



Munich Personal RePEc Archive

Conditioning prices on search behaviour

Armstrong, Mark and Zhou, Jidong

University College London (UCL)

January 2010

Online at <https://mpra.ub.uni-muenchen.de/19985/>
MPRA Paper No. 19985, posted 15 Jan 2010 14:14 UTC

Conditioning Prices on Search Behaviour*

Mark Armstrong
University College London

Jidong Zhou
University College London

January 2010

Abstract

We consider a market in which firms can partially observe each consumer's search behavior in the market. In our main model, a firm knows whether a consumer is visiting it for the first time or whether she is returning after a previous visit. Firms have an incentive to offer a lower price on a first visit than a return visit, so that new consumers are offered a “buy-now” discount. The ability to offer such discounts acts to raise all prices in the market. If firms cannot commit to their buy-later price, in many cases firms make “exploding” offers, and consumers never return to a previously sampled firm. Likewise, if firms must charge the same price to all consumers, regardless of search history, we show that they sometimes have the incentive to make exploding offers. We also consider other ways in which firms could use information about search behaviour to determine their prices.

Keywords: Consumer search, oligopoly, price discrimination, high-pressure selling, exploding offers, costly recall.

1 Introduction

In a variety of markets in which consumers sequentially search through firms' options, sellers know something about their potential customers' search behaviour. For example, a seller may be able to distinguish visitors who come to the store for the first time from those who have returned after a previous visit. Salesmen may try to elicit information from a customer about how many other sellers she has already sampled and how many more she plans to sample. In online markets, sellers may be able to monitor the entire search process of a potential customer.

*We are grateful to Simon Anderson, Marco Haan, Preston McAfee, Meg Meyer, Andrew Rhodes and Asher Wolinsky, and to the Economic and Social Research Council (UK) and the British Academy for funding assistance.

There are two broad reasons why a firm may wish to condition its price on a consumer's search behaviour. First, search history may reveal relevant *information* about a consumer's tastes. For example, a seller may wish to charge a higher price to those consumers who have already investigated other sellers, because their decision to keep searching indicates they are unsatisfied with previously sampled products when search is costly. (If consumers have heterogenous search costs, it may also reveal that such consumers have relatively low search costs, which may by contrast induce a seller to set a low price.) Second, firms may have a *strategic* reason to condition its price on search history, which is to deter a potential consumer from going on to search rival offers. For example, if a firm can distinguish new visitors from consumers who have previously visited, a firm can offer new customers a relatively low price or force them to buy its product now or not at all, which makes customers more reluctant to investigate rival offers. This paper investigates whether setting prices conditional on search behaviour, if feasible, can emerge as an equilibrium outcome, and, if so, how this practice affects market performance.

Our underlying framework is a duopoly sequential search model with horizontally differentiated products in which consumers search both for price and product fitness. Each firm has three sources of demand: consumers who sample its product first and buy immediately ("fresh consumers"), consumers who sample the firm first, go on to sample the rival's product, but eventually come back ("returning consumers"), and consumers who sample the rival first but go on to buy from the firm in question ("incoming consumers"). In the standard search model, firms cannot distinguish between these three groups and so offer the same price to all their visitors. In our models, firms are able to recognize and price discriminate between at least some of these different consumer groups.

In our main model, presented in section 3, firms can distinguish returning consumers from other consumers (i.e., they can tell old faces from new faces), and can charge them a different price. We show that in equilibrium the price for returning consumers is higher than that offered to others, or equivalently the consumers who visit a firm for the first time will be offered a "buy-now discount". Consumers are therefore more reluctant to sample the rival firm's product. Due to this effect, even the discounted buy-now price is higher than the non-discriminatory price. As such, this form of price discrimination lowers both consumer surplus and total welfare. When the search cost is relatively high, price discrimination will lead to excessively high prices even from the firms' perspective, and the ability to offer buy-now discounts reduces industry profits. In section 3.3 we relax our assumption that firms commit to their buy-later price when consumers make their first visit. When firms can commit to an upper bound on buy-later prices (e.g., by displaying a credible "regular price" on the shop shelves) then the full commitment outcome is still achieved. When firms have no commitment power,

we argue that they are often forced to make “exploding offers”, i.e., new visitors are forced to buy immediately or not at all. In section 3.4, we present a related model in which firms cannot price discriminate but can artificially inflate a consumer’s cost of returning. (Thus, only the “strategic effect” mentioned above is relevant.) In this situation, we derive conditions under which firms choose to make exploding offers.

In section 4, we investigate other ways in which firms can condition their prices on search behaviour. We first suppose that firms can distinguish all three groups of consumers and so, in principle can set three different prices. We show that firms in equilibrium will set the same (high) price to returning and to incoming customers, and offer a discount only to fresh customers. Thus, firms discriminate equally against those customers who have sampled the rival, regardless of whether they are incoming or returning customers. We then consider the case in which each firm can distinguish consumers who sample it first (i.e., fresh and returning consumers) from those who sample its rival first (i.e., incoming consumers). Since incoming consumers must be unsatisfied with the rival’s product, in equilibrium firms charge them a high price.

