equilibrium: different consumers search different numbers of options, even if their search costs are the same (because some are luckier in finding satisfactory matches early). Moreover, higher search costs induce less search, which in turn by increasing friction entails less effective competition and hence higher equilibrium prices. Another result ties together nicely the standard product differentiation literature with the Diamond Paradox reasoning. In the standard product differentiation setting (with no search frictions), a higher variance of consumer tastes for products (i.e., more heterogeneity of products) leads to higher equilibrium prices as individuals have more pronounced idiosyncratic tastes and so individual product demands are more inelastic. The presence of search costs can cause equilibrium prices to initially decrease as the degree of product differentiation rises. The key here is that more product differentiation leads to more search, which leads to more competition as more firms are

³Because the high search cost types never search again in equilibrium, one might argue that there is no real search in this model.

⁴Varian’s (1980) model of sales is described in the next section. Weitzman’s key results are in section 4.4.

⁵Different prices can be generated by allowing firms to have different production costs or vertical qualities, modulo a significantly higher consequent complication of the analysis. Another important channel for generating price dispersion is to have ordered consumer search, as discussed at length below.

effectively brought into the contest for a consumer. At first (i.e., for low product heterogeneity), this causes prices to fall down from the high Diamond levels (where there is no search at all). For higher heterogeneity consumers are searching sufficiently that they know about a significant fraction of firms. Then this is like the model without search costs, and more differentiation begets higher prices through greater product loyalty.

In many contexts—from geographical markets through online shopping—search is both sequential and systematically ordered. Ordering gives prominence to firms searched early on, and can quite drastically influence equilibrium pricing. Because earlier firms get more traffic they are usually more profitable, so firms are willing to expend resources (bid for prime locations or high positions in search engine auctions) for better positions. Equilibrium orders depend on differential pay-offs at different positions. When firms are intrinsically heterogeneous and asymmetric (e.g. when they command different distributions of consumer valuations), the equilibrium order is driven by profitability differences. However, profit differences do not necessarily reflect consumer welfare differences. Thus the order of presentation of options resulting from an auction might be unattractive to buyers who might consequently be put off from participating on a platform (Anderson and Renault, 2016).

We devote the last subsection within each section to ordered search, so as to discuss its impact in the various contexts of product heterogeneity, seller heterogeneity, and buyer heterogeneity. Contemporaneous work by Armstrong (2016) both surveys and extends the pricing theory in this important direction for the case of product heterogeneity, and we refer the reader to Armstrong’s excellent paper for further elaboration.⁶

One organizing theme we stress throughout is the impact of reducing search costs, as predicted by various approaches to modeling pricing with search costs. Our motivation, as per the start of this Introduction, is the internet experience and its impact on markets. We also take a broader perspective by engaging the endogeneity of the consumer (and firm) participation decision, recognizing the two-sided nature of many search markets. The thickness of the market may be greater with lower search costs, and must be accounted for to fully appreciate the benefits of search cost reductions.

2 Price dispersion

The central argument in Stigler (1961) is that price dispersion may persist because consumers are not aware of all prices and therefore cannot properly arbitrage price differences. Here we present various settings where costly search by buyers generates price dispersion resulting from the price mixed strategies used by firms. We focus more specifically on settings where such price dispersion arises in spite of a homogenous population of sellers

⁶One particularly stimulating contribution of Armstrong’s paper is the reformulation of sequential search as a static discrete choice problem without search frictions.

and a homogenous population of buyers. Baye, Morgan, and Scholten (2006) provide a fine survey of models of price dispersion with a homogeneous product, and emphasize the role of a clearinghouse that publishes prices. To fix ideas about how mixed strategies arise when consumers are imperfectly informed about prices we start, however, with a simple setting that allows for some heterogeneity on the demand side.

2.1 Imperfect buyer information yields price dispersion

Consider a market comprised of $n > 1$ firms with zero production costs selling a homogeneous product to a continuum of consumers with unit demand. A consumer's valuation is denoted $r > 0$. The total measure of consumers is $m > 0$, out of which there is a measure $\sigma \in (0, m)$ of shoppers who observe all prices at no cost. Shoppers therefore always buy at the lowest price in the market. The remaining consumers observe only one price for free and face a prohibitively high search cost if they wish to observe an additional price quote. These captive consumers therefore always buy from the firm whose price quote they get for free, provided that its price does not exceed r . They are equally shared among the firms, and $\gamma = \frac{m-\sigma}{n}$ denotes the measure of captive consumers per firm. Hence a firm is guaranteed a profit of γr , which would result from charging the monopoly price r and selling exclusively to captive consumers.

A standard Bertrand undercutting argument shows that there cannot be a pure strategy equilibrium where shoppers pay a price that strictly exceeds marginal cost: a firm sharing the shoppers' demand with some other firm(s) at a price above marginal cost could always profitably slightly undercut that price to capture the entire shoppers' demand. However, pricing at marginal cost cannot be an equilibrium in the current setting because a firm charging that price would earn zero profit, whereas it can always guarantee a strictly positive profit by giving up selling to shoppers altogether and extracting the entire surplus from its captive consumers.

The literature, following Varian (1980), has focused on characterizing the unique symmetric mixed strategy equilibrium of this game. As will be seen below, the characterization of this equilibrium turns out to be quite relevant to the analysis of more elaborate search frameworks. The undercutting argument mentioned above can be used again to show that the price distribution can have no atom. Furthermore, a firm charging the highest price in the support of the price distribution will sell to shoppers with probability zero, therefore it is optimal for that firm to charge the captive consumers' valuation, r . The Bertrand intuition that all profits from sales to shoppers are competed away carries over here, and a firm's equilibrium profit is equal to its fall-back profit from selling to captive consumers alone, γr . However, a firm's price cannot fall below the level at which selling to all shoppers plus its captive consumers equals the fall-back profit. Thus, the minimum equilibrium price, \underline{p} , solves $(\gamma + \sigma)\underline{p} = \gamma r$, so $\underline{p} = \frac{\gamma}{\gamma + \sigma}r$. The equilibrium price distribution is obtained by ensuring that a firm maximizes its profit at all prices within the support and hence earns the same profit at all those prices. Appendix A shows that the equilibrium price

distribution function is given by

$$(1) \quad F(p) = 1 - \left[\frac{\gamma}{\sigma} \left(\frac{r}{p} - 1 \right) \right]^{\frac{1}{n-1}}.$$

As should be expected, prices are stochastically higher (i.e. $F(p)$ is lower over the entire support, reflecting first order stochastic dominance) if either the consumer valuation, r , or the relative share of captive consumers (reflected in $\frac{\gamma}{\sigma}$) are higher.

The comparative statics with respect to the number of firms is however less intuitive, and depend upon whether the new firm brings its new captive consumers with it when entering. If it does so, the fraction of shoppers in the consumer population falls. This is the situation described in Rosenthal (1980). Alternatively, the fraction of shoppers can be held fixed, and an entrant could just garner its share of the captive population. This is the situation germane to the search context we elaborate upon later. In both cases, there are two relevant expected prices to track. The first is the expected price per firm generated from $F(p)$: this is the price paid by captive consumers. The second is the expected minimum price, which is the price paid by the shoppers. Notice that more firms means more draws from the price distribution, which per se drives a lower price. However, for both cases (whether total shoppers or total captives are constant) $F(p)$ rises, so that what happens to the expected shopper price depends on which effect dominates.

First consider the situation analyzed by Rosenthal (1980), which is equivalent to letting the population of captive consumers increase so as to keep γ unchanged as more firms enter the market.⁷ Then prices are first-order stochastically higher with more firms (to see this note that for prices in the interior of the support (\underline{p}, r) , the bracketed term is in $(0, 1)$ so $F(p)$ falls). This reflects the logic of mixed strategies which requires that, as the number of competitors increases, a firm must find it as profitable to drop its price below r to attract shoppers as to charge r to its captive consumers. Profit from the latter choice is independent of the number of firms, so there is more reason to focus on the captive (or “loyal”) ones because there is less chance to be the low-price seller when there are more rivals. Hence the expected price paid by captive consumers rises. Rosenthal (1980) shows that this effect dominates the effect of more draws, so the price to shoppers goes up too.⁸ Baye, Morgan, and Scholten (2006) note that this might be expected because the shoppers are getting smaller as a fraction of the market.

Now suppose that the shopper population remains constant so that γ is decreases in n . Janssen and Moraga-González (2004) show that the expected captive price rises (although the distributions are not first-order stochastically ordered). However, as shown by Morgan, Orzen, and Sefton (2006), the expected shopper price falls with entry,⁹ as the shoppers get

⁷Rosenthal (1980) allows for general demand; we have specialized to rectangular demand.

⁸Rosenthal (1980) also notes the supports of the equilibrium prices do not change with n .

⁹A quick confirmation of this result follows from the analysis in Appendix D. There we have

relatively more important to individual firms, and the effect of more price draws dominates. Thus the prices move in different directions. As we shall see in subsection 5.2, when the reservation price is endogenously derived from optimal search behavior, the expected shopper price is independent of n .

The price dispersion that arises because of the captive consumers' limited information should naturally feed into the incentives of these consumers to search on. This can only be captured in a setting where search costs are low enough in order for non shoppers to have a genuine choice of whether to search or not. This possibility is explored in subsection 5.2 below. In the remainder of this section, we discuss how price dispersion may arise in settings where all consumers are identical and the heterogeneity in consumer information arises endogenously.

2.2 Simultaneous search and the coordination problem

A first view of how consumers make their search decisions, which is consistent with Stigler's original analysis (Stigler, 1961), is to assume that consumers commit to getting several price quotes before deciding where to buy. For instance, each consumer could choose to request price quotes that would be sent back in due time. To analyze this situation, we modify the setting above by assuming that all consumers get a first price quote for free, and then incur a cost $s > 0$ for each additional price quote requested. The set of firms' prices observed by a consumer is selected randomly.

First remark that there is no equilibrium¹⁰ where all consumers request at least one additional price quote. Indeed, if all consumers observed at least two prices, then all firms would price at marginal cost: each firm competes *à la* Bertrand for each consumer with all other firms whose price is observed by that consumer. Expecting this though, no consumer would be willing to incur the cost of observing a second price. Hence, in equilibrium, there must be some consumers who see only one price. Importantly, no matter how small the search cost, and akin to the Diamond paradox noted in the introduction, there is always an equilibrium where all consumers see only one price and firms charge the monopoly price. Charging the monopoly price is clearly profit maximizing if consumers observe only one price, and it is optimal for consumers not to incur any search cost if they expect the same price at all firms.¹¹

that the expected shopper price is constant under entry in a situation where the reservation price rises with entry. We also know that a higher reservation price raises prices. Hence, when the reservation price is exogenously fixed, as in the Varian (1979) model, the shopper price falls with entry.

¹⁰In this section we do not dwell on the equilibrium concept although we should be noted that we always assume that consumers have passive beliefs after observing a price deviation. Somewhat more formal discussions can be found in subsections 3.2 and 4.2.

¹¹This would not be an equilibrium if the first price quote was costly because consumers would

Burdett and Judd (1983) explore the possibility of equilibria where consumers search beyond one firm and prices are dispersed. They consider symmetric equilibria where all firms choose the same probability distribution over prices. They show that there is no equilibrium where consumers observe more than two prices. Relevant equilibria are then those where consumers mix between observing only one price, where ν_1 denotes the probability of such a choice, and requesting a second price quote, which they do with complementary probability $\nu_2 = 1 - \nu_1$. The firms' pricing behavior depends on these probabilities and must be such that consumers are indifferent between the two choices in equilibrium.

To characterize a firm's profit maximizing behavior note first that the only consumers that are relevant to a firm's pricing decision are those who observe that firm's price. There are $\frac{2m\nu_2}{n}$ such consumers who also obtain a price quote from some competing firm and $\frac{m\nu_1}{n}$ consumers who only see the firm's price and are therefore in an analogous situation to captive consumers in the preceding subsection. Furthermore, consumers observing two prices are in the same position as "shoppers" in our analysis above when there are only 2 firms in the market. Now, for each consumer, the firm competes with at most one other competitor, so that it is readily seen that its profit expression given its expectation about that other firm's behavior is exactly that of subsection 2.1 with two firms, $n = 2$, with $\gamma = \frac{m\nu_1}{n}$ and $\sigma = \frac{2m\nu_2}{n}$. As shown in Appendix A the equilibrium price distribution function is therefore

$$(2) \quad F(p) = 1 - \left[\frac{\nu_1}{2\nu_2} \left(\frac{r}{p} - 1 \right) \right].$$

The value of ν_1 , and hence that of ν_2 , are obtained from the requirement that a consumer should be indifferent between observing only one price, thus paying the expected price Ep , and incurring search cost s to see a second price with a corresponding expected payment Ep_{min} , where p_{min} is the minimum of two prices. We therefore need ν_1 to solve

$$(3) \quad Ep - Ep_{min} = s.$$

Burdett and Judd (1983) show that the left hand side, which measures the benefit from search, is quasiconcave in ν_1 and zero at both ends of the $[0, 1]$ interval.¹² Intuitively, if ν_1 is close to zero so that nearly all consumers observe two prices, prices are close to marginal cost with a high probability and the expected minimum of two prices does not differ much from the expected price. Towards the other extreme, if ν_1 is close to 1 so that almost all consumers are captive, prices are very likely to be close to the monopoly value r and so there is not much benefit from seeking a second price quote. Hence, if the search cost s is

not be willing to incur the cost, expecting to be held up at the monopoly price. However, if consumer demand is price sensitive as in Burdett and Judd (1983), then this equilibrium survives as long as the search cost is not too large.

¹²Armstrong, Vickers, and Zhou (2009b) show that it is concave. These authors analyze a generalized version of the Burdett and Judd (1983) model.

not too large, there are two values of ν_1 and therefore two equilibria. One has low search intensity (ν_1 large) and high prices (although this one is unstable: see the discussion in Fershtman and Fishman, 1992); the other one has high search intensity (ν_1 low) and low prices.

An interesting property of these equilibria is that the number of sellers has no impact on the competitiveness of the market. It is quite intuitive that, because each firm expects to be competing with at most one other firm for each consumer, it behaves as if it were in a duopoly market, independent of the total number of competitors. Anderson *et al.* (1992) find an analogous result for a simultaneous search model with horizontal product differentiation, where search uncovers the consumer's match with the firm's product.¹³ They find that, as the search cost increases, consumers sample fewer firms, and the equilibrium price is merely the full information oligopoly price for a market with a number of competitors equal to the sample size selected by consumers.¹⁴

The setting in Burdett and Judd (1983) is coherent with that of Stigler (1961). Although Stigler argued that price dispersion existed because consumers could not compare all prices, he did not explicitly analyze how the price distribution endogenously results from the consumer's search behavior. Performing the full equilibrium analysis shows that, although price dispersion and search may emerge as an equilibrium outcome, there may also be a one-price equilibrium with no search beyond a first (monopoly) price quote.¹⁵ In order for the outcome conjectured by Stigler to arise, it is necessary that firm and consumer anticipations are coordinated: consumers search because they expect price dispersion and firms use mixed price strategies because they expect some consumers to observe more than one price.

Such a coordination problem raises the question of the role of intermediaries. Do they facilitate coordination? Do they improve social welfare? The work by Baye and Morgan (2001) provides some answers in a framework that shares many similarities with the one we have just analyzed.¹⁶ In their setting, consumers can join an internet platform by incurring a subscription fee, in which case they observe all prices posted by participating

¹³The next section discusses such search environments in detail for sequential search.

¹⁴Contrary to results in Burdett and Judd (1983), the number of firms sampled can increase beyond 3 and equals the total number of firms in the market if the search cost is low enough. This is because the incentive to search is determined by the exogenous match value distribution rather than by an endogenous price distribution.

¹⁵Armstrong, Vickers, and Zhou (2009b) show that a price cap in a Burdett-Judd type of market reduces the incentives to search (the LHS of (3) above is reduced) and this reduces the fraction of consumers who get two quotes. In turn, this actually increases the average price paid in the market.

¹⁶See also Baye, Morgan, and Scholten (2006) which embeds the broader context and provides useful empirical evidence.

firms. If they stay out, consumers can buy from a local firm. The two differences with the Burdett and Judd (1983) are that: (i) the number of prices observed by the consumer if she incurs the search cost is not chosen by the consumer but is determined by the firms' participation decisions; (ii) firms incur an advertising cost (set by the platform) to post a price on the platform, so this is a model of costly price advertising. Baye and Morgan characterize an equilibrium where all consumers choose to join the platform and firms play a mixed strategy in which they either join the platform and pick a price according to a distribution with support up to the monopoly price, or else they do not advertise and charge the monopoly price. Obviously, this two-sided market setting does not eliminate the coordination problem: as is standard, agents on each side participate only if they expect sufficient participation by agents on the other side (this is the classic "failure to launch" problem: more generally, coordination issues are addressed in the voluminous literature on the "chicken and egg" problem, following Caillaud and Jullien, 2003). Furthermore, Baye and Morgan show that the intermediary can worsen social welfare if individual consumer demand is not sufficiently price elastic. In a specification where demand elasticity is small enough, it is never socially optimal that such a platform is created if it involves any setup cost, because the social welfare benefit from the resulting decrease in price is too small. Yet, an intermediary can profitably create such a platform through which it extracts some of the producer and consumer surplus through the subscription fee and the advertising fee.

Whereas the results presented so far may be viewed as a partial validation of Stigler's conjecture, they assume that search is "simultaneous", which may hold when there are response lags as with getting builders' quotes for one's deck extension, but does not seem consistent with many modern actual search environments (including internet search). As we now show, viewing search as sequential can lead to a drastic rejection of Stigler's analysis.

2.3 Sequential search and the Diamond paradox

Assume now that, in the setting outlined in the previous subsection, consumers may choose whether or not to incur the search cost s to obtain a new price quote after observing a firm's price (up to the point where they have observed all firms' prices). First, it is immediate that the argument provided above (for simultaneous search) also shows that there always exists a monopoly pricing equilibrium at which there is no search when search is sequential. The drastic difference between the two environments arises when we look for the possibility of other equilibria.

For the simple unit demand specification, there is no scope for pricing above the consumer valuation r . Consider now a situation where firms charge prices strictly below r . If this were an equilibrium, consumers know that the lowest such price, denoted here \underline{p} , is the lowest price they can hope for and will therefore always buy if they encounter that price. Now, a firm charging \underline{p} could increase its price slightly by some amount less than $\min\{s, r - \underline{p}\}$: no consumer would choose to search on, anticipating that the best she could find is another firm charging \underline{p} . Hence there cannot be an equilibrium where \underline{p} is charged

in equilibrium.¹⁷ The argument can readily be generalized to price elastic demand as long as monopoly profit is single peaked.¹⁸

This result has come to be called the Diamond paradox after Diamond (1971). It is puzzling on two counts. First, monopoly pricing prevails even if the search cost is very small or the number of competitors in the market is very large. Second, consumers never search in equilibrium. Both predictions are at odds with casual observation. The rest of this chapter discusses how the introduction of some heterogeneity among agents on either side of the market might resolve or temper the puzzle.

As we noted in our discussion of simultaneous search, the monopoly pricing equilibrium cannot exist if individual demand is completely inelastic and the first price quote is costly. As noted by Stiglitz (1979), this also holds true for sequential search and (as we explain below) even if consumers have heterogenous valuations. Thus search costs, no matter how small, lead to market unraveling. However, an active market can exist if individual demand is sufficiently price elastic so that consumer surplus at the monopoly price is strictly positive.

2.4 Price advertising

In his seminal article, Butters (1977) argued that price advertising is an important channel through which the Diamond paradox can be overcome. He also makes the point that the cost of advertising combined with buyer imperfect information generates price dispersion.¹⁹ To illustrate the point, consider a simple setting where each firm may choose to inform *all* consumers about its price at some fixed cost A , and the advertising decision and price are chosen simultaneously by firms. If a consumer is informed about a firm's price through advertising, she may buy from it at no additional cost. Otherwise, we retain the setting above except that there are only two firms.²⁰

First, if both firms advertised with probability 1, then the outcome would be Bertrand competition with marginal cost pricing, which would yield a profit of $-A$, whereas a firm can ensure at least zero profit if it does not advertise. Second, if neither firm advertises, the outcome is the Diamond outcome with both firms charging r .²¹ However, if $A < \frac{mr}{2}$,

¹⁷For mixed strategies, think of \underline{p} as the minimum of the support of a firm's mixed strategy; the argument goes through for a firm charging a price in the neighborhood of \underline{p} .

¹⁸See Renault (2016) for a sketch of the proof.

¹⁹See Renault (2016) for a detailed discussion of the main contributions in the literature following Butters (1977).

²⁰Janssen and Non (2008) consider a similar advertising technology in a more elaborate setting.

²¹Equivalently, the model applies in a non-search setting when r is an exogenous common posted price set by firms and is therefore the default price if consumers get no ad. Then advertising a price corresponds to offering a discount. As we show below, there is a positive probability of not

each firm could profitably deviate by advertising and slightly undercutting r to capture the entire market. Hence, in a symmetric equilibrium firms necessarily mix between advertising and not advertising. We now describe a symmetric equilibrium.

We seek an equilibrium where, consumers expect a firm that does not advertise to charge the monopoly price r . A firm must be indifferent between not advertising and charging a price of r , and advertising some price p in the support of the equilibrium pricing strategy conditional on advertising. Because no advertised price will exceed the valuation r , a firm that does not advertise sells only if the competitor does not advertise either (with the tie breaking rule that, when indifferent, a consumer picks the advertised price). We show in the Appendix that equilibrium profit is A and the probability of advertising a price at or below p is given by

$$(4) \quad F(p) = 1 - \frac{2A}{mp},$$

for $p \in [\frac{2A}{m}, r)$. Hence firms do not advertise (charging price r) with probability $\frac{2A}{mr}$. Note that a firm that does not advertise could not earn more by charging a price below r , so this characterizes a Nash equilibrium. This equilibrium has the interesting property that firm profit does not depend on search costs but rather is entirely determined by the advertising cost. This result does not hold in the more elaborate model of Robert and Stahl (1993), where each consumer faces a non-trivial search problem because she is not reached by an ad from all firms that advertise: she must therefore update her beliefs about the pricing behavior of those firms from which she has received no ad. In that setting, a lower search reduces prices and profits. However, prices do not tend to marginal cost as the search cost goes to zero. By contrast, driving advertising costs to zero does lead to marginal cost pricing, which is also the case in our simple setting.

Finally and importantly, although this equilibrium exhibits some price dispersion, consumers never search in equilibrium. This is also the case in the analysis of Robert and Stahl (1993) despite the influence of the search cost on prices. As we argue below, it is actually quite challenging to avoid such an outcome in a setting with homogeneous products unless we allow for heterogeneity on both sides of the market. Such two-sided heterogeneity arises quite naturally in settings with horizontally differentiated products, to which we now turn.

3 Matching products to consumers

By considering consumer search in markets for horizontally differentiated products, Wolinsky (1984, 1986) initiated a very novel and fruitful approach. Up to now, and in line with

advertising, and therefore a positive probability of transactions at the “regular” price. Anderson, Baik, and Larson (2016) analyze a similar model with heterogeneous consumer valuations, and render endogenous the posted prices too.

the search literature up to the 1980s, our central question has been the emergence of price dispersion, which is needed to create some motive for consumers to engage in costly search in a market for a homogeneous product. Horizontal product differentiation introduces an alternative motive to search, which results from the consumer's desire to find a satisfactory taste match with the purchased product. We start with a first look at the optimal search problem in the simple case where all the alternatives are *a priori* equally attractive to search.

3.1 Optimal search behavior

A consumer chooses one among n alternatives, where alternative i is expected to yield utility u_i , where u_i is a random variable i.i.d. across alternatives: it could be for instance that all products are sold at the same price and match utilities are i.i.d. or, as in the previous section, a homogeneous product is sold by firms playing a symmetric mixed strategy in prices. Let G denote the distribution function of u_i . The consumer may uncover her true utility with alternative i by incurring a search cost $s > 0$. At any time, she may enjoy any of the alternatives for which she has learned her utility, with no additional cost: this is the so-called free recall assumption.²²

Consider first the choice between enjoying some known utility u and incurring the search cost to find out about the utility of a single alternative u_1 . If the consumer searches, alternative 1 will be preferred only if $u_1 > u$, and the consumer will otherwise enjoy utility u from the free recall assumption. It is therefore optimal to search if and only if

$$(5) \quad E(\max\{u_1 - u, 0\}) = \int_u^{+\infty} (x - u)dG(x) > s,$$

(throughout the chapter, our tie-breaking rule is that a consumer does not search when indifferent). The left-hand side of the inequality is zero for u larger than the maximum of the support of u_1 and it is strictly decreasing in u , from $+\infty$ to zero, for u less than the maximum of u_1 (the derivative with respect to u is $G(u) - 1 < 0$). Hence, if \hat{u} is the value of u at which (5) holds with equality, then, the consumer should search alternative 1 if and only if $u < \hat{u}$.

Now assume, without loss of generality, that the consumer searches through the n alternatives from option n to option 1 in decreasing order of the index i . If, after searching all alternatives up to alternative 2, the consumer holds utility u as her best option, then she searches firm 1 if and only if $u < \hat{u}$. Now consider the decision whether or not to search alternative $i > 1$ if the best utility held thus far is u , and assume that it is optimal to search $i + 1$ if and only if $\max\{u, u_i\} < \hat{u}$. If $u \geq \hat{u}$, the consumer anticipates she will not search beyond alternative i and, by construction, it is not optimal for her to sample alternative i alone. Next, if $u < \hat{u}$, it would be desirable for the consumer to search alternative i

²²See Kohn and Shavell (1974) for a formal treatment of an analogous problem.

even if she anticipates she would stop there and, if she does search beyond alternative i , it is because the expected utility from searching on exceeds $\max\{u, u_i\}$. The consumer should therefore sample alternative i . This shows by induction that, for any $i = 1, \dots, n$, the consumer should search alternative i if and only if the best utility she currently holds is $u < \hat{u}$. In other words, the consumer's optimal behavior is "myopic" in the sense that she always behaves as if there were only one alternative left to be sampled.

As an illustration consider again the case where the consumer has valuation v for some homogeneous product and her uncertainty concerns the price charged by firms $i = 1, \dots, n$, which play the same mixed strategy. Then $u_i = v - p_i$ and, letting $r = v - \hat{u}$ be the consumer's reservation price, from the definition of \hat{u} , r must solve $E(\max\{r - p_i, 0\}) = s$. If, as is typically the case, firms never charge a price larger than r in equilibrium so that $r - p_i \geq 0$ with probability 1, then r solves $r - Ep_i = s$.

Finally, it is readily seen that the same search rule applies if there are infinitely many alternatives. The consumer then solves a dynamic programming problem where the state is the best utility uncovered thus far, u . If u is such that the consumer prefers enjoying u to searching on, then she knows that if she searches one more alternative she will stop searching afterwards. Hence, u is such that the consumer would not wish to search one more alternative so that $u \geq \hat{u}$. Furthermore, if $u < \hat{u}$, so the consumer would wish to search with one more alternative left, as in the finite horizon case, she wants to search all the more if she has the option of searching on beyond the next alternative.

We next embed this optimal search behavior in a model of price competition in a market for horizontally differentiated products.

3.2 Equilibrium and comparative statics

Suppose now there are n products sold by n firms with identical constant marginal costs c . There are m consumers per firm with unit demand, where a consumer's utility from buying product i at price p_i is $u_i = \mu\epsilon_i - p_i$: ϵ_i , $i = 1, \dots, n$ are i.i.d. across products and consumers (the consumer index is dropped to ease notation) and $\mu > 0$ is a scaling parameter reflecting the degree of taste and product heterogeneity (the limit case $\mu = 0$ corresponding to a homogeneous product). The random terms ϵ_i have distribution function F and density f with support $[a, b]$. The outside option of not buying has utility zero. As is standard,²³ firms only know the distribution of ϵ_i . They select prices simultaneously in a first stage.

Next, assume that consumers do not initially observe their realization of ϵ_i with the various products nor the prices charged by firms. They may however find out about both through sequential search before making a purchase. We focus here on the case of random search, which can only be rationalized if consumers have identical expectations about all

²³See Perloff and Salop (1985) and Anderson *et al.* (1992).

firms' pricing behavior, as in section 2 above. Furthermore, we look for an equilibrium where all firms charge the same price p^* so that, on the equilibrium path, observing one firm's price does not provide any information about the price of the remaining firms: in an equilibrium where different firms charge different prices (though consumers do not know *a priori* which firm is charging which price), the search process would involve learning because a consumer observing, say a low price with one firm, would infer that remaining firms are charging higher prices. The equilibrium concept is perfect Bayesian equilibrium and we impose the additional restriction that consumers hold passive beliefs about prices of firms remaining to be searched if they observe a deviation in price (this restriction would actually apply if we used the more restrictive sequential equilibrium concept because firms choose prices simultaneously and a firm's pricing should not signal what the firm does not know, as shown by Fudenberg and Tirole, 1991). We exclude the coordination failure equilibrium where firms are expected to charge such high prices that consumers choose not to initiate search (unless it is the only possible outcome).

The analysis of optimal search above can be applied to describe consumer behavior. Because all firms charge price p^* in equilibrium, a consumer's utility from buying firm i 's product is $u_i = \mu\epsilon_i - p^*$. Letting G be the distribution function of u_i , the reservation utility \hat{u} solves

$$(6) \quad \int_{\hat{u}}^{+\infty} (u_i - \hat{u}) dG(u_i) = \mu \int_{\hat{x}}^b (\epsilon_i - \hat{x}) f(\epsilon_i) d\epsilon_i = s,$$

where $\hat{x} = \frac{\hat{u} + p^*}{\mu}$. As can be seen from (6), the parameter \hat{x} is exogenously determined by the model's fundamentals: the search cost s , distribution F and scale parameter μ . The value $\mu\hat{x}$ is the consumer's reservation utility associated with uncovering her match with a product at cost s if the product was available for free: it may be derived graphically from the inverse demand for the product noting that, by integration by parts, $\int_{\hat{x}}^b (\mu\epsilon_i - \mu\hat{x}) f(\epsilon_i) d\epsilon_i = \mu \int_{\hat{x}}^b (1 - F(\epsilon_i)) d\epsilon_i$. The middle term in (6) is strictly decreasing in \hat{x} for $\hat{x} \in (-\infty, b]$, so \hat{x} is decreasing in search cost s and increasing in μ , reflecting the decrease in the incentive to search if search costs are higher or the uncertainty about the match realization is lower.

The implication of the consumer's search behavior for a firm's demand is straightforward. If some firm i charging price p_i competes with the option of searching another firm, then the consumer chooses to buy product i if and only if $\mu\epsilon_i - p_i \geq \hat{u} = \mu\hat{x} - p^*$, which happens with probability $1 - F\left(\hat{x} + \frac{p_i - p^*}{\mu}\right)$.

As was evidenced by Wolinsky (1986), the monopolistic competition version of this model provides a very tractable and elegant setting to analyze imperfect competition with consumer search: with an infinite number of sellers, no consumer ever "comes back" to a firm sampled earlier. Next we apply this setting, drawing on Anderson and Renault (1999), to investigate the economics of consumer search with product match heterogeneity. Under monopolistic competition the analysis is greatly simplified because the only competition a firm faces arises from a consumer's option to search on. Indeed, the other two potential sources of competition are the outside option and all the other firms. However, if the

consumer has found it optimal to start searching in the first place, the outside option is dominated by continuing to search. Similarly, if search has been preferred to purchasing some product in some round of search, then search will be preferred indefinitely. Hence, conditional on the consumer ever reaching firm i in her search, the probability that she buys product i is $1 - F\left(\hat{x} + \frac{p_i - p^*}{\mu}\right)$. In equilibrium, where $p_i = p^*$ for all i , the probability that a consumer searches on after any round of search is $F(\hat{x})$. Because there are m consumers per firm and search is random, the demand for firm i is given by

$$(7) \quad D(p_i, p^*) = m \frac{1 - F\left(\hat{x} + \frac{p_i - p^*}{\mu}\right)}{1 - F(\hat{x})}.$$

In order for this market to exist, consumers should search in the first place, so we need $\hat{u} = \hat{x} - \frac{p^*}{\mu} > 0$. Hence, the conditional probability of a purchase is less than monopoly demand $1 - F\left(\frac{p_i}{\mu}\right)$. This means that the firm is up against an alternative for the consumer, search, which is more attractive than the outside option, so the inverse demand is correspondingly lower.