There has apparently been little research exploring the question of how firms might make use of knowledge of consumer search behaviour or search history to refine their pricing decisions. Nevertheless, as we discuss next, our paper relates to several strands of the industrial organization literature. In the recent literature on ordered search, consumers sequentially search through options according to some publicly known order, which automatically reveals part of their search history to sellers. Armstrong, Vickers, and Zhou (2009) consider a model where consumers search for both price and product fitness (as in this paper) and where consumers consider one “prominent” firm first. In equilibrium the prominent firm sets a lower price than the others: firms further back in the search order know that any customer who reaches them does not particularly value the previous offers made, and so such firms can exploit their market power by setting a higher price. (This paper is discussed further in section 4 below.) Arbatskaya (2007) analyzes a related model where consumers care only about price but differ in their search cost. She obtains the reverse finding, which is that firms further back in the search order set *lower* prices. Intuitively, this is because these firms are only visited by consumers with relatively low search costs and so face more competition.¹

Related effects can be obtained in models of simultaneous, as opposed to sequential, search. For instance, in a differentiated product setting, Gal-Or, Gal-Or, and Dukes (2007) suppose that the buyer can announce from how many sellers she is soliciting quotes, and sellers can condition their price on that information. As is intuitive, in many cases (see their *Corollary 1*) firms offer a lower price when they know they

¹More distantly related is the idea of a “stigma effect” of long-term unemployment. For instance, Vishwanath (1989) supposes that potential employers can observe how many times a worker has failed to generate a job offer, and can condition wage offers on this variable.

must compete against more rivals. Deck and Wilson (2006) conduct an experimental study of a related market. Firms (played by experimental subjects) sell a homogenous product, and consumers (played by price-minimizing robots) are exogenously divided into several groups according to the number of firms they sample. Each firm knows how many other firms a consumer samples and can condition the price on that information. In their experiment, more informed consumers are treated better when sellers can price discriminate relative to when sellers do not observe a consumer’s search behaviour.

In our main model firms offer a “buy-now” discount, and this acts to increase the cost of returning to a previous firm. As such, our analysis is related to models of search with costly recall. Janssen and Parakhonyak (2008) study the optimal stopping rule when consumers care only about price and must incur a cost to return to a previous firm. This stopping rule is significantly more complicated than when return is costless. When there are more than two firms, a consumer’s stopping rule is non-stationary, her reservation surplus level is not independent of her previous offers, and a consumer may end up buying a product which is more expensive than others she has sampled. (This is the pragmatic reason why we focus on a duopoly model, where these complexities are absent.) However, in their model, the supply side is passive and sellers’ prices are randomly drawn from an exogenous distribution. Our paper is also related to models of search with endogenous recall. Daughety and Reinganum (1992) make the point that the extent of consumer recall may be endogenously determined by firms’ equilibrium strategies. In their model, the instrument that a firm can use to influence consumer recall is the length of time that it will hold the good for consumers at the quoted price. (In our main model, that instrument is the buy-now discount.) In contrast to our assumption that a consumer can discover a seller’s return cost only after investigating that seller, Daughety and Reinganum suppose that sellers can announce their recall policies to the population of consumers before search begins.

Our paper relates also to situations with high-pressure selling. Broadly speaking, “high-pressure selling” can be classified into at least two types. One type aims to “persuade” (maybe boundedly rational or socially nervous) consumers to buy unnecessary or over-priced products such as some extended warranties, some payment protection insurance, and so on.² Another type of high-pressure selling aims to reduce the consumer’s incentive to search further. For example, the seller can commit (or at least claim) that the current price offer—or even the availability of the product—will expire if consumers do not buy immediately. Sellers of financial products may sometimes attempt this sales technique, especially if they are in the home, and the practice is also used in some specialized labour markets.³

²See, for instance, Chu, Gerstner, and Hess (1995) for a model of this type of persuasive high-pressure selling.

³One of the authors encountered an in-home salesman of financial products, and when he said he

Firms often benefit from the reduction of consumer search intensity, since this usually softens price competition. In our model, the buy-now discount or the exploding offer serves this purpose. Alternatively, Ellison and Wolitzky (2008) present a model in which consumers have convex search costs (i.e., a consumer’s incremental search cost increases with her cumulative search effort). Then if a firm increases its in-store search cost (say, by making its price harder to comprehend), it will make further search more costly, and thus less likely. Though very different, the two models lead to a similar result that, even if the exogenous search cost tends to zero, the market may not be frictionless in equilibrium.

De los Santos (2008) presents a rare empirical study of individual consumer search behaviour prior to making a purchase, using data from online book purchases. In this particular market there is no evidence of price discrimination of the kind we study in this paper (as far as we know), but the patterns of search behaviour are nevertheless relevant for our work. In his data, De los Santos (2008, section 4) finds that three-quarters of consumers search only one retailer before making their purchase, i.e., using our terminology, fresh demand makes up 75% of total demand. Of the remaining consumers who search at least twice, approximately two-thirds buy from the final firm searched (“incoming demand”) and one-third go back to a firm searched earlier (“returning demand”). De los Santos also finds that the initial search is non-random (unlike our main model below), and one firm (*Amazon.com*) was sampled first by about two-thirds of all consumers making a purchase. (We discuss the impact of having one firm more prominent in section 5 below.)

Our analysis of buy-now discounts is also somewhat related to the emerging literature on auctions with a “buy now” price (see Reynolds and Wooders (2009), for instance). Online auctions such as those run by *eBay* sometimes offer bidders the option to buy the item immediately at a specified price rather than enter an auction against other bidders. In these situations, a seller has one item to sell to a number of potential bidders, and so a bidder needs to pay a *high* buy-now price in order to induce the seller from going on to “search” for other bidders by running an auction, whereas our model involves sellers offering a *low* buy-now price so as to induce a buyer from going on to search for other sellers. Common rationales for buy-now prices in auctions are impatience or risk-aversion on the part of bidders, neither of which is needed in our framework with costly search.