If the distribution of the random utility terms ϵ_i satisfies the increasing hazard rate property (i.e., log-concavity of $1 - F(\cdot)$) then, as shown in Appendix B, profit is quasiconcave in price p^i and the equilibrium price has the following simple closed form solution

$$(8) \quad p^* = c + \frac{\mu[1 - F(\hat{x})]}{f(\hat{x})}.$$

Because \hat{x} is decreasing in search cost s and the hazard rate is $\frac{f}{1-F}$, the increasing hazard rate property (log-concavity) also implies that price increases in the search cost. It tends to marginal cost c if $s = 0$ so that $\hat{x} = b$, provided that $\frac{1-F(\hat{x})}{f(\hat{x})}$ tends to zero as \hat{x} tends to b (consumers keep on searching forever)²⁴, and it increases as s increases to a level such that $\hat{x} - \frac{p^*}{\mu} = 0$, where p^* is the monopoly price (at this point, search is no more attractive than the outside option). For larger search costs, the consumer gives up searching all together.²⁵ When this is the case, the market collapses (a finding we noted earlier for homogeneous products when search generates insufficient surplus at the monopoly price). Notice that this knife-edge tipping result is smoothed if consumers have heterogeneous opportunity costs for initiating search (or analogously, heterogeneous outside options as in de Cornière, 2016). Then, for given s , only those consumers with low enough “entry” costs start the search process. There is then another effect (in addition to the price effect) from changing s . This

²⁴As discussed by Anderson and Renault (1999), the limit condition on the inverse hazard rate is related to the condition derived by Perloff and Salop (1985) in order for price to go to marginal cost as the number of firms becomes infinite.

²⁵This analysis assumes that monopoly price is larger than μa so the value of \hat{x} at which price reaches its monopoly level is strictly above a .

is a volume effect from the number of consumers participating in the market: a lower s both increases participation directly by decreasing costs of later search and indirectly through decreasing price (see subsection 5.3 for more results on the impact of changes in search costs on consumer participation).

Regarding product and taste heterogeneity, captured by μ , the impact of a change is ambiguous. On the one hand, an increase in heterogeneity has a direct positive impact on the price captured by the direct inclusion of μ in the price expression. This reflects the standard increase in market power associated with more product differentiation if consumers are perfectly informed. However, as was pointed out above, more product differentiation implies a lower \hat{x} and hence induces more search. This intensifies competition and lowers prices (which is the case under the increasing hazard rate property). Intuitively, the first effect dominates if consumers are likely to sample many sellers so the situation is close to perfect information: this is the case if μ is large so consumers' incentive to find out about their match is high. By contrast, if product and taste heterogeneity is limited (μ small) then an increase in heterogeneity may induce a drop in price because the increased search activity of consumers intensifies price competition. We show in Appendix B that, for a low enough μ , if $f(a) = 0$ then price must be decreasing in μ : the gist of the argument is that for the values of μ such that \hat{x} is close to a (and these values are bounded away from zero as long as the search cost is strictly positive) price goes to infinity as \hat{x} goes to a .

What do we learn from the above analysis about how markets are impacted by the development of the internet? Typically, this development has drastically reduced the level of search costs. The above model predicts a fall in the market price resulting from an increase in search activity leading to better consumer information. By contrast, in the homogeneous product settings we have discussed thus far, although price falls when search costs fall, the equilibrium search behavior is unaffected by changes in search costs (e.g., if search is sequential, consumers stop after one price quote).²⁶ As we explain below, this property remains valid for many search models with a homogeneous product, even if agents are heterogeneous. Because search costs are reduced, the additional information acquired by consumers is not necessarily associated with more resources devoted to search (e.g. more time spent searching). We actually show in Appendix B that, with the increasing hazard rate property for match values, the total search cost incurred by consumers goes down. This result, however, is obtained while keeping the supply side unaffected, in particular regarding the nature and diversity of products available in the market. We return to the issue of endogenous product choice in subsection 4.3. To capture the impact of an increased product variety due to the increased number of sellers, we now briefly discuss the oligopoly setting.

The analysis of (symmetric) oligopoly is a little more complex, but retains the qualitative

²⁶With simultaneous search, the probability that consumers sample a second price increases only in the high search intensity equilibrium. It actually decreases in the low search intensity equilibrium. In any case the maximum number of prices sampled remains 2.

properties of the monopolistic competition model.²⁷ The key difference from monopolistic competition is that some consumers reach the end of the sample set of firms, and “come back” to the one they like most or do not buy at all. This (rather appealing) feature implies that the demand facing a firm comprises two components, from those consumers passing through for the first time, and from those who have not found a good enough match to stop without sampling the full set. The latter behave like consumers with full information in a standard product differentiation context, modulo the proviso that their valuations for all goods (in equilibrium) are below the stop threshold.

The additional driver of the equilibrium price is the competition for those who may come back. The more firms there are, the fiercer the competition (as per standard discrete choice models), because a firm has to beat the best of n other options including the outside option. Hence, the price is higher when there are fewer firms. Put another way, individual demand is more inelastic with fewer firms because there are fewer substitutes. Anderson and Renault (1999) show this intuition is borne out in a covered market under the increasing hazard rate assumption.²⁸ They also consider endogenous entry and find that search costs exacerbate the over-entry that characterizes such markets with perfect consumer information:²⁹ higher search costs make entry more profitable by increasing market power, and they reduce the benefit of having more variety because consumers sample fewer products. The reduced search costs resulting from online shopping should be expected to mitigate this inefficiency. Furthermore, it is likely that the online technology reduces entry costs so that the number of sellers in the market increases despite the weaker market power. This increased product variety is desirable. There is also a market expansion effect due to increased buyer participation when consumers are heterogeneous regarding the decision to start search, as discussed above and further in subsection 5.3. The welfare implications of this market expansion are non-trivial: if the volume effect dominates the price effect, then profitability of entry increases. Moreover, as we explain in subsection 5.3, if search costs are heterogeneous the price effect might also be positive.

Finally, Zhou (2014) provides some caveats to these conclusions. He considers a multiproduct search market in which consumers search for more than one good, and that a search reveals the match values (and prices) of all goods sold in the store. For simplicity, suppose that valuations for goods are independent, so they are neither complements or substitutes, per se. Notice that the assumption of free recall implies that a consumer will never buy one product and keep searching for others, for there is a chance of finding a better product later, and she can always come back costlessly. This feature implies that reducing the price of one product at a firm will also raise the demand for the firm’s other products, which Zhou calls the *joint search effect*: a firm can thus induce consumers to stop and buy

²⁷Anderson and Renault (1999) show the comparative statics results described above for monopolistic competition still hold with the increasing hazard rate property and a covered market.

²⁸We are not aware of any general results when allowing for an outside option.

²⁹See Anderson *et al.* (1995) for the perfect information case.

its other products, rendering products complements. If search costs were to decrease, the joint search effect is weakened and, as he shows, this can dominate the standard effect that works through prolonging search, and actually cause prices to rise.³⁰

3.3 Mergers and cartels

Moraga-González and Petrikaitė (2013) study the impact of a merger in an oligopoly with consumer search and differentiated products. There are two main insights. First, anticipating that insider firms will charge a higher price than outsiders, consumers begin searching through the latter first.³¹ This order of search penalizes the merged entity by reducing the consumer base to which it has access as compared to the pre-merger situation where consumers search randomly. If search costs are large enough, this adverse effect may actually outweigh the benefit from coordinated pricing afforded by the merger, so the merged entity's profit ends up being lower than the participating firms' pre-merger joint profit. This result contrasts with previous literature on mergers with price competition. This type of "merger paradox" is typically expected to arise in a context of quantity rather than price competition. Here the disincentive to merge results from the impact of the merger on consumer search behavior rather than from its impact on the strategic interaction between firms.

The second insight is that the above may only be a short run consequence of the merger. In the long run, the merged entity might be able to reorganize its commercial activities so as to facilitate the consumers' access to information about its products and prices. When searching that firm, consumers can then obtain information about their match with a wide range of products by incurring a low search cost. This may then induce consumers to search through the insider products before the outsider products. This prominence of insider products may lead the merged firm to charge lower prices than outsiders. Because outsiders are searched last, their consumer base shrinks and for high enough search costs their profit is below its pre-merger level. Finally, thanks to search cost economies, consumer welfare as well as social welfare may be increased by the merger.

These results underscore that allowing for search costs brings new considerations to the analysis of mergers. Much of the standard merger analysis contrasts the anti-competitive effect of mergers with the pro-competitive benefits due to the exploitation of economies of scale by the merged firm. Here the long term benefit of the merger stems from the search costs saved by consumers, who may find out about all the products sold by the merged firm while incurring the search cost once. The results also show that the prediction about the impact of the merger on prices may critically depend on how the consumers' search

³⁰Schulz and Stahl (1996) have a product differentiation model in which prices rise with entry when search reveals all match values (and prices are known in advance).

³¹Ordered search is discussed in considerably more detail in the next sub-section and all the subsequent section finales.

behavior is affected.

Janssen and Moraga-González (2008) explore the impact of mergers in a market for a homogeneous product where consumers have heterogeneous search costs. Their setting is similar to that of Stahl (1989), described in subsection 5.2: some of their findings confirm the results in Moraga-González and Petrikaitė (2013). In particular, they also find a merger paradox and the possibility that a merger with no efficiency gain in production may benefit consumers, although the underlying mechanisms are quite different.

Petrikaitė (2016) looks at the factors facilitating collusion in such a setting. The main insight is that a search cost increase has two countervailing effects on cartel stability.³² On the one hand, it decreases the deviation gain because fewer consumers are aware of the deviation. On the other hand, it makes the punishment milder because a higher search cost means more market power.³³ In total though, she shows that the first effect dominates and cartel stability is enhanced if search costs are larger.

3.4 Ordered search

In keeping with the corresponding literature, we have treated search as random thus far (with the exception of the merger analysis of the previous sub-section). Even when search is sequential, we have assumed that the order of search is chosen randomly across options. This has the convenient property of allowing for symmetric equilibria, but there are natural circumstances when sequential search follows an order (which might sometimes differ across consumers). For example, consumers might travel to closer stores before further ones, so geography intervenes (see for example the description in section 5.4 of the model of Arbatskaya, 2007). More subtly, there may be equilibria at which consumers are expected to follow a particular order, and firms price accordingly. Indeed, as argued by Armstrong (2016), the symmetric equilibrium (of Wolinsky, 1986, and Anderson and Renault, 1999) is unstable in the sense that if more consumers choose a particular firm first, then, insofar as that firm's price is lower (which is a key point that we develop below), then other consumers would also want to search the more popular firm first. This tips the equilibrium naturally to one of ordered search. If, in addition, being searched earlier is more profitable,

³²The working definition of cartel stability here is based on the critical discount factor analysis from simple trigger strategy equilibria in infinitely repeated games. Full collusion is sustained as long as the discount factor, $\delta \geq \hat{\delta} = \frac{\pi^D - \pi^C}{\pi^D - \pi^N}$ where the superscripts denote Deviation, Cartel (or Collusion), and Nash (i.e., punishment) payoffs. Greater stability is interpreted as lower $\hat{\delta}$, which is achieved if the excess of deviation over collusive profits goes down relative to deviation over punishment. The condition can be phrased alternatively in terms of differences over punishment profits as $\hat{\delta} = 1 - \frac{\pi^C - \pi^N}{\pi^D - \pi^N}$.

³³These effects are quite analogous to increasing product differentiation in a perfect information setting.

this introduces the possibility that firms might pay for more prominent positions on websites (as with position auctions), or on supermarket shelves (as with slotting allowances), or closer to consumers (as with geographical locations). Such paying for prominence (in the terminology of Armstrong and Zhou, 2011) may involve a variety of firm practices as discussed in Armstrong (2016), such as advertising a low price (we briefly return to price directed search in the conclusion) or non-price advertising where consumers coordinate their search on the seller which advertises the most, as in Haan and Moraga-González (2011).

Taking the leading example of the differentiated products model of section 3.2 (match values are heterogeneous, but i.i.d., as per Wolinsky 1986) we suppose now that prices are expected to differ across firms. Then the consumer will clearly follow the order of increasing expected prices, although we still need to determine her stopping decision. We argued above in section 3.2 that if all prices are expected to be the same, then she searches on if the best match thus far is less than some threshold \hat{x} ; moreover, this is true independent of the number of firms left. That is, she uses a myopic rule, and additional equally attractive search options do not make search more appealing. As should be apparent intuitively, this logic extends to when prices differ. If firm i has the next lowest expected price, and that price is p_i^* , then the consumer searches if her best utility with firms she has already visited is less than $\hat{x} - p_i^*$. The prospects from searching firm i cannot improve when the later options are increasingly less attractive.

This case already covers most of the models in the literature on ordered search with product heterogeneity. Because consumers search by lowest prices first, consistency requires firms searched earlier do indeed want to price lower. To see the tension, consider a simple duopoly case. Let us compare the two firms' price elasticity of demand if they charge the same price. Insofar as the first firm sampled gets all traffic (the first bite of the apple), then its demand is larger, which all other things equal implies a lower price elasticity. However, as we now argue, its demand derivative is larger (in absolute terms), and hence the source of the potential ambiguity. It is instructive into the nature of pricing to see why demand derivatives are different by considering the simpler case of a covered market. Then, in the absence of search the derivatives are the same in a duopoly (each firm loses the marginal consumers to its rival from a price hike). With firm 1 being visited first, we can split the consumers into 2 groups. One group buy directly from firm 1; they stop before finding out their match with firm 2. The rest of the consumers sample both firms, and are thus fully informed. All consumers observe a price rise by firm 1, and so it loses consumers to firm 2 from both the fully informed and those who erstwhile stopped at it. But only searchers observe an (unanticipated) price hike from firm 2. That is, firm 2 does not lose consumers over the transom between it and its rival at the margin of discouraging them from searching, whereas firm 1 does. Therefore 1's demand derivative is larger, as claimed.

Extending the reasoning to several firms, the demand of a firm searched later is smaller, but is less sensitive to price changes because the firm cannot communicate a price drop: consumer decisions to stop are driven by (rationally) expected prices, but not actual ones. The earlier is a firm in the search order, the more consumers see a price change, but the larger is demand. However, as shown in Armstrong (2016), an increasing hazard rate in

the product match distribution (log-concavity) is a sufficient condition for demand to be more elastic, and hence for price to be lower earlier in the search order.³⁴ When a consumer arrives at a later firm, the firm can infer that she does not like the earlier products: she searches even if she expects later prices to be higher. This gives the later firms extra market power.

The possibility of bidding for prominence adds yet a further requirement on the equilibrium. In addition to prices rising with the consumer search order, it should also be the case that profits fall. Otherwise, if this condition were revoked at some point in the order, firms would want to wait for later positions, and not pay a premium to be earlier. The requirement is actually necessarily satisfied from a simple revealed preference argument. Indeed, a firm can always choose to charge the equilibrium price of its successor in the queue. In this symmetric product setting, this necessarily yields a larger profit for the earlier firm. Hence if it chooses to price lower, it must be earning more profit than the next firm.

Unfortunately, for reasons similar to those we discussed in the case of random search, there is no existence argument relying on general properties on the match distribution (such as an increasing hazard rate). The early articles on the topic, such as Armstrong, Vickers, and Zhou (2009a) or Zhou (2011), used a uniform match distribution. As with random search, the monopolistic competition setting provides a simple fix. It is then easy to characterize a symmetric price equilibrium. Because price is expected to be the same at all firms, a consumer’s search behavior is identical to that derived for random search. For reasons analogous to those discussed in subsection 3.2, the infinite number of firms implies that each firm only competes with the consumer’s option to continue search. Because the consumer’s search prospects are stationary, all firms charge the same price: although firms placed earlier sell more, their demand derivative is also larger exactly by the same proportion, so elasticity is the same at all positions. As with random search, the equilibrium price is given by (8). In this case, then consumer welfare and aggregate firm profits remain unaltered, so market performance is neutral. However, the first firm searched has more equilibrium sales than the second, etc., so that profits decrease with exponential decay through the order (all prices are the same, and each firm after the first one serves a fraction $F(\hat{x})$ of its predecessor’s demand). This opens up the possibility that firms would want to pay for prominence, which we reconsider in the next section where firms are assumed to differ *ex ante*.

4 Seller heterogeneity

Much attention has been devoted to the theoretical possibility of price dispersion that results solely from imperfect consumer information (see section 2). In practice, it is likely that sellers that charge different prices face different costs and/or demand.

³⁴Rhodes and Zhou (2016) obtain a similar result in a variant with two firms, one of which is multiproduct.

4.1 Searching for an efficient seller

A (somewhat obvious) first point is that introducing efficiency differences across sellers naturally generates price dispersion. Following Reinganum (1979), assume a market where m consumers have identical individual downward sloping demands, $d(p)$ at price p , for a homogeneous product sold by a continuum of firms. Let firm i have constant marginal costs, c_i , and these differ across firms according to a continuous distribution on an interval $[\underline{c}, \bar{c}]$, with $\underline{c} > 0$. Assume $d(p)$ is well behaved so that for all i , firm i 's monopoly profit $md(p)(p - c_i)$ is single-peaked in p , and the corresponding monopoly price is $p^m(c_i)$. A well-behaved demand also ensures that monopoly price is strictly increasing in marginal cost. Consumers can only observe prices through random sequential search, where $s > 0$ is the search cost.³⁵ The marginal cost distribution is common knowledge but consumers do not know its realization for each firm.