Although there is not a great deal of work on conditioning prices on search history,

wished to think about the offer and get back to the salesman, the salesman claimed it was his last day in his current job. Roth and Xing (1994, page 1001) document some examples of high-pressure job offers. For instance, in the market for judicial clerkships, some judges use exploding offers which would be withdrawn if they are not accepted in some very short time, or even during the telephone conversation itself.

there is a substantial literature about how firms can use the information of consumer *purchase* history to refine their pricing decisions. See, for instance, Hart and Tirole (1988), Chen (1997), Fudenberg and Tirole (2000), and Acquisti and Varian (2005). These models often predict that a firm will price low to a customer who previously purchased from a rival (or consumed the outside option in the case of monopoly), since such a customer has revealed she has only a weak preference for the firm’s product. In the current paper, though, we will see that a firm will price high to a customer who previously *searched* a rival, since that customer has revealed she has relatively weak preferences for the rival’s product.

2 A Model and Preliminary Results

Our underlying model of consumer choice is based on Wolinsky (1986), simplified to duopoly.⁴ (See Anderson and Renault (1999) for a further development of Wolinsky’s model.) There are two firms $i = A, B$ in the market, each supplying a horizontally differentiated product at zero production cost. A consumer’s valuation of product i , u_i , is a random draw from some common distribution with support $[0, u_{\max}]$ and with cumulative distribution function $F(\cdot)$ and density $f(\cdot)$. We suppose that the realization of match utility is independent across consumers and products. In particular, there are no systematic quality differences across the products. Each consumer has a unit demand for only one product, and for expositional simplicity the number of consumers is normalized to *two*.

Consumers initially have imperfect information about prices and match utilities. But they can gather this information through a sequential search process: by incurring a search cost $s \geq 0$, a consumer can visit a firm and find out its price and match value.⁵ After sampling the first firm, a consumer can choose to buy at this firm immediately or continue on to investigate the second firm. If she continues to search, she can buy from the second firm or return to buy from the first firm, or neither. If she returns to buy from the first firm, she does not incur a further search cost.⁶ We focus on symmetric

⁴Note that if there were infinitely many firms, there is no role for conditioning prices on search history. Even if consumers can costlessly return to a previously sampled product, when consumers have unlimited options it is a standard result in search theory that consumers optimally do not return to previous options, but always prefer to search for a new option. Note also that this duopoly model applies when there are more than two sellers if, for some reason, consumers never consider more than two options.

⁵If the search cost is zero, we require that consumers nevertheless consider products sequentially.

⁶Our analysis could straightforwardly extend to situations where consumers have an intrinsic cost of returning to a previously sampled firm, but few new insights emerge with that extension. The main exception is in section 3.3, where we consider the situation where firms cannot commit to their buy-later price and show that the presence of even a small returning cost will shut down the market for returning customers altogether.

situations in which half the consumers sample firm A first, while the other half first sample B .⁷ Each firm has three sources of demand: consumers who sample its product first and buy immediately (“fresh consumers”), consumers who sample its product first, go on to investigate the rival firm, but eventually come back (“returning consumers”), and consumers who sample its rival first but go on to buy from the firm in question (“incoming consumers”).

Firms maximize their profits, and set prices simultaneously based on their expectation of consumer behavior. However, their pricing decisions depend on how much information they have about consumers’ search history. If all consumers are indistinguishable, as in Wolinsky (1986), each firm must set a uniform price. But if some groups of consumers are distinguishable, firms may be able to price discriminate. There are a number of different kinds of price discrimination which could be considered. In our main model, analyzed in section 3, we suppose that firms can distinguish returning consumers from other consumers, but cannot distinguish incoming from fresh consumers. For instance, this might apply when a consumer needs to interact with a seller to discuss specific requirements which reveal the consumer’s identity, and it is therefore clear if a consumer has previously consulted the seller. (Examples might include job offers, medical insurance, or home improvements.) Likewise, in online markets sellers might place “cookies” on a consumer’s computer which reveal whether the consumer has visited previously. In such cases, it is often the case that a seller does not know whether the consumer has also interacted with rival sellers. In section 4, we consider the alternative situation in which a firm does observe whether a consumer has visited a rival.

When a firm sets a higher price to returning customers than to customers on their first visit, a customer faces a returning cost if she chooses to leave and then return to the firm, since the price she has to pay will then rise. Similarly, in our model of high-pressure selling in section 3.4, firms impose a non-monetary returning cost. As a necessary ingredient for each of these models, we next derive the optimal stopping rule for consumers in a search model with an arbitrary (monetary or non-monetary) returning cost, which may be firm specific and is denoted τ_i .

A consumer’s optimal stopping rule balances the surplus of the first sampled product and the expected surplus from searching on. Denote by

$$V(p) \equiv \int_p^{u_{\max}} (u - p) dF(u) - s \tag{1}$$

the expected surplus of sampling the second product if a consumer expects that the second firm will charge price p and she surely does not return to buy the first product. Define

$$a_\tau \equiv z_\tau + \tau, \tag{2}$$

⁷We discuss the impact of one firm being more prominent in section 5 below.

where z_τ solves

$$V(z_\tau) \equiv \tau . \quad (3)$$

Note that both z_τ and a_τ decrease with τ . Indeed,

$$\frac{da_\tau}{d\tau} = \frac{-F(z_\tau)}{1 - F(z_\tau)} . \quad (4)$$

As we show next, a_τ determines the consumer's reservation utility level when she anticipates a returning cost τ if she leaves and then returns to a firm. That is, if the two firms charge the same tariffs, a consumer will investigate the second firm when the match utility of the first product is below a_τ . In more detail:

Lemma 1 *Consider a consumer who encounters firm i first and is offered the price p_i and match utility u_i . Suppose the consumer expects that firm j will charge p with $V(p) > 0$ and expects that she will incur a returning cost τ_i if she returns to firm i after sampling firm j (in addition to firm i 's price p_i).*

(i) *If $\tau_i < V(p)$, she will search on to firm j if and only if*

$$u_i - p_i < a_{\tau_i} - p , \quad (5)$$

and she will return to firm i after sampling firm j if $u_i - p_i - \tau_i > \max\{0, u_j - p_j\}$, where p_j is firm j 's actual price.