Consider now the following strategies. A consumer searches at least one firm and keeps on searching as long as the best price observed so far strictly exceeds some reservation price r . Firm i charges $p^m(c_i)$ if and only if $p^m(c_i) \leq r$, and charges r otherwise. Let \tilde{c} be the marginal cost such that $p^m(\tilde{c}) = r$ (which exists as long as r is not too large). Optimal search requires that r solves $r - E_{c_i < \tilde{c}} p^m(c_i) = s$ (note that, due to the continuum of firms assumption, a consumer draws from the same price distribution at any round of search). As r increases from $p^m(\underline{c})$ to $p^m(\bar{c})$, the left hand side increases from 0 to $p^m(\bar{c}) - E p^m(c)$. For s in this range, r defines an optimal strategy for consumers given the firms' pricing (as long as consumer surplus at all those monopoly prices exceeds s). The proposed strategies for firms are profit maximizing because all prices are below r , so that all consumers buy from the first firm sampled. Thus each firm should maximize its profit on the mass of consumers $\frac{m}{n}$ visiting it first, subject to the constraint that it should not price above r . For $s > p^m(\bar{c}) - E p^m(c_i)$, the equilibrium entails all firms charging their monopoly prices.

Although introducing cost heterogeneity generates price dispersion that makes search potentially appealing, it does not substantially overturn the Diamond paradox. Only the least efficient firms are forced to price below their monopoly prices. Furthermore, as the search cost tends to zero, the equilibrium prices all tend to the lowest monopoly price, rather than marginal cost (which might have been expected in a market where consumer information about prices is almost free). Note that inefficient firms may remain active even if the search cost is very small. Regarding consumer search behavior, there is no search beyond the first firm in equilibrium, no matter how small the search cost.

We next return to the monopolistic competition model with heterogeneous products of subsection 3.2, assuming now that each firm's marginal cost is distributed identically and independently across firms on support $[\underline{c}, \bar{c}]$ with distribution function H . Consumers do not know firm i 's marginal cost, but the cost distribution is common knowledge. Suppose

³⁵Bénabou (1993) and Bénabou and Gertner, R. (1993) allow for heterogeneity in both production and search costs.

that if firm i has cost c_i , it is expected to charge price $p(c_i)$.³⁶ Then a consumer's utility from buying product i is the random variable $u_i = \epsilon_i - p(c_i)$ (we here assume $\mu = 1$ to simplify notation). The analysis of subsection 3.1 applies, so we may define the threshold utility \hat{u} above which a consumer stops searching. In equilibrium, a firm's pricing decision depends on this threshold, so we denote the equilibrium price of a firm with cost c_i by $p^*(c_i, \hat{u})$. Again letting G denote the distribution function of u_i , \hat{u} solves

$$(9) \quad \int_{\hat{u}}^{+\infty} (u_i - \hat{u}) dG(u_i) = \int_{\underline{c}}^{\bar{c}} \int_{\hat{u} + p^*(c_i, \hat{u})}^b (\epsilon_i - p^*(c_i, \hat{u}) - \hat{u}) dF(\epsilon_i) dH(c_i) = s.$$

As shown in Appendix B, profit maximization implies

$$(10) \quad p^*(c_i, \hat{u}) = c_i + \frac{1 - F(p^*(c_i, \hat{u}) + \hat{u})}{f(p^*(c_i, \hat{u}) + \hat{u})}.$$

Appendix B also shows that the increasing hazard rate property for the match distribution ensures that a firm reacts to a higher reservation utility by charging a lower price, although it adjusts its price by an amount which is less than the reservation utility change. It follows that $p^*(c_i, \hat{u}) + \hat{u}$ is strictly increasing in \hat{u} . This is because log-concavity of demand (implied by the increasing hazard rate) induces a pass-through of demand or cost shifts less than unity (Anderson and Renault, 2003).³⁷ Hence the middle term in (9) is strictly decreasing in \hat{u} (for each realization of c_i the integral with respect to ϵ_i is strictly decreasing). It follows that \hat{u} is uniquely defined and strictly decreasing in search cost s : an increased search cost makes search less attractive.

The above analysis shows that the predictions obtained with homogeneous firms regarding the impact of search costs readily extend to a setting where firms selling differentiated products have different marginal costs: lower search costs induce more search and hence lower prices. However, this setting provides new insights for very high or very low search costs. Because high-cost firms charge higher prices, the critical match value at which consumers prefer searching to buying their product, $p(c_i, \hat{u}) + \hat{u}$, reaches b (as s decreases) earlier for high-cost (less efficient) firms than for low-cost (more efficient) firms. Hence, less efficient firms become inactive if search costs are sufficiently low.³⁸ For very high search costs however, efficient firms exert a positive externality on (and suffer a negative externality from) inefficient firms. This is because the search cost level at which \hat{u} reaches zero (at which point the market collapses) depends on the entire distribution of marginal costs. Hence, this critical search cost value would be higher if the market comprised only efficient firms, and lower if consumers anticipated that all firms are inefficient.

³⁶This framework is a simplified version of that analyzed by Bar-Isaac *et al.* (2012) to which we return in subsection 4.3.

³⁷Weyl and Fabinger (2013) provide further results on pass-through in oligopoly.

³⁸If these inefficient firms drop out when they can no longer sell, the value of \hat{u} is affected because the upper bound on marginal costs is endogenously decreasing in \hat{u} .

Finally, allowing for cost heterogeneity provides some insight into the impact of search costs on price dispersion. With homogenous products, dispersion is reduced with lower search costs to the extent that more firms charge the reservation price. By contrast, with horizontal product differentiation, as long as the reduction in search costs has no impact on the mix of firms (in terms of marginal costs) participating in the market, there is no clear reason why price dispersion should be reduced: whether inefficient firms lower their prices more than low cost firms in response to an increase in \hat{u} depends on properties of the match value distribution, beyond the increasing hazard rate restriction.³⁹ However, if search costs drop so much as to drive inefficient firms from the market, then there is a clear pressure towards reducing the range of prices in the market.

4.2 Learning about sellers' efficiency

Although it is convenient to assume that firm characteristics are idiosyncratic, doing so leaves out an important feature of consumer search, which concerns how consumers can learn about market conditions. Unfortunately, accounting for such learning complicates the analysis considerably. Here we describe the main insights that can be derived from simple specifications and also point out the added difficulties associated with this type of analysis.⁴⁰

Consider again a market for a homogeneous product where sellers have different marginal costs. By contrast with our previous analysis, assume now that these costs are positively correlated due to common shocks such as changes in input prices. As in Bénabou and Gertner (1993), think of a simple duopoly model where consumers first observe one price and then decide whether or not to incur a search cost to find out about the other price. Provided that price is monotonically increasing in cost, a high price observed by the consumer in the first round leads her to revise her expectations upwards regarding the price she will find at the second firm. She will all the more revise her price expectations unfavorably if the uncertainty regarding the common shock is large relative to the uncertainty regarding the idiosyncratic cost differences among firms. This is because more variability in the common cost component means a higher correlation between the firms' costs. This *per se* provides an incentive for firms to charge higher prices when global cost uncertainty increases. However, as emphasized by Bénabou and Gertner (1993), an increased uncertainty and an increased correlation in costs also impact the consumers' incentives to search, which in turn affects the firms' pricing. In particular, increased cost uncertainty resulting in more variability in prices makes search more attractive, which induces more competition and pushes prices

³⁹For a uniform taste distribution (i.e., linear demand), the equilibrium price is $p^*(c_i, \hat{u}) = \frac{1}{2}(1 + c_i - \hat{u})$ so that all prices respond the same way to a change in \hat{u} .

⁴⁰Our presentation leaves out a fairly substantial literature that looks at learning and price dynamics in search models such as Dana (1994), Fershtman and Fishman (1992), Fishman (1996) or Tappata (2009). These studies typically assume homogeneous products and simultaneous search.

downwards. Bénabou and Gertner find that a higher global cost uncertainty increases prices if search costs are high but has the reverse effect otherwise.

Note that if there is no idiosyncratic heterogeneity in firm costs and firms play a symmetric pure strategy equilibrium, consumers learn about the common shock perfectly with only one price observation. This would be the case for instance in the model with product and taste heterogeneity of section 3 if there were some cost uncertainty although all firms would have the same cost. Indeed, Janssen and Shelegia (2016), who consider a duopoly version of the Wolinsky (1986) setting, find that the equilibrium price is higher when consumers are uncertain about marginal cost than if there is no such uncertainty. They find however that, for a high enough search cost, price with cost uncertainty is decreasing in the search cost, whereas the price with no uncertainty is still increasing, and both prices become equal to the monopoly price if the search cost is high enough so consumers stop searching all together. Janssen *et al.* (2011) also explore the role of common cost uncertainty in a market for a homogeneous product with symmetric firms. In their setting, there is however price dispersion due to mixed strategies resulting from some heterogeneity in consumer information.⁴¹ Because of the mixed pricing strategies, consumers cannot learn about the cost shock perfectly. Still, they find the expected price is larger with cost uncertainty.⁴²

As emphasized by Janssen and Shelegia (2016), when comparing the cost uncertainty setting with the standard setting where marginal cost is common knowledge, we should be cautious about what we assume about consumer beliefs. Cost uncertainty naturally leads us to assume that, in a pure strategy equilibrium, any price that is in the support of equilibrium prices is interpreted as a sufficient statistics for the marginal cost. Thus, when a firm deviates upwards, it expects consumers to infer from the new price that the other firm is also charging a higher price. By contrast, the standard approach with no cost uncertainty presented in section 3 assumes consumers have passive beliefs. However, if consumers expected the slightest positive correlation in marginal costs, the justification for passive beliefs would no longer be valid. Janssen and Shelegia show that the equilibrium behavior with cost uncertainty is analogous to the equilibrium behavior in the standard model where consumers have “symmetric” rather than passive beliefs, so they always expect the price at the second firm to equal the price at the first firm they visit.

An important methodological remark is in order here. Once there is learning, the standard optimal search analysis described in subsection 3.1 no longer applies. First, even if the optimal search behavior involves a reservation utility, it should evolve over the search process and depend on the search history because learning affects a consumer’s expectations. Second, as first noted by Rothschild (1974), an even more drastic implication of learning is

⁴¹They use the setting of Stahl (1989) discussed below in subsection 5.2.

⁴²Note though that the equilibrium they characterize only exists for certain parameter configurations. In particular, the search cost should be large enough relative to the uncertainty in marginal production costs.

that the optimal rule may not be based on a reservation utility level.⁴³

We next give a brief discussion of endogenous choice of product design.

4.3 Product choice and search

In our discussion of the impact of the search cost reduction associated with the development of internet search, we pointed out that it is critical to take into account supply side adjustments, in particular regarding the types of products that are available for sale. In the search framework with horizontal differentiation and heterogeneous costs of subsection 4.1, a drop in search costs shifts a firm’s inverse demand downwards, due to the improved attractiveness of search reflected in a higher \hat{u} . Firms might then want to adjust their product designs appropriately. The analysis in Johnson and Myatt (2006) provides a simple setting through which this choice can be investigated.

Here we borrow from Bar-Isaac *et al.* (2012) to understand the change in product mix that is induced by lower search costs. Johnson and Myatt (2006) describe a firm’s choice whether to make its product widely appealing to a large customer base or instead very appealing to a small niche of customers who are willing to pay a lot for it. This choice can be depicted as the firm choosing to rotate its inverse demand by making it flatter or taller (with a flatter curve corresponding to a larger customer base). They show that the optimal choice is always extreme, so a firm chooses either the flattest or tallest inverse demand. Whether a firm chooses flat or tall depends on whether its marginal cost is high or low relative to the vertical position of the inverse demand that reflects how much buyers value the product “on average”. If marginal cost is low, then it is optimal for the firm to sell a lot and it will select the mass market product design.⁴⁴ If, on the contrary, marginal cost is high, the firm prefers to sell less and selects the alternative niche market product design. Then, this analysis leads us to expect that there is some threshold cost such that, all firms with marginal cost above the threshold select a niche product whereas firms with marginal cost below the threshold select a mass market product. Bar-Isaac *et al.* (2012) characterize such an equilibrium in an environment where firms differ by product quality (where a high quality is equivalent to a low marginal cost in our setting). Now if the search cost decreases, leading to an increase in \hat{u} , each buyer’s willingness to pay for the product (as opposed to searching on) drops, making the niche product strategy more attractive. As a result the product mix involves more niche products.⁴⁵

In terms of our analysis in section 3, choosing a niche product corresponds to choosing

⁴³See Janssen *et al.* (2016) for an equilibrium search analysis without reservation price search rules.

⁴⁴These results recall those of Lewis and Sappington (1994).

⁴⁵Bar-Isaac *et al.* (2012) show that the same predictions can arise in a setting with *ex ante* homogeneous firms, where firms mix between the two product designs.

a product with a large μ , reflecting a larger heterogeneity in match values. Because such products are more attractive to search (\hat{x} being increasing in μ), the change in the product mix increases search activity beyond the level that would ensue if product design was exogenous. This is well reflected in the approach taken by Larson (2013). In accord with the results noted above, he shows that firms' equilibrium choices are extreme. He emphasizes the feedback loop in the equilibrium: high reservation values encourage search and so encourage firms to choose highly differentiated (niche) products, which in turn further encourages search. Likewise, low reservation values encourage low differentiation. Larson (2013) also argues that common specifications of consumer preferences are likely to give an asymmetric equilibrium with niche and generic products coexisting.

4.4 Ordered search

Thus far in this section, although we have allowed firms to differ in marginal costs and product quality, we have postulated that consumers search randomly. This requires that the characteristics of a firm remain unknown to a consumer until she visits it. Further assuming an infinite number of sellers allows for describing the optimal search rule as a simple stationary stopping rule. Yet, if firms differ in costs or quality, we would expect them to try to make this information known if it is favorable. A central theme here is that paying for prominence is a means of credibly transmitting such good news. We first reconsider the optimal search problem from a more general perspective than the one of subsection 3.4.

Weitzman (1979) delivers a very clean characterization of optimal search behavior. Suppose that options can differ by the distribution of the utility that they are expected to yield, so that we write G_i as the utility distribution for option i . We retain the independence of the distributions, but we can dispense with their being identical. We also retain costless recall, and search cost s per search.⁴⁶ As Weitzman (1979) shows, the solution to the consumer's problem has a simple algorithmic condition that determines what to do. Surprisingly perhaps, the solution is not to search in order of highest expected value (although this is true in the simplest cases), but is not much more complex: Weitzman's result is gratifyingly simple.

Define for each option a score value \hat{u}_i by

$$(11) \quad \int_{\hat{u}_i}^{+\infty} (u_i - \hat{u}_i) dG_i(u_i) = s,$$

which generalizes (6). The consumer's optimal search is to go through these scores in decreasing order until the next option delivers a score below what is currently held: this may involve costless recall, which is to go back and select whatever previously held option

⁴⁶Weitzman (1979) does allow for different costs per search, as well as discounting of options searched later, with the same characterization we give below.

is best. These scores are myopic reservation values, where \hat{u}_i is the highest utility currently held such that searching option i is optimal. The rule can readily be proved by backward induction similarly to the argument in subsection 3.1.⁴⁷ There are important cases where the reservation values are ranked in an unambiguous manner. First, a distribution has a higher score if options are ordered in terms of first order stochastic dominance, as should be expected. One simple case is when match distributions are i.i.d. and only prices differ. Second, a wider spread yields a higher score when all options are mean preserving spreads of one another: this reflects that it is more appealing to search riskier options. In line with the previous section, this means that consumers would choose to search sellers of niche products first (*ceteris paribus*).

For the most part, the existing literature on ordered search with firm heterogeneity has considered products with different qualities, so that there is a first order stochastic dominance ordering of match value distributions. Even in this simple case, the characterization of a price equilibrium is quite challenging. To see why, consider again the monopolistic competition version of the model in Wolinsky (1986), which has proved in many instances to be quite tractable. Suppose, as in Bar-Isaac *et al.* (2012) that match values for product i are given by $v_i + \epsilon_i$, where ϵ_i is the idiosyncratic match utility that differs across buyers, whereas v_i is some quality measure that is evaluated in the same manner by all. In contrast with the analysis in subsections 4.1 or 4.3, consumers know v_i and may therefore devise an optimal search order, which will be the same for all because they are *ex ante* identical. The optimal order does not necessarily follow the order of v_i because it also takes into account the prices expected by consumers at each firm, p_i for firm i . With no loss of generality, assume the reservation values defined by (11) for each firm i decrease in index i , $\hat{u}_1 > \hat{u}_2 > \hat{u}_3 \dots$. Then consumers start search with firm 1 and move on to firm 2 and so on as long as they hold a utility which is less than the next reservation value. Hence, a consumer searches firm 2 if $\epsilon_1 < \hat{u}_2 - v_1 + p_1$. Among such consumers, those whose match with firm 1 is $\epsilon_1 > \hat{u}_3 - v_1 + p_1$ (and there are such consumers because $\hat{u}_3 < \hat{u}_2$) will not want to search firm 3, and will pick whichever firm they like better between firms 1 and 2. Hence, there are consumers coming back despite the infinite number of firms. This precludes the simple existence arguments based on the increasing hazard rate property we have used thus far for monopolistic competition. Besides, the characterization of equilibrium pricing must deal with an infinite number of demand terms reflecting the populations of returning consumers after additional rounds of search, compounded with the asymmetries resulting from the search order and the heterogeneity in product qualities.

Armstrong *et al.* (2009a), in their extension where firms have different qualities,⁴⁸ overcome these obstacles by supposing that only one firm obtains a prominent position and then the other firms (an infinite number) are searched randomly. Then a consumer who

⁴⁷See Armstrong (2016) for an alternative proof of the result and for some more discussion of Weitzman's rule.

⁴⁸Their specification of quality is different from the one we have used above.

decides to search after visiting the prominent firm enters a world like that described in subsection 4.1 so she keeps on searching until she obtains a utility above her reservation value, which is now stationary. Then no consumer ever returns to any firm. They show that the firm with the highest quality has the highest willingness to pay for prominence. Another route is to simplify the model by only considering a duopoly as in Song (2016).⁴⁹ Still the analysis is quite involved even when using uniform distributions of matches.

An alternative is to follow Chen and He (2011) or Athey and Ellison (2011) and assume that all prices are exogenous and identical.⁵⁰ They suppose that firms only differ in the probability that a consumer is “matched” with the product, i.e. the probability that the consumer wants to buy. Let β_i denote that probability for firm i then, normalizing price and the consumer population to 1, a firm’s profit is merely β_i times the probability that it is visited. Taking a two firm example and assuming firm 1 is visited first, firm 1 earns β_1 and firm 2 earns $(1 - \beta_1)\beta_2$. Firm 2 is willing to pay $\beta_1\beta_2$ to be visited first rather than second. This is symmetric and therefore both firms have the same incremental value from being ranked first. This means that consumers have no reason to expect that the firm ranked first (say on a web page) is the most attractive and that they should therefore start with that firm. One solution is to assume that firms only know their own probability. For instance if probabilities are drawn from i.i.d. distributions with mean λ , then firm i ’s incremental value for being ahead is $\beta_i - (1 - \lambda)\beta_i = \lambda\beta_i$, which is increasing in probability β_i . As explained in subsection 5.4, different search costs for different consumers can also fix this problem.

Anderson and Renault (2016) introduce a demand system that captures some key features of pricing with ordered search while allowing for multiple dimensions of *ex ante* product heterogeneity. As in Athey and Ellison (2011) or Chen and He (2011), consumers stop search upon finding a product with which they are matched. However, this behavior is not exogenously assumed but here results from the firms’ pricing decisions. This is because the demand function ensures that, conditional on being matched, a consumer has a high enough willingness to pay for the product that the firm always finds it profitable to charge a price low enough that she does not search. Firm pricing has the key property that firms price in such a way that consumers find it optimal to follow whatever search order is assumed. The demand system allows for three dimensions of heterogeneity. First, as in the simple fixed price setting above, products differ in how popular they are (the probability of matching a consumer’s need). Second, they differ in how valuable they are to consumers, which is captured by the minimum valuation of a consumer who is matched with the product. Finally, they differ in the heterogeneity of tastes captured by a measure of the thickness of the upper tail of the match distribution. The latter engages the characterization of the optimal search

⁴⁹Song (2016), more generally, supposes that one match distribution is a mean-preserving spread of the other, making it more appealing to search. He shows that joint profits are maximized if consumers search first the more spread distribution.

⁵⁰Chen and He (2011) assume that all products that a given consumer cares about are perfect substitutes, so the Diamond paradox ensues and all firms charge the monopoly price endogenously.

rule in Weitzman (1979) in quite a general manner deeper than simply first order stochastic dominance (as was previously considered). One insight into quality signaling by paying for prominence is that, with endogenous pricing, a firm selling a less popular product may be willing to pay more to be ahead in the search order.

5 Buyer heterogeneity

Buyer heterogeneity has been a much more active and fruitful research direction than seller heterogeneity. More specifically, search cost heterogeneity is a dimension of the search equilibrium problem that has delivered significant progress in our understanding of the impact of search costs on competition. We start though by discussing demand heterogeneity.

5.1 Demand heterogeneity with a homogeneous product.

It might be expected that demand heterogeneity could overcome the Diamond paradox. Some firms could specialize in selling small amounts to consumers with high valuations, while others could charge low prices and cater to more people, including some consumers with lower valuations. We now explain why this is not a very promising research agenda for looking at the impact of reduced search costs.

First, following Stiglitz (1979), it is fairly well known that, if the first search is costly and consumers have (possibly heterogeneous) unit demands, then the market necessarily unravels if consumers with identical search costs must search sequentially to find out about prices.⁵¹ To recapitulate the argument, let us again denote the lowest candidate equilibrium price by \underline{p} . A firm charging that price knows that all consumers entering the market necessarily have a valuation of at least $\underline{p} + s$, because they were willing to incur the first search cost. The firm could then deviate by charging a slightly higher price and not lose any customers. Hence there can be no such \underline{p} , and therefore no equilibrium. It is straightforward to use the same line of argument to establish that the unique equilibrium has all firms charging the monopoly price if the first search is free and monopoly profit is single peaked; moreover, no consumer searches beyond the first firm. To the best of our knowledge, it has not been investigated whether or not this result generalizes to price sensitive demands and/or a monopoly profit that is not single-peaked. However, we argue below that, although it is possible to devise situations where the monopoly price may not sustain, the resulting predictions are not substantially different from those of the Diamond paradox.

It is fairly immediate that then there may be single-price equilibria different from the monopoly price if monopoly profit is not single-peaked. Consider a local maximum of monopoly profit at some price p , such that no local maximum at a lower price yields a

⁵¹The argument here does not require homogeneous products, just that consumers know their valuations before search.

larger profit. Because p is a local maximum, there exists $h > 0$ such that price p maximizes profit on $[p, p + h]$. Now, if there is a finite number of consumer types and if search cost s is small enough, then a price hike with a magnitude of at least h starting from p will result in a drop in surplus larger than s for all consumers. Hence it is possible to sustain an equilibrium with n firms where each firm charges p . In such an equilibrium, no consumer searches and each firm earns one n th of the monopoly profit at price p . By construction, a downward deviation is unprofitable. Furthermore, any price increase such that consumers choose not to search must be less than h and thus does not increase profit. Then we have an equilibrium if the first price quote is for free. Furthermore, if individual demands are price elastic and s is less than the minimum consumer surplus at price p , consumers will participate in the market and the equilibrium is sustained even if the first price quote costs s . Although this shows that it is possible to obtain equilibrium pricing below the monopoly price (if that price corresponds to a local maximum above p), the candidate equilibrium, like the Diamond outcome, involves no search and pricing that is unaffected by changes in the search cost or the number of firms.

Multiple peaks in monopoly profit also allow price dispersion in the form of a mixed strategy equilibrium. This is the case even without demand heterogeneity. Consider a single consumer type and a continuous monopoly profit. Assume there exists some local maximum p less than the monopoly price but yielding more profit than any price below p . Because p is below the monopoly price, continuity of profit requires that there exists at least one price r at which profit is increasing in price and is the same as at price p . Take the lowest such price so that no price in $[p, r]$ generates more profit than prices r and p . Letting S denote the consumer's surplus, assume now that $s < S(p) - S(r)$. Then there exists $\alpha \in (0, 1)$ such that $\alpha(S(p) - S(r)) = s$. This defines a symmetric mixed strategy equilibrium with support $\{p, r\}$ where firms charge price p with probability α and the consumers' search behavior is characterized by a reservation price of r . Note though that consumers do not search. Furthermore, as search cost s decreases, equilibrium prices do not change and the probability α of a low price decreases, so expected prices rise.

Demand heterogeneity is an obvious source of multiple peaks in monopoly demand even if individual demands are well behaved. More importantly, it creates new possibilities for price mixed strategies even if monopoly profit is single peaked. Suppose there are two consumer types with price sensitive demands. Because the two types have different surplus functions, in a symmetric mixed strategy equilibrium, they will have different reservation prices, r_h and r_ℓ , with $r_h > r_\ell$. This opens up the possibility that firms mix between prices in the interval $(r_\ell, r_h]$ and prices below r_ℓ . Type ℓ consumers would then search when getting a first price quote above r_ℓ . This in turn might make it profitable for firms to price below r_ℓ to prevent such search. However, individual demands should be chosen appropriately to make sure that type r_ℓ consumers' demand is sufficiently elastic at prices below r_ℓ so firms would choose to price strictly below r_ℓ with positive probability: otherwise r_ℓ could not be a reservation price. Assuming that this can be achieved, we conjecture the equilibrium would exhibit a comparative statics pattern with respect to search costs similar to the one described in the homogeneous consumer case with multiple peaks: with firms

being more likely to charge high prices if the search cost is smaller. The idea is that, if the search cost became smaller, then the probability of a low price below r_ℓ should become sufficiently small in order for r_h to remain sufficiently high as compared to r_ℓ so a firm finds it profitable not to prevent consumers with reservation price r_ℓ from searching. The example below illustrates a similar logic in a somewhat different environment.

Interestingly, Diamond himself (Diamond, 1987) provides an intriguing twist whereby demand heterogeneity can overcome the Diamond paradox. The reason this works is because he assumes search costs are delay costs instead of a cost per search (we have been supposing the latter throughout). It could be for instance that consumers are hit by direct response advertising while they are surfing on the web or watching TV as in Renault (2016): the delay is then determined by the frequency of advertisements from firms selling competing products. The following simplified version of his model illustrates. Consumers have unit demands and there are only two consumer types: a measure m_ℓ of them have valuations v_ℓ and a measure m_h have valuation v_h , $0 < v_\ell < v_h$, $m_\ell + m_h = 1$. Production costs are zero. In a standard sequential search setting with cost per search $s > 0$, the paradox would sustain because the lowest equilibrium price could not be below v_ℓ , and hence low valuation consumers would never search. Suppose instead that the only cost associated with searching is due to having to wait before being able to consume the product, so the corresponding surplus is discounted by a factor $\delta \in (0, 1)$. Then, a low valuation consumer will not buy if she runs into a price strictly above v_ℓ , but she will buy if she later runs into a product priced at v_ℓ (we assume that she buys if indifferent) and waiting involves no cost because her surplus was zero to begin with. We seek a symmetric two-price equilibrium, with firms mixing between v_ℓ and $p_h > v_\ell$ and the probability of pricing low is α . The higher price cannot exceed v_h . Furthermore, in order to generate some positive profit, it must be low enough to stop a high valuation consumer from searching on in hope of getting the lower price (here we assume that if a consumer finds a high price after searching she buys from the second firm searched). Hence, it must be less than some reservation price r at which a consumer is just indifferent between searching and not searching. This reservation price is solution to

$$(1 - \delta)(v_h - r) = \delta\alpha(r - v_\ell).$$

The left-hand side represents the cost of searching associated with having to consume the product later if the consumer ended up not getting a lower price. The right-hand side is the expected benefit of search, measured by the discounted expected decrease in price. Thus $r = \frac{v_h - \delta(v_h - \alpha v_\ell)}{1 - \delta(1 - \alpha)}$. Because r is less than v_h and demand is perfectly inelastic it is optimal for a firm choosing the higher price to charge r (where again a consumer buys if indifferent). In particular, because high valuation consumers do not search in equilibrium, undercutting r would not be an optimal deviation. A high valuation consumer buys from a firm if and only if she starts out getting a quote from that firm, whereas a low valuation consumer buys from a firm if that firm charges v_ℓ and, if she starts off at that firm or the other firm is charging r . Search is random, so each firm initially gets half of the consumers.

Appendix C establishes that, if parameters are such that v_h is the monopoly price, as long as there is not too much discounting (i.e. δ is large enough), then there exists a unique

$\alpha^* \in (0, 1)$ that characterizes a mixed strategy equilibrium in prices. There are also two pure strategy equilibria: one where both firms charge the monopoly price v_h and one where they both charge v_ℓ . If the delay between the two periods is reduced so δ increases, r decreases whereas, for a given probability α , profit associated with the low price increases, because the value of selling to additional consumers in the second period is increased. As a result, the equilibrium involves a lower probability α^* that firms charge a low price. This setting thus provides a simple theoretical underpinning for the empirical findings by Ellison and Fisher Ellison (2014) that the price of rare used books sold online is higher than it is for the same items sold off line: interested buyers enjoy searching for such products so they are all shoppers but they enjoy getting hold of them as early as possible and, in an equilibrium where high prices and low prices coexist, the probability of finding a low price must be low in order for high valuation buyers to be willing to pay a high price while they expect to get frequent price quotes.⁵²

We next turn to search cost heterogeneity.

5.2 Search cost heterogeneity

In this subsection we concentrate on settings with a homogeneous product. Salop (1977) shows that a “noisy” monopolist can effectuate price discrimination by selling the same good at multiple outlets and, by setting different prices at each, discriminate against those with high search costs. As he says (p. 393) “dispersion acts as a costly device for sorting consumers into sub-markets to permit price discrimination.” An obvious first question is whether search cost heterogeneity can overcome the Diamond paradox for single-product firms. As we now argue, the answer is mixed.⁵³

A first crucial remark is that the Diamond paradox does not require that all consumers have the same search costs. All the arguments go through if there is a strictly positive lower bound on search costs. In order to obtain pricing below the monopoly price (and possibly some price dispersion), it is necessary that the search cost distribution has a fat enough tail in the neighborhood of zero. Some analysis of search market equilibrium with general search cost distributions can be found in Rob (1985) for monopolistic competition, and Stahl (1996) for oligopoly. It is quite challenging to get substantive results, especially for oligopoly. Still, a few takeaways from the analysis in Stahl (1996) are worth noting. First, it is possible to have a one-price equilibrium at a price below the monopoly price. This is

⁵²The empirical analysis in Ellison and Fisher-Ellison (2014) is structural and relies on a very different model, which is closer to that of Baye and Morgan (2001).

⁵³An early example is the tourists-and-natives model of Salop and Stiglitz (1977), which is closed with free-entry of firms facing U-shaped average costs. A fraction of firms are bargains, pricing at minimum average cost, and found by lucky tourists and all natives (shoppers). The rest are rip-offs, pricing at the reservation price, and serving those unlucky tourists who happen upon them first, and for whom search costs are too high to countenance seeking a bargain.

because, if the density of search costs at zero is positive, and if all firms are expected to charge some price p^* , then a firm's demand at p^* is very elastic for upward changes in price because all consumers with search costs near zero would want to switch to other sellers if the firm deviates upwards. Still the monopoly pricing equilibrium always exists as long as there is no atom at zero search cost. This equilibrium is unique if the density at zero is zero. Importantly, if the search cost distribution is atomless and its hazard rate increases over its support, then any symmetric Nash equilibrium is in pure strategies, implying that no consumer searches beyond the first firm encountered, expecting the same price at all firms. Hence (at least if we restrict attention to symmetric equilibria), there can be search in equilibrium only with some mass points in the search cost distribution (in particular at zero) or possibly if the distribution has a non-increasing hazard rate. Little is known about the properties of such equilibria with dispersed prices, except in the rather special case we discuss below in more detail.

A simple way to avoid a strictly positive lower bound on search costs is to assume as in subsection 2.1 that there is an atom of shoppers with zero search costs and the rest of the consumer population shares the same search cost $s > 0$. This simple modification of the basic framework of Varian (1980) is analyzed by Stahl (1989). The model is solved assuming that shoppers search through all the firms even if they expect the same price everywhere, and non-shoppers can observe a first price for free. This atom of shoppers implies an atomless equilibrium distribution of prices if we restrict attention to equilibria where all firms play the same strategy.⁵⁴ This is because, if there were an atom at \hat{p} in the price distribution, then a firm charging \hat{p} could lower its price slightly, thus increasing the probability that it captures the entire shopper population by a strictly positive amount and therefore earn a strictly larger profit than at \hat{p} , which contradicts the assumption that \hat{p} is played with a positive probability in equilibrium.