(ii) *If $\tau_i \geq V(p)$, she will search on to firm j if and only if*

$$u_i - p_i < V(p) ,$$

and she will never return to firm i after sampling firm j (regardless of firm j 's actual price p_j).

(All omitted proofs can be found in the Appendix.)

Note that when $\tau_i = V(p)$, the cut-off surplus levels in cases (i) and (ii) coincide. Note also that in case (i) some consumers will indeed return to firm i after investigating firm j . Those consumers with match utility just below the threshold in (5) will go on to investigate j , and if j 's match utility is low (below p_j say), then the consumer still obtains positive utility if she returns to i .⁸

When there is no returning cost ($\tau_i = 0$), Lemma 1 implies that the consumer will sample the second firm whenever $u_i - p_i < a_0 - p$.⁹ This is the stopping rule in the search

⁸Formally, when $u_i \approx a_{\tau_i} + p_i - p$ so that (5) is nearly satisfied, then $u_i - p_i - \tau_i \approx a_{\tau_i} - \tau_i - p = z_{\tau_i} - p$, which is positive since $V(\cdot)$ is a decreasing function and by assumption $V(z_{\tau_i}) = \tau_i < V(p)$.

⁹Notice that the stopping rule also applies for a negative returning cost $\tau_i \geq -s$. (If $\tau_i < -s$, consumers will always sample the second product.) We show later that this does not occur in equilibrium. In addition, if $\tau_i < 0$ (which might occur if a firm's buy-later price \hat{p}_i is lower than its buy-now price p_i) then a consumer has an incentive to leave firm i and return, even if she has no intention of investigating the rival. If this kind of consumer arbitrage behavior—of stepping out the door and then back in again—cannot be prevented, then setting $\hat{p}_i < p_i$ is equivalent to setting a uniform price \hat{p}_i , and so without loss of generality we can assume firms are constrained to set $\tau_i \geq 0$.

model with costless recall as described in Wolinsky (1986). Since a_τ decreases with the returning cost τ , it follows that if the returning cost increases, a consumer becomes less willing to investigate the rival firm. (The utility from sampling the rival falls with τ since there are fewer profitable opportunities to return to the initial firm.) The case where $\tau_i < V(p)$ is illustrated in Figure 1, where the shaded area is the region where the consumer buys neither product. (Recall that p is the price which consumers anticipate firm j will charge, while p_j is the price that firm j actually offers.) When τ_i exceeds $V(p)$, the returning cost is so great that the consumer never returns to firm i once she leaves it and returning demand vanishes. Hence, the continuation payoff for sampling the second firm is $V(p)$, and she will sample both products whenever $u_i - p_i < V(p)$.




Figure 1: Pattern of demand for consumers who sample firm i first ($\tau_i < V(p)$)

In the remainder of the paper, an important consideration is whether a higher cost for returning to a firm acts to boost or reduce that firm's total demand. First, it is clear that a firm's returning cost τ_i cannot affect its incoming demand, since incoming consumers buy immediately or not at all. Second, Figure 1 shows that increasing τ_i boosts a firm's fresh demand. Third, the figure also shows that the firm's returning demand contracts with a higher τ_i . Thus, the impact on the firm's total demand is unclear *a priori*. The next result discusses the net impact in some special cases:

Lemma 2 *Suppose that firm j 's actual incoming price p_j is equal to the anticipated price p . Then an increase in firm i 's returning cost τ_i will:*

(i) *increase the firm's total demand if the density f is increasing with u ;*

- (ii) decrease the firm's total demand if the density f is decreasing with u ;
- (iii) leave the firm's total demand unchanged if u is uniformly distributed.

The intuition for this result is as follows. A fresh consumer who buys at firm i immediately has a higher valuation of the product than a returning consumer. When f is increasing, there are more consumers with higher valuations in the population, and so the effect of a higher recall cost on increasing fresh demand is more significant than its impact on returning consumers. When f is decreasing, the opposite is true. When utilities are uniformly distributed, the net impact of the returning cost on total demand is zero, and an increase in τ_i causes demand to be substituted from returning to fresh demand in a one-for-one manner.

As discussed earlier, a firm has two broad reasons to condition its offer on a consumer's search behaviour: an informational motive based on what it learns about the consumer's preferences, and a strategic motive to influence the consumer's decision to engage in further search. The strategic motive is essentially captured by the impact of the return cost on a firm's total demand. Thus, if the density is increasing, a firm has an incentive to make it hard for consumers to return to it, as the resulting gain in fresh demand outweighs the loss of returning customers. As we discuss further in section 3.4, this implies that when the density is increasing, firms have an incentive to force potential consumers to buy "now or never".