Stahl's original paper allowed for downward-sloping individual demands. However, if we assume unit demands for consumers, we can engage the results of the Varian (1980) model presented in subsection 2.1. In that model, there is an exogenous reservation price, r , and "captive" consumers are exogenously allocated symmetrically to firms. We can now take that framework and endogenously determine r as the price at which a consumer with a search cost, s , will be indifferent as to searching again, given the expected prices of firms: and we break indifference so that the consumer does *not* search. Then, the "captive" consumers of the Varian model are analogous to the costly search types who randomly select a particular firm to start at. Because the price distribution is atomless, if the highest price was strictly above r , a firm charging that price could not sell and so we know that, as in subsection 2.1 r is the largest price. Then, as explained in subsection 3.1 the reservation price is the sum of search cost and expected price.

⁵⁴Varian's (1980) model has multiple asymmetric equilibria too, with any number of two or more firms playing a mixed strategy, and the rest playing r (Baye, Kovenock, and de Vries, 1992, Kocas and Kiyak, 2006). This feature carries through to the Stahl-type analysis with endogenous r .

Drawing on Janssen, Moraga-González, and Wildenbeest (2005), we show in Appendix D that the equilibrium reservation price (for the model with unit consumer demand) is uniquely determined and is proportional to s , with the equilibrium distribution of prices given by (1) in subsection 2.1.⁵⁵ Janssen, Moraga-González, and Wildenbeest (2005) show that the equilibrium reservation price here unambiguously rises with n , which reflects the property that the price per firm rises with n in the sense of first-order stochastic dominance. We also show in Appendix D the striking result that the expected minimum price is independent of n (Janssen, Pichler, and Weidenholzer, 2011). Even though consumers face higher prices at each individual firm, the greater number of options exactly cancels this out. The implications of more competition are therefore different for different consumer groups (Stahl, 1989, also pointed out tensions in his setting with price sensitive demand). The shoppers are unaffected, but the others are strictly worse off (because the first search is expected to find a higher price to start with, and no search ensues). These results are usefully compared to those for Varian’s (1980) model, in which the reservation price is fixed exogenously. As we pointed out earlier, the shoppers are better off there because the minimum price falls (while the others are worse off, as here). The difference in results comes from the endogenous reservation price, which rises with n in the search context and so brings up the minimum price.

This model has become a very useful workhorse for analyzing search equilibrium in markets for homogeneous products. In particular, it does capture the increased competitiveness in a market resulting from a lower search cost because r is decreasing in s so a decrease in search costs induces a new price distribution which is stochastically dominated in the first order by the old one. Although this decrease in prices arises because search becomes more attractive, it does not result from consumers having better information about the competing prices: in equilibrium, non-shoppers know only one price and shoppers know all prices independent of the level of the search cost. Thus, this setting fails to deliver usable predictions about the evolution of consumer search behavior and hence of their information from improved information technologies.⁵⁶

We next consider heterogeneity of search costs when consumers search for a good match.

5.3 Intensive and extensive margins

We now allow for search cost heterogeneity to see how equilibrium prices react to lower search costs. From the analysis in section 3, we expect consumers to search more when

⁵⁵Specifically, they show that $r = \frac{s}{1-\alpha}$, where $\alpha = \int_0^1 \frac{1}{\frac{\sigma}{m-\sigma}ny^{n-1}+1} dy$.

⁵⁶Bénabou (1993) argues that this unappealing feature can be overcome by introducing marginal cost heterogeneity as in Reinganum (1979). Low search cost consumers keep on searching until they find a firm with a low enough cost. However, as in Reinganum’s setting, equilibrium prices for low cost firms are monopoly prices and bringing the search cost distribution down to zero cannot bring prices below the monopoly price of the most efficient firm.

search costs less, inducing more competition and hence lower prices. This, however, overlooks the impact of lower search costs on a consumer’s decision to search actively. If the search costs go down for consumers with high search costs to induce them to become active in new markets, the search cost distribution of participating consumers in these markets can rise to induce prices to rise. Janssen *et al.* (2005) illustrate this point within the framework of Stahl (1989), described in subsection 5.2. They show that if the search cost of non-shoppers is large enough, they are indifferent between participating in the market or not. Their equilibrium participation probability falls as their search costs decrease,⁵⁷ so that the share of shoppers in the market decreases resulting in higher prices. We now discuss, following Moraga-González *et al.* (2016), how this idea can be fruitfully generalized within a search framework with horizontal differentiation *à la* Wolinsky (1986).

Consider again the monopolistic competition setting of subsection 3.2, which is now extended to include a continuum of consumers with search costs distributed on $[\underline{s}, \bar{s}]$, $\underline{s} > 0$. Heterogeneity in search costs impacts the firm’s demand in a non-trivial way. For each search cost level, s , there is a different threshold match $\hat{x}(s)$ that determines the consumer’s demand. Thus a firm’s demand integrates the probability of a purchase over all values of $\hat{x}(s)$ so the standard existence argument based on an increasing hazard rate for the match distribution does not work here. However, as Moraga-González *et al.* show, this consumer heterogeneity can be dealt with using a fairly standard argument. The intuition for this is as follows. Recall from subsection 3.2 that under monopolistic competition, each firm only competes with the option of searching on, as represented by the threshold $\hat{x}(s)$. This is just as if each firm faced a single competitor for whose product the consumer’s match is $\hat{x}(s)$. Hence the arguments from Caplin and Nalebuff (1991) showing equilibrium existence in an oligopoly model where matches have log-concave densities can be applied here to show existence as long as $\hat{x}(s)$ has a log-concave density.⁵⁸

To build some intuition about the comparative statics regarding changes in the search cost distribution, consider a simple example with only three search cost levels, which are initially $s_{\ell 1}$, $s_{m 1}$ and $s_{h 1}$ and change to $s_{\ell 2}$, $s_{m 2}$ and $s_{h 2}$. The fraction of consumers at each level are first $\pi_{\ell 1}$, $\pi_{m 1}$, $\pi_{h 1}$, and then $\pi_{\ell 2}$, $\pi_{m 2}$, and $\pi_{h 2}$, for low, intermediate, and high search costs respectively. The new search costs are stochastically lower than the original ones (in terms of first order stochastic dominance). Furthermore, assume that only consumers with search costs of at most $s_{m 1}$ engage in search initially. Because consumers of types ℓ and m have lower search costs after the change, they search more and this *per se* induces downward pressure on prices. This is what Moraga-González *et al.* call the *intensive margin*. If type h consumers remain out of the market after the drop in search costs, then only the intensive margin is in play, and the price goes down. If, however, search costs for some type h consumers falls by enough that they search at the new equilibrium, then the composition

⁵⁷This property is rather a vagary of the mixed strategy participation indifference condition.

⁵⁸Moraga-González *et al.* do not however provide general joint conditions on the search cost distribution and the match distribution guaranteeing that $\hat{x}(s)$ has a log-concave density.

of the participating consumers changes in a way that might upset the stochastic decrease in search costs for active searchers. This change in the *extensive margin* might pull prices upwards.

Which margin (extensive or intensive) dominates depends on how the search cost distribution changes. To illustrate, suppose the only change is a drop of the high search cost (as in the Janssen *et al.* example discussed above), so that s_{h2} is now low enough to induce search by the high types. Then search costs for the participating consumers unambiguously increase in a stochastic sense, so that the extensive margin dominates, leading to an increase in price. Alternatively, suppose instead that search values do not change but consumer fractions do. If some of the high types shift to the middle type with no other change (so $\pi_{m2} > \pi_{m1}$ and $\pi_{\ell 2} = \pi_{\ell 1}$) then the new distribution for participating consumers is stochastically higher. Once again, the impact of the extensive margin dominates. In order for the intensive margin to be the dominant factor it is necessary that there is a sufficient shift of intermediate search cost consumers to low search cost so that $\frac{\pi_{\ell 2}}{\pi_{m2}} > \frac{\pi_{\ell 1}}{\pi_{m1}}$ (which ensures that the proportion of low search cost consumers among searching consumers has increased). Thus, the intensive margin will dominate if the shift of consumers towards the lower values of search costs dominates the shift in favor of higher search cost values. This can be captured by comparing the two distributions in terms of a monotone likelihood ratio property. Moraga-González *et al.* show that, if the likelihood ratio between the final distribution and the initial distribution⁵⁹ is increasing, then the stochastic drop in search costs will induce a price increase, whereas if it is decreasing, then lower search costs will result in a price fall.⁶⁰

As noted by Moraga-González *et al.*, the study by Hortaçsu and Syverson (2004) on the US mutual fund market in the late 1990s provides an empirical illustration of how the extensive margin may lead to higher prices in a market where search costs fall due to the development of the internet. They find that prices increased and, according to their estimates, search costs dropped at the lower percentile of the distribution but actually rose at the upper percentiles. Hortaçsu and Syverson explain this by the arrival of new high search cost households in the mutual fund market, whose search costs went down so much that they found it worth looking for investment opportunities.

5.4 Ordered search

We treat here ordered search when buyers have heterogeneous search costs. Let all consumers have unit demands for a homogeneous good with common valuation v , assumed to be large. Firms are searched in an exogenous order, the same for all consumers (think geography, for example, with all consumers starting at the same firm on a street, and getting further away from the start point with successive firms searched). What we shall show is

⁵⁹This is defined as the ratio of the two densities.

⁶⁰Their condition implies this in the case of a stochastic decrease in search costs.

that consumers with lower search costs will search longer (i.e., further). Notice before we start that the only way consumers will want to keep searching is if they expect lower prices later. This means that, as per the geography example, consumers cannot choose a later firm without going past an earlier one – if instead a consumer could reach any firm with an identical cost of search, she would want to skip immediately to the end.⁶¹ The analysis that follows is based on Arbatskaya (2007): we specialize the search cost distribution to a uniform one, with $s \in [\underline{s}, \bar{s}]$, and we set $\bar{s} = \underline{s} + 1$.

To illustrate, suppose there are three firms with zero production costs and they are searched in order 1 through 3. Consumers who observe p_2 at firm 2 will go through to firm 3 if they expect a price discount there greater than their search cost. That is, letting p_i^* be firm i 's equilibrium price, they continue if $p_2 - p_3^* > s$. Now consider a consumer observing p_1 at firm 1. She will search on to firm 2 while expecting to buy from it at the equilibrium price p_2^* if $s < p_1 - p_2^*$. Hence, in order for firm 2 to be active in equilibrium, prices must satisfy $p_1^* - p_2^* > p_2^* - p_3^*$: else, any consumer who would move on to firm 2 expecting to buy at firm 2 would end up preferring buying from firm 3. The argument easily generalizes to n firms and Arbatskaya (2007) shows that price differences decrease for all active firms (which must be the earliest firms in the sequence). Thus it is necessarily the consumers with the highest search costs that stop and buy at a given firm. It also means that, in equilibrium, the search decision is myopic in the sense that a consumer moves on to the next firm if and only if it would be optimal to do so even if the next firm was the last. In general, if consumers do not search in the optimal order characterized by Weitzman (1979) as is the case here, optimal search is not necessarily myopic. Myopia here is an equilibrium property in the sense that it would not be optimal to search myopically if price differences were not decreasing.

Further note that firm 3 cannot be active in equilibrium. If it expected any visit, it would just undercut firm 2's price and could thus capture all visiting consumers: then it would not be optimal for any consumer to search on all the way to firm 3 because no price discount can be expected there. Consider now firm 2's problem assuming that all consumers with search costs in excess of $\underline{s} + \Delta$ with $\Delta > 0$ stop at firm 1. Firm 2 cannot affect this decision to stop at firm 1, which depends on its expected price rather than on its actual price.⁶² Because $s \geq \underline{s}$, firm 2's profit is $p_2(\underline{s} + \Delta - (p_2 - p_3^*))$ for $(p_2 - p_3^*) \geq \underline{s}$ and $p_2\Delta$ otherwise. The solution must be to set $p_2^* = p_3^* + \underline{s}$ so that further search is deterred. Hence the price elasticity of demand at that price for a price increase must be at most -1, that is, $p_3^* + \underline{s} \geq \Delta$. The lowest expected equilibrium price for firm 3 such that firm 2 chooses to

⁶¹Notice that the Stahl model with ordered search has the same property: it can be readily shown for duopoly that the second firm has a lower equilibrium expected price than the first one: if it were possible for consumers to choose which firm to sample first, they would choose the second one, upsetting the equilibrium pricing.

⁶²Arbatskaya (2007) also analyzes the case when all prices are observed (so that a firm influences both margins). There is some literature on observed price deviations, notably Carlson and McAfee (1982).

deter further search is $p_3^* = \Delta - \underline{s}$, and the associated profit-maximizing price for firm 2 is $p_2^* = \Delta$. Any $p_3^* \geq 1 - \underline{s}$ constitutes an equilibrium, with $p_2 = p_3^* + \underline{s}$ so that there is a continuum of equilibria. Following Arbatskaya (2007), we choose the lowest possible value for p_3^* consistent with no search at the final firm.⁶³ The last firm, though inactive, performs a policing role on its predecessors.

Because we seek an equilibrium where firm 2 is active, firm 1 must be setting a price above the kink in its demand corresponding to the price at which all consumer would give up search. This means that it selects a price that yields a unit price elasticity. Its demand being $1 - \Delta$, its equilibrium price is $p_1^* = 1 - \Delta$. Hence the search cost such that a consumer is indifferent between buying from firm 1 and searching is $p_1^* - p_2^* = 1 - 2\Delta$ and, because firm 2 sells to all consumers with search costs below that threshold we have $\Delta = 1 - 2\Delta - \underline{s}$ which yields

$$(12) \quad \Delta = \frac{1 - \underline{s}}{3}.$$

It follows that firm 2 can be active only if $\underline{s} < 1$. Note further that if \underline{s} is more than $\frac{1}{4}$, $p_3^* = \Delta - \underline{s} < 0$, which is optimal because no consumer visits firm 3, although it is a weakly dominated strategy. Interestingly, a drop in the minimum search cost has diverging impacts on the two prices: it increases the price of the second firm and decreases that of the first firm.

The insight that the maximum number of active firms in equilibrium decreases if the minimum search cost increases extends to the n firm case. An equilibrium where $n - 1$ out of n firms are active is characterized by $n - 1$ market shares $\Delta_1, \dots, \Delta_{n-1}$ with $\Delta_i > 0$ for all i . Because pricing by firms 1 through $n - 1$ must set the price elasticity of demand to -1 , we have $p_i^* = \Delta_i$. These market shares constitute an equilibrium if and only if they satisfy $\sum_{i=1}^{n-1} \Delta_i = 1$ and $\Delta_i = \underline{s} + \sum_{j=i+2}^{n-1} \Delta_j + 2\Delta_{i+1}$ for $i = 1, \dots, n - 2$. The second equality merely means that consumers who buy from firm i are those who do not buy from earlier firms and have a search cost larger than $p_i^* - p_{i+1}^* = \Delta_i - \Delta_{i+1}$. It also guarantees the equilibrium property that price differences should decrease and should be equal to some threshold search cost lying above \underline{s} . It is readily seen that the larger is n is the lower \underline{s} should be in order for the two constraints to hold simultaneously. Furthermore, for $\underline{s} = 0$, it is always possible to specify a sequence of market shares satisfying the two constraints for any number of active firms: the second constraint allows for deriving all market shares from the market share of the last active firm Δ_{n-1} so that, taking it appropriately small ensures that the sum of market shares adds up to 1.

What is in evidence here for $\underline{s} > 0$ is a “natural oligopoly” (or finiteness) property of the ordered search model. Natural oligopolies were earlier described by Shaked and Sutton (1983) for vertical differentiation models. Even as fixed costs tend to zero, only a finite

⁶³Alternatively, its price might be determined by local demand conditions. Or, indeed, one might include some shoppers too, and the firms far back would only encounter them, and so they would set price equal to marginal cost.

number of firms might survive in equilibrium because lower qualities cannot turn a profit due to their disadvantage. So it is here that firms too far down the search order cannot survive profitably given the pricing behavior of their predecessors.⁶⁴

We now return to the ranking of sellers with heterogeneous purchase probabilities and illustrate how it can be interacted with search cost heterogeneity. As discussed in subsection 4.4, Athey and Ellison (2011) consider a setting where all firms charge the same exogenous price and differ in the probability that their product matches a consumer's need. They study how the search order can emerge as the outcome of an auction where firms bid for positions on a web site. Purchase probability is private information to the product's seller. As was pointed out in subsection 4.4, this assumption implies that higher probability firms have a higher incremental value for being ahead in the search order. Consumers should indeed expect them to bid more and hence, to be placed higher on the web page: so the proposed order is the optimal search order. Athey and Ellison introduce heterogeneity in search costs, which also induces a positive relation between purchase probability and seller willingness to pay for being positioned early. To see this, consider again the two firm example with an exogenous price where firms know each other's probability and assume now that there is only a fraction $\alpha \in (0, 1)$ of consumers whose search cost is low enough so that they want to check out the other product if they are not satisfied with the first product sampled. Then firm 1's incremental value for being ahead of firm 2 is $(1 - \alpha)\beta_1 + \alpha\beta_1\beta_2$, which is indeed larger than the incremental value of firm 2 for being in front of firm 1 if $\beta_1 > \beta_2$.