3 "Buy-Now" Discounts

Consider the situation in which firms are able to recognize returning consumers and charge them a different price. Each firm then potentially sets two prices: p_i is firm i 's price to those consumers who visit it for the first time (comprising both the fresh and incoming consumers), while \hat{p}_i is its price offered to a consumer when she returns to the firm. An interpretation is that each firm sets a "regular" (or "buy-later") price \hat{p}_i and offers the first-time visitors a "buy-now" discount $\tau_i \equiv \hat{p}_i - p_i$. We assume that a firm can commit to \hat{p}_i when it offers new visitors the buy-now price p_i . (We discuss the impact of more limited commitment later in section 3.3.)

For simplicity, from now on we focus mostly on the case where match utility u_i is uniformly distributed on $[0, 1]$, so that $F(u) \equiv u$ and expression (2) becomes

$$a_\tau = 1 + \tau - \sqrt{2(s + \tau)}.$$

As we saw in Lemma 2, this assumption essentially shuts down a firm's "strategic" motive to make return more costly, since a firm's total demand is unaffected by its return cost with a uniform distribution. This greatly simplifies the algebra and enables us to check second-order conditions for the candidate equilibrium. However, our main

result that firms will choose to offer a buy-now discount continues to hold for more general distributions for the match utility, provided that the density function for u does not decrease “too fast”. (See Table 1 below for illustrations of the equilibrium when the distribution of match utility is non-uniform.)

To ensure an active market we assume a relatively small search cost $0 \leq s \leq \frac{1}{8}$. (If $s > \frac{1}{8}$ then there is no equilibrium in which consumers sample even one firm’s offer.) We also focus on symmetric equilibria where both firms offer the same price pair (p, \hat{p}) .

3.1 Equilibrium prices

For convenience, we analyze the model in terms of the buy-now price p and the total returning cost $\tau = \hat{p} - p$ (rather than in terms of p and \hat{p}). Suppose firm j adopts the equilibrium strategy (p, τ) and consumers expect that both firms will offer these prices. If firm i unilaterally deviates to (p_i, τ_i) , with $\tau_i \leq V(p)$, its profit is¹⁰

$$p_i Q_T + \tau_i Q_R, \quad (6)$$

where Q_T is firm i ’s total demand and Q_R is the portion of demand from its returning customers. (The firm obtains revenue p_i from all its customers, plus the incremental revenue $\hat{p}_i - p_i = \tau_i$ from its returning customers.) These demand functions can be derived by calculating the areas of the various regions in Figure 1 to yield¹¹

$$Q_T = \underbrace{1 - (p_i + a_{\tau_i} - p)}_{\text{fresh demand}} + \underbrace{\frac{1}{2}(a_{\tau_i} - \tau_i)^2 - \frac{1}{2}p^2}_{\text{returning demand } Q_R} + \underbrace{a_{\tau}(1 - p_i) - \frac{1}{2}(a_{\tau} - p - \tau)^2}_{\text{incoming demand}}. \quad (7)$$

Note that the firm’s returning demand does not depend on its buy-now price p_i over the relevant range. (By examining Figure 1, we see that varying p_i simply shifts the region of returning demand uniformly to the left or right.) Note also that the firm’s total demand Q_T does not depend on its buy-now discount τ_i . (This follows from part (iii) of Lemma 2.) Thus, a firm’s profit in (6) is additively separable in its buy-now price p_i and its buy-now discount τ_i .

In particular, firm i will choose its returning cost τ_i to maximize $\tau_i Q_R(\tau_i)$, the extra revenue from its returning consumers. We deduce that given the equilibrium buy-now price p , the buy-now discount τ maximizes the incremental industry profit from returning customers, which is

$$\tau((a_{\tau} - \tau)^2 - p^2) = \tau(z_{\tau}^2 - p^2). \quad (8)$$

¹⁰When $\tau_i \geq V(p)$, the returning demand disappears and fresh demand becomes $1 - [p_i + V(p)]$, and so the profit will be independent of τ_i . Hence, our restriction to $\tau_i \leq V(p)$ is without loss of generality.

¹¹To calculate firm i ’s incoming demand use Figure 1 to calculate firm j ’s incoming demand, but permute the two firms, i.e., replace p_j with p_i and replace p_i with p . (A consumer who visits firm j first expects that firm i will offer the equilibrium price p , although firm i in fact has price p_i .)

It is clear (e.g., by using a revealed preference argument) that the value of τ which maximizes (8) is a decreasing function of p .¹² In equilibrium, as we will show, we have $p < a_0 = z_0$ provided that $s < \frac{1}{8}$. Since z_τ is a decreasing function of τ , it follows that the profit generated from returning consumers in (8) is negative whenever $\tau < 0$, and strictly positive for τ in a neighborhood above 0. We deduce that firms in equilibrium will choose a strictly positive buy-now discount if $s < \frac{1}{8}$. Note that a firm has an incentive to offer a positive buy-now discount even if its rival does not. Moreover, returning demand falls to zero if $\tau \geq V(p)$, i.e., if $z_\tau \leq p$, and from (8) we can deduce that in equilibrium $p < z_\tau$ whenever $s < \frac{1}{8}$. From Figure 1, it follows that there is indeed a region of returning demand in equilibrium. That is to say, although firms *could* force consumers to buy “now or never” by choosing a very large τ , it is more profitable for them to induce at least some consumers to return after sampling the rival’s offer and to extract some additional revenue τ from them.