In addition, and maybe more importantly, heterogeneity in search costs also yields a higher total industry profit if firms are ordered according to a decreasing probability ranking whereas the search order would not impact total profit if it were the case that all consumers searched on until finding a satisfactory product. Total profit is given by $\beta_1 + \alpha(1 - \beta_1)\beta_2$ with firm 1 searched first: for $\beta_1 > \beta_2$, this is strictly larger than with firm 2 in the top position if and only if $\alpha < 1$. Obviously, an ordering of firms with higher purchase probabilities earlier is also what is preferred by consumers, because it minimizes expected search costs (and it must arise in this fixed price setting, so consumer search behavior is consistent with the firms' bidding for prominence on the web page). The analysis of Athey and Ellison (2011) therefore yields an alignment of preferences between firms and consumers regarding the search order. By contrast, the results in Anderson and Renault (2016) where prices are endogenous point rather to a systematic misalignment. In particular, if firms only differ by the probability of a match with the consumer's tastes, total profit maximization requires that firms with a *lower* match probability should get more prominence because early positions are priced lower. However, consumers prefer to reach high match probabilities early. Not only does this reduce total search costs, but early products are also sold at lower prices.

⁶⁴The result for $\underline{s} = 0$ parallels what obtains in the vertical differentiation setting for a standard uniform distribution of valuations for quality.

6 Conclusions

We began this Chapter by noting that research on search costs and firm pricing has really taken off with the advent of the internet. Search costs are very visible in searching for goods and services online. And yet, search costs were prevalent before the internet. It might be argued that search costs were responsible for rather limited access to variety before the internet age, and a concurrent stifling of price competition. These costs were hithertofore quite hidden. Perhaps ironically, it is only when search costs have fallen that economists (Nobel prize-winners aside!) became much more aware of them (one exception is the analysis of job search, although the equilibrium analysis initiated by Diamond, 1982, only picked up in the 1990s). Along with the opening of markets has come a supply response that has greatly increased the variety of goods and services that are readily accessed.⁶⁵

We have surveyed various models of market pricing with search costs. We have noted the great achievements (much credit should be ascribed to Weitzman, 1979, for these) in the description of individual search behavior. But closing the loop and endogenizing firm pricing and realistically describing market equilibrium has proved a bit more cumbersome, though progress continues apace. Models differ by the degree of heterogeneity they assume about individuals, and firms. For simplicity below, we divide the main approaches into the product differentiation approach following Wolinsky (1986) and the heterogeneous search costs approach involving mixed strategies for prices, following Stahl (1989) and building on Varian (1980).

Models can be judged according to various criteria. One is their success in matching empirical regularities. These include price distributions, and consumer search behavior. The Stahl model succeeds in generating a price distribution (which is a truncated Pareto for the inelastic demand case), although critics of mixed price strategy equilibrium are quite common. However, the model does not generate search in equilibrium.⁶⁶ The Wolinsky model, by contrast, involves pure price strategies (arguably a strong point per se). Under fully random search, price dispersion would have to come from asymmetries in firms' costs or qualities (which are a bit cumbersome to manage in the model), but ordered search does deliver dispersion. On the consumer side, different consumers have different search patterns, with lucky ones finding acceptable matches earlier, while others go to the end and return.⁶⁷

⁶⁵See Goldmanis et al. (2010) for empirical evidence on how travel agencies and bookstores have adapted.

⁶⁶The model starkly depicts two groups of consumers. One group does not search at all in equilibrium, while the other already knows all prices. Semantically, it could be argued that the latter group does (without cost) search through all products.

⁶⁷This is less likely the bigger the number of options. All comeback consumers reach the end: extra heterogeneity of firms and/or consumers would be needed to get some consumers to return to an earlier option before reaching the end (as can be envisaged from the Weitzman, 1979,

We can also look to the models to pull together some results about the impact of reduced search costs. In the Wolinsky (1986) model, a lower search cost raises the search threshold, so consumers become more choosy. They search longer, on average. The *total* search cost falls though under monopolistic competition for log-concave match distributions. In this case market prices too fall as consumers search longer. When the number of active consumers is fixed, the price drop drives profits to fall too. This reduces total product variety in the long run so globalization can streamline (and bankrupt) local market offerings.⁶⁸

However, there is another important channel through which lower search costs might raise total variety. This is the two-sided market effect that we just suppressed by taking the number of consumers as fixed. Because lower search costs increase expected surplus from market participation, more consumers will want to look, so delivering a virtuous circle (i.e., bilateral positive externalities from participating) that induces more firms to want to partake too; thus, lower search costs generate a thickening of the market all around.⁶⁹

The two-sided market effects just discussed are one example of the externalities inherent to search markets. For another example, in the Stahl model there are externalities between the consumer groups. All consumers are better off the more informed consumers there are, as price competition for the informed is fiercer. But, the more consumers there are whose search is costly, the worse off are all. Anderson and Renault (2000) consider a variant of the Wolinsky model in which some consumers know their match values already, while others need to search to find them. The former immediately check out the firm that they like most, which means that their demand is more inelastic than for the other consumers. The more of them there are, the higher are equilibrium prices. These informed types impose a negative externality on all (and the uninformed imbue a positive externality). Introducing a private cost to getting informed (which differs across individuals), there is thus a negative externality from investing in information acquisition, and so *too much* information is gathered on search for product matches.

Internet search is inherently sequential, and frequently ordered to boot. Models of equilibrium pricing for such markets, with concurrent competition for positions in the search order, are only just being developed. Relatedly, the internet also greatly facilitates access to price information (probably more so than to information about product attributes).

framework).

⁶⁸Some of these contrasting effects are also at play in the (suitably extended) Stahl (1989) model. A lower search cost (for those with positive search costs) disciplines firms more and intensifies price competition (even though the actual search is unchanged, it is the threat of search that drives prices down), although lower costs per search do not impact total search costs because there is no search in equilibrium. The lower price level improves the welfare for both consumer types. Firms though suffer lower profits.

⁶⁹This can hold too in an extended version of the Stahl (1989) model with endogenous consumer participation. With lower search costs, profits could rise in the short run and more firms enter in the long run.

Consumers can thus easily compare prices and firms can use pricing to direct search, with consumers being enticed to start search with the cheapest. How this affects price competition is an exciting and challenging research question: Armstrong (2016) discusses some of the recent contributions on this topic.⁷⁰ As internet commerce develops into more different variations and opportunities, the economics of search will develop along with it. It promises to be an exciting research trajectory.

Appendix A: Price Dispersion

A1: Model with exogenous population of captive consumers.

Let $\pi(p)$ denote a firm's expected profit when it charges a price p and all other $n - 1$ firms follow the symmetric price distribution $F(p)$. Either all other firms price above p and the firm serves all shoppers and its captives, or else at least one sets a lower price and the firm just serves its captives. Hence

$$\pi(p) = p \{ (\gamma + \sigma) (1 - F(p))^{n-1} + \gamma (1 - (1 - F(p))^{n-1}) \}.$$

Setting this equal to the profit earned by pricing at r , $\pi(r) = \gamma r$, and rearranging yields the equilibrium price distribution as given in the text.

A2: Simultaneous search model

Recall that ν_1 is the probability a consumer gets one quote (so that each of the n firms has an equal chance), and ν_2 is the complementary probability that she gets two (and then chooses the one with the lower price). Hence

$$\pi(p) = p \left(\nu_1 \frac{m}{n} + 2\nu_2 \frac{m}{n} (1 - F(p)) \right).$$

Standard arguments show that r is in the price support and the support has no atoms. Setting $\pi(p)$ equal to the profit earned by pricing at r , $\pi(r) = r\nu_1 \frac{m}{n}$, and rearranging yields the equilibrium price distribution as given in the text. As per the text, the equilibrium values of ν_1 and ν_2 are determined by the condition that the consumer be indifferent between getting one or two quotes.

⁷⁰Price posting is also a key ingredient of the directed search models initiated by Peters (1984) and widely used in the analysis of the labor market (although Peters' original model was applied to a product market). The other key ingredient is that firms have a limited capacity, in contrast to the settings we have discussed where capacity is unlimited and marginal production costs are constant.

A3: Price advertising model

We characterize an equilibrium where a firm that does not advertise charges r and advertised prices have support $[\underline{p}, r]$ where we take the convention that a consumer chooses the advertised price in case of a tie with an unadvertised one. There will be no advertising if $\frac{2A}{m} \geq r$, so assume this condition does not hold. If advertising, a firm gets profit

$$\pi(p) = pm(1 - F(p)) - A,$$

where we define $F(p)$ as the probability that a price is advertised at or below p , and so $F(r)$ is the probability of advertising. If a firm does not advertise, it gets $\frac{r}{2}m(1 - F(r))$ where $(1 - F(r))$ is the probability the other firm does not advertise (and hence the symmetric equilibrium probability). From above, $\pi(r) = rm(1 - F(r)) - A$, which we equate to the not-advertising profit to find

$$1 - F(r) = \frac{2A}{rm},$$

and the equilibrium profit is consequently A . The equilibrium probability of an advertised price at or below p is then

$$F(p) = 1 - \frac{2A}{mp}$$

for $p \in [\frac{2A}{m}, r]$.

Appendix B: Horizontal differentiation and search in monopolistic competition.

B1: Profit function quasiconcavity and symmetric equilibrium

Given the demand function from the text, we have firm i 's profit as

$$\pi_i(p_i) = (p_i - c)m \frac{1 - F\left(\hat{x} + \frac{p_i - p^*}{\mu}\right)}{1 - F(\hat{x})},$$

which is positive for $p_i \in (c, \mu(b - \hat{x}) + p^*)$, and, since any candidate p^* exceeds c , this interval is non-empty. First note that if $c < \mu(a - \hat{x}) + p^*$ then profit is linearly increasing in p_i for $p_i \in (c, \mu(a - \hat{x}) + p^*)$ because all consumers reaching firm i stop at it ($F\left(\hat{x} + \frac{p_i - p^*}{\mu}\right) = 0$). Over the interval $p_i \in (\max\{c, \mu(a - \hat{x}) + p^*\}, \mu(b - \hat{x}) + p^*)$, we can write

$$\frac{d\pi_i(p_i)}{dp_i} = \frac{mf\left(\hat{x} + \frac{p_i - p^*}{\mu}\right)}{1 - F(\hat{x})} \left\{ \frac{1 - F\left(\hat{x} + \frac{p_i - p^*}{\mu}\right)}{f\left(\hat{x} + \frac{p_i - p^*}{\mu}\right)} - (p_i - c) \right\}.$$

The term outside the parentheses is positive as long as demand is positive, so consider the terms inside. The first (ratio) term is positive, and decreasing in p_i from the increasing hazard rate property (log-concavity of $1 - F(\cdot)$), from which we subtract the mark-up term, $(p_i - c)$, which is zero at $p_i = c$ and linearly increasing in p_i . Hence, the profit derivative is at first positive then it is negative. Profit is therefore quasi-concave. The profit-maximizing best response to a candidate $p^* > c$ is therefore either at $\mu(a - \hat{x}) + p^*$ or else where the term in parentheses is zero. The former cannot constitute a symmetric equilibrium, so the symmetric equilibrium is where the term in parentheses is zero at $p_i = p^*$, which is the expression given in the text.

B2: Behavior of equilibrium price for μ low enough

First recall from (6) that $\mu \int_{\hat{x}}^b (\epsilon_i - \hat{x}) f(\epsilon_i) d\epsilon_i = s$ defines the match-value stopping-rule, \hat{x} . Then the value $\underline{\mu} = \frac{s}{\int_a^b (\epsilon_i - a) f(\epsilon_i) d\epsilon_i} > 0$ is the value of μ that just entails the consumer stopping immediately, even at the worst possible match, a . So, for $\mu < \underline{\mu}$ the consumer stops immediately. The equilibrium price (from the text) is $p^* = c + \frac{\mu[1-F(\hat{x})]}{f(\hat{x})}$, which is finite as long as $\hat{x} > a$. However, as $\mu \downarrow \underline{\mu}$, this price tends to $c + \frac{\mu}{f(a)}$, which tends to infinity for $f(a) = 0$, which implies p^* must be decreasing over some non zero measure set in the neighborhood of $\underline{\mu}$.

B3: Total search costs

We wish to determine whether or not total search costs spent increase or decrease when s falls. For example, is more time spent searching? The equilibrium chance of a consumer stopping (and therefore buying) at any firm, conditional on reaching it, is $1 - F(\hat{x})$. The expected number of searches is then⁷¹

$$E = \frac{1}{1 - F(\hat{x})},$$

so that the equilibrium total search cost is

$$sE = \frac{\mu \int_{\hat{x}} (1 - F(x)) dx}{1 - F(\hat{x})}.$$

Assuming that $1 - F$ is log-concave, then the integral here is also log-concave (by the inheritance property of log-concave functions under integration: see e.g. Caplin and Nalebuff,

⁷¹The expected number of searches is (dropping temporarily the dependence of F on \hat{x}) $E = (1 - F) + 2(1 - F)F + 3(1 - F)F^2 + \dots$. The series $1 + 2F + 3F^2 + \dots$ is the sum of the series $S = 1 + F + F^2 + \dots$ plus F times the series S plus F^2S etc., i.e. $\frac{S}{1-F}$, and likewise $S = \frac{1}{1-F}$, so that $E = \frac{1}{1-F}$.

1991). Then the ratio $\frac{-(1-F(\hat{x}))}{\int_{\hat{x}}(1-F(x))dx}$ should be decreasing in \hat{x} , so that sE is decreasing in \hat{x} . Then, note from the search rule that⁷²

$$\frac{d\hat{x}}{ds} = \frac{-1}{\mu(1-F(\hat{x}))}$$

so that sE increases in s .

B4: Equilibrium analysis under asymmetric costs

Firm i 's demand is proportional to $1 - F(p_i + \hat{u})$, so firm i 's profit is proportional to

$$\pi_i(p_i) = (p_i - c_i)(1 - F(p_i + \hat{u})).$$

The argument in the first sub-section of this Appendix applies, *mutatis mutandis*, to show that π_i is quasi-concave in p_i , so that an interior solution (i.e., with $p^* \in (a - \hat{u}, b - \hat{u})$) solves the first-order condition

$$p^*(c_i, \hat{u}) = c_i + \frac{(1 - F(p^*(c_i, \hat{u}) + \hat{u}))}{f(p^*(c_i, \hat{u}) + \hat{u})}.$$

The increasing hazard rate property (that demand, $1 - F(\cdot)$, is strictly log-concave) implies that $k = -\left[\frac{(1-F(x))}{f(x)}\right]' > 0$. Hence, using the implicit function theorem that the cost pass-through rate, $\frac{dp^*(c, \hat{u})}{dc} = \frac{1}{1+k} \in (0, 1)$. Similarly, $\frac{dp^*(c, \hat{u})}{d\hat{u}} = \frac{-k}{1+k} \in (-1, 0)$, as claimed in the text.

Finally, briefly consider non-interior solutions, and note the profit derivative (from the quasi-concave profit function) is

$$\frac{d\pi_i(p_i)}{dp_i} = -(p_i - c_i)f(p_i + \hat{u}) + (1 - F(p_i + \hat{u}))$$

A firm will choose to sell to all comers if this is negative where $p_i + \hat{u} = a$: rearranging, the condition for this to happen is

$$c_i \leq a - \hat{u} - \frac{1}{f(a)}.$$

On the other side, a firm will not want to sell at all if the profit derivative is positive where $p_i + \hat{u} = b$; rearranging, this holds if $c_i \geq b - \hat{u}$. In this case the firm cannot capture even its most eager customer by pricing at marginal cost.

⁷²Notice too that the change in the expected number of searches wrt s is $\frac{dE}{ds} = \frac{-f(\hat{x})}{(1-F(\hat{x}))^3} < 0$. The “demand” for search therefore slopes down.

Appendix C: when search costs are waiting costs or search with discounting.