One can show that (8) is concave in τ in the (relevant) region $0 \leq \tau \leq \frac{1}{8}$ (and is decreasing in τ for larger τ).¹³ From (8) and noting that $dz_\tau/d\tau = -1/(1 - z_\tau)$, the first-order condition for equilibrium τ given p is

$$z_\tau^2 - p^2 - \frac{2\tau z_\tau}{1 - z_\tau} = 0 . \quad (9)$$

Turning to the equilibrium choice of buy-now price p , note that firm i ’s total demand in (7) is linear in p_i , and so its profit is concave in p_i . Therefore, the first-order condition for p_i to be optimal is also sufficient. Each firm’s equilibrium total demand is $Q_T = 1 - p(p + \tau)$. That is, a consumer will leave the market without buying anything if and only if she neither buys at the second firm nor wants to go back to the first one. Using this fact, one can see that the first-order condition for the equilibrium buy-now price p , given τ , is

$$\frac{1}{p} - p = 1 + z_\tau + 2\tau . \quad (10)$$

The right-hand side of (10) is at least $\frac{3}{2}$.¹⁴ Since the left-hand side of (10) is decreasing in p , it follows that the solution to this first-order condition satisfies $p \leq \frac{1}{2}$. (In particular, since $a_0 = z_0 \geq \frac{1}{2}$, it follows that $p < a_0 = z_0$ as claimed previously.) Moreover, for $0 \leq \tau \leq \frac{1}{8} - s$, which will be the relevant range of τ , the right-hand side of (10)

¹²More precisely, the firm’s optimal choice for τ_i decreases with its rival’s buy-now price.

¹³Expression (8) is decreasing in τ if and only if $z_\tau^2 - p^2 + 2\tau z_\tau z'_\tau < 0$, which is certainly true if $z_\tau + 2\tau z'_\tau < 0$. But since $z'_\tau = -1/(1 - z_\tau)$, this is true if $2\tau > z_\tau(1 - z_\tau)$, which in turn is always true if $\tau > \frac{1}{8}$. Thus, profit decreases in τ for $\tau > \frac{1}{8}$. Next, suppose that $\tau \leq \frac{1}{8}$. Differentiating the first-order condition once more shows that (8) is concave in τ if and only if $\tau < 2z_\tau(1 - z_\tau)^2$, i.e., if $\tau/(s + \tau) < 4(1 - \sqrt{2(s + \tau)})$. But when $s, \tau < \frac{1}{8}$ the right-hand side of this is greater than 1, and so the claim follows.

¹⁴Note that $z_\tau + 2\tau$ is a convex function which is minimized by setting $\tau = \frac{1}{8} - s$, which makes the right-hand side of (10) equal to $\frac{7}{4} - 2s$. Since $s \leq \frac{1}{8}$, the claim follows.

is decreasing in τ , and so the buy-now price p in (10) is an increasing function of τ . Intuitively, a buy-now discount increases search frictions in the market, which in turn allows firms to charge a higher price.

As an aside, note that expression (10) reveals the equilibrium price in the alternative situation where firms choose a uniform price p and consumers have an *exogenous* returning cost τ for returning to a previously-sampled firm. (In this situation, firm i 's profit is modified from (6) to be $p_i Q_T$, where we set $\tau_i = \tau$ in (7). The first-order condition for the equilibrium price is therefore still (10).) As far as we know, this is only example where equilibrium price is derived in an oligopoly model with a returning cost.¹⁵

The equilibrium strategy (p, τ) is then found by solving the pair of nonlinear equations (9)–(10), which can typically be done only numerically. For instance, when $s = 0$, so that there are no intrinsic search frictions in the market, solving these equations shows that $p \approx 0.45$ and $\tau \approx 0.06$ and hence a buy-later price of $\hat{p} \approx 0.51$ (which is actually slightly above the monopoly price of 0.5). In this example, although the market has no intrinsic search frictions, firms in equilibrium impose “tariff intermediated” search frictions on consumers via the buy-now discount, which here is about 12% of the buy-later price. By contrast, in a market with $s = \frac{1}{8}$, which is the highest intrinsic search cost which induces consumers to participate, one can check that the (exact) solution to this pair of equations is $p = \frac{1}{2}$ and $\tau = 0$, so that there is no buy-now discount. (When $s = \frac{1}{8}$, search costs are so high that consumers will accept the first offer which yields them a non-negative surplus. In particular, there are no returning consumers even with costless recall.)

More generally, the equilibrium buy-now discount τ decreases with the search cost s . That is, the higher is the intrinsic search cost, the less incentive firms have to deter consumers from searching on. This can be seen from Figure 2A below, which depicts how the buy-later price $\hat{p} = p + \tau$ (the upper solid curve) and the buy-now price p (the middle solid curve) vary with s . As is expected, the buy-now price increases with the search cost. Less expected is the observation that the buy-later price depends non-monotonically on s (and is always above the monopoly price).

Further details of the equilibrium are provided in this formal result.¹⁶

Proposition 1 *Given $s \in [0, \frac{1}{8}]$, the system of (9) and (10) has a solution $(p, \tau) \in [0, \frac{1}{2}] \times [0, \frac{1}{8} - s]$, and the solution satisfies $p \leq z_\tau$.*

¹⁵As mentioned in the Introduction, Janssen and Parakhonyak (2008) examine consumer behaviour in a model where consumers incur a recall cost. Although their model is in other respects more general than ours, they do not model firms’ price choices.

¹⁶As discussed, τ decreases with p in (8) and p increases with τ in the range $0 \leq \tau \leq \frac{1}{8} - s$ in (10), and so there is a *unique* equilibrium with $(p, \tau) \in [0, \frac{1}{2}] \times [0, \frac{1}{8} - s]$. Numerical simulations suggest that in the wider region of $(p, \tau) \in [0, 1]^2$ there are no other equilibria.

This analysis assumes that firms cannot distinguish fresh from incoming customers and so must charge both groups the same price. In section 4 we suppose instead that firms can distinguish all three customer groups, in which case it turns out that firms will offer incoming consumers the same price as the returning consumers, a price which is higher than that offered to fresh customers. Thus, if feasible, firms have an incentive to discriminate against customers who have investigated the rival firm (i.e., the returning and incoming customers), rather than discriminate between returning customers and those customers who come to the firm for the first time.

3.2 Comparison with uniform pricing

In the more familiar search model where firms can only charge a uniform price p_0 , consumers face no returning cost and so expression (10) implies that the first-order condition for the price is

$$\frac{1}{p_0} - p_0 = 1 + a_0 . \quad (11)$$

(Unlike the discriminatory prices, this uniform price can be explicitly computed.) Recall that the price which solves (10) increases with the returning cost τ for $\tau \in [0, \frac{1}{8} - s]$. We can therefore deduce the following result immediately:

Proposition 2 *Price discrimination over returning consumers leads to higher prices, i.e., $p_0 \leq p \leq \hat{p}$.*

That is, even the discounted buy-now price in the discriminatory case is higher than the uniform price, and the ability to offer discounts for immediate purchase drives up *both* prices.¹⁷ The intuition for this comes in two steps. First, a firm's total demand is inelastic with respect to changes in the buy-now discount (perfectly so when match utility is uniformly distributed). Since a firm makes additional profit of τ_i from its returning customers, it surely wishes to set a positive buy-now discount. Second, the buy-now discount adds to the intrinsic search frictions in the market, and this allows firms to charge a higher price. (Relative to the uniform-price case, consumers become less willing to search on, and so the firms' demand is less price elastic.) Figure 2A depicts the three prices, where from the bottom up the three curves represent p_0 , p and \hat{p} , respectively. In particular, when $s = \frac{1}{8}$ the search cost is so high such that no consumers want to search beyond the first sampled firm, there is no returning demand in either regime, and so all three prices coincide.

In the discriminatory case, each firm's equilibrium demand is $1 - p\hat{p}$. This is clearly lower than the counterpart $1 - p_0^2$ in the uniform-pricing case since $p, \hat{p} \geq p_0$. Hence, this

¹⁷It is not unusual that the ability to price discriminate in oligopoly leads to falls in all prices, but cases where all prices rise are less familiar.

form of price discrimination excludes more consumers from the market. In addition, one can show that fresh demand rises, and returning demand falls, when this form of price discrimination is used. This is illustrated in the case $s = 0$ in Table 2 below.

However, whether price discrimination leads to higher profit depends on the magnitude of the search cost. In the non-discriminatory case, each firm's profit is $\pi_0 = p_0(1 - p_0^2)$. In the discriminatory case, from (6) each firm's profit is

$$\pi = p(1 - p\hat{p}) + \frac{1}{2}\tau(z_\tau^2 - p^2) .$$

Figure 2B shows how π_0 (the dashed curve) and π (the solid curve) vary with the search cost s . We see that price discrimination leads to higher profit (i.e., $\pi > \pi_0$) only if the search cost is relatively small. When the search cost is relatively high, price discrimination leads to prices which exclude too many consumers. In these cases, firms are engaged in a prisoner's dilemma: when feasible an individual firm wishes to offer a buy-now discount, but when both do so industry profits fall.¹⁸




Figure 2A: Prices and search cost




Figure 2B: Profits and search cost

Finally, we observe that aggregate consumer surplus and total welfare (measured by the sum of consumer surplus and profit) fall when firms discriminate in this manner. Consumer surplus falls since both prices rise compared to the non-discriminatory regime. (Even if $\hat{p} = p$, i.e., if firms charged returning consumers the buy-now price, consumers would obtain lower surplus in the price-discrimination case since $p \geq p_0$. The buy-later premium $\tau \geq 0$ only adds to their loss.) As far as total welfare is concerned, relative to uniform pricing, price discrimination not only induces suboptimal consumer search (i.e., consumers on average cease their search too early due to the

¹⁸This bears some similarities to situations with competitive bundling. There, a firm often has a unilateral incentive to offer consumers a discount for buying two products rather than one, but when all firms do this industry profits fall. However, in contrast to the current case where the discount relaxes competition and drives prices up, with bundling the discount intensifies competition and drives prices down. For instance, see Armstrong and Vickers (2010) for more details.

buy-now discount), but also excludes more consumers from the market, both of which harm efficiency.¹⁹

Our analysis of buy-now discounts so far has assumed that the match utility is uniformly distributed. It is possible to derive equilibrium prices in non-uniform examples by calculating the measure, rather than simply the area, of the regions in Figure 1. As a robustness check on our results, we report numerical calculations for the equilibrium tariff in examples where the density function f is linear rather than constant. Specifically, suppose that the density takes the form $f(u) = 2\beta u + 1 - \beta$, where $\beta \in [-1, 1]$, so that the density function is a straight line with slope 2β passing through the point $(\frac{1}{2}, 1)$. Table 1 reports the equilibrium prices for various values of β , assuming that the search cost s is zero.