Given the equilibrium strategy played by its competitor, a firm is indifferent between the two prices if

$$(13) \quad \frac{m_h}{2} \left(\frac{v_h - \delta(v_h - \alpha v_\ell)}{1 - \delta(1 - \alpha)} \right) = \frac{\alpha}{2} v_\ell + (1 - \alpha) ((1 + \delta)m_\ell + m_h) \frac{v_\ell}{2}.$$

where firms are assumed to have the same discount rate as consumers.

The left hand side is strictly decreasing and convex in α : it is equal to $\frac{m_h v_h}{2}$ for $\alpha = 0$ and $\frac{m_h}{2}(v_h - \delta(v_h - v_\ell))$ for $\alpha = 1$. The right hand side is linear and decreasing in α from $((1 + \delta)m_\ell + m_h) \frac{v_\ell}{2}$ for $\alpha = 0$ to $\frac{v_\ell}{2}$ for $\alpha = 1$. Because v_h is the monopoly price we have $\frac{m_h v_h}{2} > \frac{v_\ell}{2}$. Assume m_ℓ is small enough so that $\frac{m_h v_h}{2} > (m_\ell + \frac{m_h}{2}) v_\ell$, and, in addition, that δ is close enough to 1 so that $\frac{m_h}{2}(v_h - \delta(v_h - v_\ell)) < \frac{v_\ell}{2}$. Then there are two symmetric pure strategy equilibria where price is v_h and v_ℓ respectively. In addition, there exists a value of α , $\alpha^* \in (0, 1)$ that satisfies (13). Furthermore, there can be only one solution: if there were a second solution, then the left-hand side of (13) would cross the right-hand side from below and, since the former is convex and the latter is linear, the right-hand side profit would have to be larger than the left-hand side profit at $\alpha = 1$, which is not the case for δ large enough.

Appendix D: search equilibrium with some shoppers.

We follow the exposition of Stahl's model (specialized to unit demand) in Janssen, Moraga-González, and Wildenbeest (2005). These authors decompose the analysis into two constituent parts. On the firm side, the equilibrium relation for the price distribution is given from the Varian model (see Appendix A). The reservation price, r , is endogenously determined from the condition that the consumers with search costs stop at first search. That is, r is determined from the condition that the expected benefit from further search be just equal to the search cost, s (so that no such consumer searches again):

$$r - \int_{\underline{p}}^r p f(p) dp = s$$

with $\underline{p} = \frac{\gamma}{\gamma + \sigma} r$ (as per the Varian model). Here the left-hand side is the expected price saved from a further search, because the integral expression is the expected price at any randomly chosen firm and can be rewritten as $\int_{\underline{p}}^r p f(p) dp = \int_0^1 p dF(p)$.

As we showed in Appendix A, the Varian model gives $(1 - F(p)) = \left(\frac{\gamma}{\sigma} \left(\frac{r}{p} - 1 \right) \right)^{\frac{1}{n-1}}$ with $\gamma = \frac{m - \sigma}{n}$ (below we break out n when pertinent in order to discuss the effects of

changing the number of firms). We now show that there is a unique r solving the model, and give its closed form solution. Indeed, set $y = 1 - F(p)$ to invert the equilibrium distribution characterization as

$$p = \frac{r}{\frac{\sigma}{\gamma}y^{n-1} + 1}$$

which we can then insert in the expected price expression to write

$$(14) \quad \int_0^1 pdF(p) = r \int_0^1 \frac{1}{\frac{\sigma}{\gamma}y^{n-1} + 1} dy,$$

and hence we have an explicit solution for the reservation price, as given in the text, namely $r = \frac{s}{1-\alpha}$ (with $\gamma = \frac{m-\sigma}{n}$ and $\alpha = \int_0^1 \frac{1}{\frac{\sigma}{\gamma}y^{n-1} + 1} dy$).

The insight of Janssen, Pichler, and Weidenholzer (2011) to find the expected minimum price, Ep_{\min} , in the market follows from the indifference property of the mixed strategy equilibrium. Charging the reservation price nets a firm $\frac{r(m-\sigma)}{n}$, which must equal its overall expected profit. The latter has two sources, the expected price on its share of consumers with positive search costs (i.e., $\alpha r \frac{m-\sigma}{n}$), plus its chance ($1/n$) of getting the s shoppers at the expected minimum price. Pulling that together yields the indifference condition as

$$r(m-\sigma) = \alpha r(m-\sigma) + \sigma Ep_{\min},$$

which rearranges to $Ep_{\min} = \frac{r(m-\sigma)(1-\alpha)}{\sigma} = \frac{(m-\sigma)s}{\sigma}$, where the last step recalls the reservation price property that $r = \frac{s}{1-\alpha}$. The expected minimum is clearly independent of n .

Finally, it is useful to draw out the simpler case for duopoly. From the analysis above, with $n = 2$, we have (by integrating) $\alpha = \frac{m-\sigma}{2\sigma} \ln\left(\frac{m+\sigma}{m-\sigma}\right)$, $r = \frac{s}{1-\alpha}$, and $(1 - F(p)) = \frac{\gamma}{\sigma} \left(\frac{r}{p} - 1\right)$, with $\underline{p} = r \frac{m+\sigma}{m-\sigma}$. The distribution equation is already in the Varian model analysis, so let us simply derive the expected price from a search, which is the left-hand side of (14). Noting from the Varian distribution result that $f(p) = \frac{\gamma}{\sigma} \frac{r}{p^2}$, we have $\int_0^1 pdF(p) = \int_{\underline{p}}^r pf(p) dp = r \frac{m-\sigma}{2\sigma} \ln\left(\frac{m+\sigma}{m-\sigma}\right)$, and hence the results above are readily verified.

References

Anderson, S. P., Baik, A., Larson, N. (2016) Price Discrimination in the Information Age: List Prices, Poaching, and Retention with Personalized Discounts. Mimeo.

Anderson, S. P., de Palma, A., Nesterov, Y. (1995). Oligopolistic Competition and the Optimal Provision of Products. *Econometrica*, 63(6), 1281-1301.

Anderson, S. P., de Palma, A., Thisse, J. F. (1992). Discrete choice theory of product differentiation. MIT Press, Cambridge, MA.

- Anderson, S. P., Renault, R. (1999). Pricing, product diversity, and search costs: a Bertrand-Chamberlin-Diamond model. *The RAND Journal of Economics*, 30(4), 719-735.
- Anderson, S. P., Renault, R., (2000). Consumer Information and Firm Pricing: Negative Externalities from Improved Information. *International Economic Review*, 41(3), 721-42.
- Anderson, S. P., Renault, R., (2003). Efficiency and surplus bounds in Cournot competition. *Journal of Economic Theory*, 113(2), 253-264.
- Anderson, S. P., Renault, R. (2006). Advertising content. *American Economic Review* 96(1), 93-113.
- Anderson, S. P., Renault, R. (2013). The advertising mix for a search good. *Management Science*, 59(1), 69-83.
- Anderson, S. P., Renault, R. (2016). Search direction. Mimeo.
- Arbatskaya, M. (2007). Ordered search. *The RAND Journal of Economics*, 38(1), 119-126.
- Armstrong, M. (2016). Ordered Consumer Search. CEPR Discussion Paper 11566.
- Armstrong, M., Vickers, J., Zhou, J. (2009a). Prominence and Consumer Search. *The RAND Journal of Economics*, 40(2), 209-233.
- Armstrong, M., Vickers, J., Zhou, J. (2009b). Consumer Protection and the Incentive to Become Informed. *Journal of the European Economic Association*, 7(2-3), 399-410.
- Armstrong, M., Zhou, J. (2011). Paying for Prominence. *Economic Journal*, 121(556), F368-F395.
- Athey, S, Ellison, G. (2011). Position Auctions with Consumer Search. *Quarterly Journal of Economics*, 126(3), 1213-1270.
- Bar-Isaac, H., Caruana, G., Cunat, V. (2012). Search, Design, and Market Structure. *American Economic Review*, 102(2), 1140-1160.
- Baye M., Kovenock, D., de Vries, C. G. (1992) It takes two to tango: equilibria in a model of sales. *Games and Economic Behavior*, 4(4), 493-510.
- Baye, M., Morgan, J. (2001). Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets. *American Economic Review*, 91(3), 454-474.
- Baye, M., Morgan, J., Scholten, P. (2006), Information, Search, and Price Dispersion. In *Handbook of Economics and Information Systems*, Vol. 1, Elsevier.
- Bénabou, R. (1993). Search Market Equilibrium, Bilateral Heterogeneity, and Repeat Purchases. *Journal of Economic Theory*, 60(1), 140-158.
- Bénabou, R., Gertner, R. (1993). Search with Learning from Prices: Does Increased Inflationary Uncertainty Lead to Higher Markups?. *Review of Economic Studies*, 60(1), 69-93.

- Burdett, K., Judd, K. L. (1983). Equilibrium price dispersion. *Econometrica*, 51(4), 955-969.
- Butters, G. (1977). Equilibrium distributions of sales and advertising prices. *Review of Economic Studies*, 44(3), 465-491.
- Caillaud, B., Jullien, B. (2003). Chicken & Egg: Competition among Intermediation Service Providers. *The RAND Journal of Economics*, 34(2), 309-328.
- Caplin, A., Nalebuff, B. (1991). Aggregation and Imperfect Competition: On the Existence of Equilibrium. *Econometrica*, 59(1), 25-59.
- Carlson, J., McAfee, R. (1983). Discrete Equilibrium Price Dispersion. *Journal of Political Economy*, 91(3), 480-93.
- Chen, Y. He, C. (2011). Paid Placement: Advertising and Search on the Internet. *Economic Journal*, 121, F309-F328.
- Dana, J. (1994). Learning in an Equilibrium Search Model. *International Economic Review*, 35(3), 745-771.
- De Cornière, A. (2016). Search Advertising. Mimeo, Oxford University.
- Diamond, P. A. (1971). A model of price adjustment. *Journal of economic theory*, 3(2), 156-168.
- Diamond, P. A. (1987). Consumer Differences and Prices in a Search Model. *Quarterly Journal of Economics*, 102(2), 429-436.
- Edelman, B, Ostrovsky, M., Schwarz, M. (2007). Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review*, 97(1), 242-259.
- Ellison, G., Ellison, S. F. (2014). Match Quality, Search, and the Internet Market for Used Books. Mimeo.
- Ellison, G., Wolitzky, A. (2012). A search cost model of obfuscation. *The RAND Journal of Economics*, 43(3), 417-441.
- Fershtman, C., Fishman, A. (1992). Price Cycles and Booms: Dynamic Search Equilibrium. *The American Economic Review*, 82(5), 1221-1233.
- Fishman, A. (1996). Search with Learning and Price Adjustment Dynamics. *Quarterly Journal of Economics*, 111(1), 253-268.
- Fudenberg, D., Tirole, J. (1991). Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2) 236-260.
- Goldmanis, M. Hortaçsu, A. Syverson, C, Emre, Ö. (2010). E-Commerce and the Market Structure of Retail Industries. *The Economic Journal*, 120(545), 651-682.
- Haan, M., Moraga-González, J. L. (2011) Advertising for attention in a consumer search model. *The Economic Journal* 121 (552), 552-579.

Hortaçsu, A., Syverson, C. (2004). Product Differentiation, Search Costs and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds. *Quarterly Journal of Economics*, 119(2), 403-456.

Hotelling, H. (1929). Stability in competition. *Economic Journal*, 39, 41-57.

Janssen, M., Moraga-González, J. L., Wildenbeest, M. R. (2005). Truly costly sequential search and oligopolistic pricing. *International Journal of Industrial Organization*, 23(5-6), 451-466.

Janssen, M., Moraga-González, J. L. (2004) Strategic Pricing, Consumer Search and the Number of Firms. *Review of Economic Studies*, 71(4), 1089-1118.

Janssen, M., Moraga-González, J. L. (2008). On Mergers in Consumer Search Markets. Mimeo.

Janssen, M., Non, M. (2008). Advertising and consumer search in a duopoly model. *International Journal of Industrial Organization*, 26(1), 354-371.

Janssen, M., Parakhonyak, A., Parakhonyak, A. (2016). Non-reservation Price Equilibria and Consumer Search. Mimeo.

Janssen, M., Pichler, P., Weidenholzer, S. (2011). Sequential Consumer Search with Incompletely Informed Consumers. *The RAND Journal of Economics*, 42(3), 444-470.

Janssen, M., Shelegia, S. (2016). Beliefs and Consumer Search. Mimeo.

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

Kocas, C., Kiyak, T. (2006). Theory and evidence on pricing by asymmetric oligopolies. *International Journal of Industrial Organization*, 24(1), 83-105.

Kohn, M., Shavell, S. (1974). The Theory of Search. *Journal of Economic Theory*, 9, 93-123.

Kreps, D. M., Wilson, R. (1982). Sequential equilibria. *Econometrica*, 50(4), 863-894.

Lach, S. (2002). Existence and Persistence of Price Dispersion: an Empirical Analysis. *Review of Economics and Statistics*, 84(3), 433-444.

Larson, N. (2013). Niche products, generic products, and consumer search. *Economic Theory*, 52(2), 793-832.

Lewis T. R., Sappington, D. E. (1994) Supplying information to facilitate price discrimination. *International Economic Review*, 35(2), 309-27.

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

Moraga-González, J. L., Sandor, Z., Wildenbeest, M. (2016). Prices and Heterogeneous Search Costs. Forthcoming, *RAND Journal of Economics*.

- Morgan, J., Orzen, H., and Sefton, M. (2006) An Experimental Study of Price Dispersion. *Games and Economic Behavior*, 54(1), 134-158.
- Perloff, J., Salop, S. (1985). Equilibrium with Product Differentiation. *Review of Economic Studies*, 52(1), 107-120.
- Peters, M. (1984). Bertrand Equilibrium with Capacity Constraints and Restricted Mobility. *Econometrica*, 52(5), 1117-1127.
- Petrikaitė, V. (2016). Collusion With Costly Consumer Search. *International Journal of Industrial Organization*, 44, 1-10.
- Renault, R. (2016). Advertising in Markets. In *Handbook of Media Economics*, Vol 1A., Elsevier.
- Reinganum, J. (1979). A Simple Model of Equilibrium Price Dispersion. *Journal of Political Economy*, 87(4), 851-858.
- Rhodes, A., Zhou, J. (2016). Consumer Search and Retail Market Structure. Mimeo.
- Rob, R. (1985). Equilibrium Price Distributions. *Review of Economic Studies*, 52(3), 487-504.
- Robert, J., Stahl II, D. O. (1993). Informative price advertising in a sequential search model. *Econometrica*, 61(3), 657-686.
- Rothschild, M. (1974). Searching for the Lowest Price When the Distribution of Prices Is Unknown. *Journal of Political Economy*, 82(4), 689-711.
- Salop, S. (1977). The Noisy Monopolist: Imperfect Information, Price Dispersion and Price Discrimination. *Review of Economic Studies*, 44(3), 393-406
- Salop, S., Stiglitz, J. (1977). Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion. *Review of Economic Studies*, 44(3), 493-510.
- Schulz N., Stahl K. (1996) Do consumers search for the highest price? Oligopoly equilibrium and monopoly optimum in differentiated-products markets. *The RAND Journal of Economics*, 27, 542-62.
- Shaked, A., Sutton, J. (1983). Natural Oligopolies. *Econometrica*, 51(5), 1469-1483.
- Song, H. (2016). Ordered Search with Asymmetric Product Design. Forthcoming, *Journal of Economics*.
- Stahl II, D. O. (1989). Oligopolistic Pricing with Sequential Consumer Search. *American Economic Review*, 79(4), 700-712.
- Stahl II, D. O. (1996). Oligopolistic Pricing with Heterogeneous Consumer Search. *International Journal Industrial Organization*, 14(2), 243-268.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3), 213-225.

Stiglitz, J. E. (1979). Equilibrium in product markets with imperfect information. *American Economic Review*, 69(2), 339-345.

Tappata, M. (2009). Rockets and feathers: Understanding asymmetric pricing. *RAND Journal of Economics*, 40(4), 673-687.

Tirole, J. (1988). *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.

Varian, H. R. (1980). A Model of Sales. *American Economic Review*, 70(4), 651-659.

Weitzman, M. (1979). Optimal Search for The Best Alternative. *Econometrica*, 47(3), 641-654.

Weyl, E. G., Fabinger, M. (2013) Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition. *Journal of Political Economy*, 121(3), 528-583.

Wolinsky, A. (1984). Product Differentiation with Imperfect Information. *Review of Economic Studies*, 51(1), 53-61.

Wolinsky, A. (1986). True monopolistic competition as a result of imperfect information. *Quarterly Journal of Economics*, 101(3), 493-512.

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

Zhou, J. (2014). Multiproduct search and the joint search effect. *American Economic Review*, 104(9), 2918-2939.