	$\beta = -1$	$\beta = -0.5$	$\beta = 0$	$\beta = 0.5$	$\beta = 1$
p_0	0.312	0.382	0.414	0.405	0.360
p	0.313	0.392	0.450	0.474	0.470
τ	0.017	0.032	0.060	0.091	0.124
$\tau/(p + \tau)$	0.05	0.075	0.12	0.16	0.21

Table 1: Equilibrium prices when utilities are not uniformly distributed ($s = 0$)

Here, the first row reports the equilibrium non-discriminatory price, the second and third rows report the buy-now price and discount, while the final row reports the buy-now discount as a fraction of the buy-later price. The central column with $\beta = 0$ is the uniform case previously discussed. As expected, when the density is increasing, the incentive to set a buy-now discount is reinforced by an additional strategic effect: as shown in Lemma 2, with an increasing density a firm's total demand is boosted if it makes it costly to return. Thus, we see that the size of the buy-now discount increases with β , both in absolute terms and as a proportion of the buy-later price. Nevertheless, over all linear density functions, the equilibrium buy-now discount remains positive. Notice also that all prices are higher with price discrimination than without discrimination, even when the density is decreasing. In particular, consumers in aggregate are worse off with this form of price discrimination.

3.3 Buy-now discounts without commitment

We discuss next whether buy-now discounts can still emerge as an equilibrium outcome if we relax the assumption that a firm can commit to its buy-later price when consumers first visit. We consider two cases: one with partial commitment in which firms can commit to a returning purchase price cap (but cannot commit to a specific price), and

¹⁹The proof of Proposition 3 in Armstrong, Vickers, and Zhou (2009) can be adapted to prove this welfare result formally.

the other with no commitment at all. The first case may be reasonable in situations where firms can post a “regular” price. For example, the price printed on the price label in a store usually has this kind of commitment power. The second case may be more relevant in bargaining situations where it is infeasible to sign a contract for the returning purchase price. In both cases, we argue that a buy-now discount can still arise in equilibrium.

Partial Commitment: Here, there is an equilibrium with the same outcome as in the full commitment case. Specifically, in this equilibrium, firms charge a buy-now price p , commit to a buy-later price cap \hat{p} , and actually charge returning consumers \hat{p} , where both p and \hat{p} take the same values as the commitment case analyzed in section 3.1. To sustain this equilibrium, we assume that all consumers believe that for any (maybe off-equilibrium) committed price cap \hat{p}_i the firm’s actual buy-later price will be \hat{p}_i .

To prove this claim, it suffices to observe that when consumers hold the above belief, they will return to the first sampled firm, say firm i , only if $u_i \geq \hat{p}_i$, where \hat{p}_i is firm i ’s buy-later price cap. This implies that firm i can have no incentive to charge them a price below \hat{p}_i . (It may have an incentive to raise the price above \hat{p}_i , but that is not permitted given that the firm commits to this cap.) This in turn fulfills consumer beliefs. Thus, the price cap can be used as a full commitment device.²⁰

No Commitment: We turn next to the case where firms can make no credible commitments about their buy-later prices when consumers first visit. Here, and unlike the rest of this paper, it makes an important difference whether or not consumers face an intrinsic returning cost when they come back to a firm. If consumers face no such cost (as we assumed in the rest of the paper) it is possible to construct equilibria of the constant markup form which are qualitatively similar to the commitment prices: consumers anticipate that a firm will set a return price $\hat{p}_i = p_i + \tau$ if that firm offers the buy-now price p_i , and given these expectations firms have no incentive to set a different buy-later price. (There is a continuum of such τ which can make up an equilibrium of this form, and which induce some consumers to return to their initial firm.) There is also an “exploding offer” equilibrium: consumers anticipate that firms offer very high prices if they return and so never return, and given this behaviour firms have no incentive to reduce their return prices.

However, if there is even a small intrinsic returning cost, say $r > 0$, it turns out that the *only* equilibrium involves exploding offers where the returning demand segment is shut down altogether. To see this, suppose in some equilibrium that each consumer forecasts that a firm’s buy-later price is $\hat{p}(p_i)$ when its buy-now price is p_i , where $\hat{p}(\cdot)$ can take any form. Suppose that the buy-now price in this equilibrium is p^* , say, and

²⁰There may exist other equilibria involving different consumer beliefs.

suppose—contrary to the claim—there is some returning demand in this equilibrium. (Formally, from Figure 1 this implies that $r + \hat{p}(p^*) - p^* < V(p^*)$, where $r + \hat{p}(p^*) - p^*$ is a consumer’s total returning cost.) But if a consumer returns to firm i after first sampling firm i and then firm j , it follows that her match utility satisfies $u_i \geq \hat{p}(p^*) + r$, since the consumer must also pay the returning cost $r > 0$. Since all its returning customers have match utility at least $\hat{p}(p^*) + r$, the firm’s optimal price for these customers must be at least $\hat{p}(p^*) + r$, which contradicts the assumption that $\hat{p}(p^*)$ was the correctly anticipated buy-later price.

Thus, when there an intrinsic returning cost, no matter how small, consumers anticipate that the buy-later prices will be so high that it is never worthwhile to return to the first firm after leaving it. In effect, firms are forced to make exploding offers, and consumers have just one chance to buy from any firm. This result is analogous to Diamond’s (1971) paradox, showing how a small search cost can cause a market to shut down. Diamond’s result relies on consumers knowing their match utility in advance, and one major advantage of Wolinsky’s formulation with *ex ante* unknown match utilities is that this paradox can be avoided. But even in our Wolinsky-type framework, the *returning* consumers know their match utility, and so the returning market fails for the same reason as the primary market failed in Diamond’s framework.

3.4 Non-monetary returning costs and exploding offers

The previous section suggested that exploding offers would often arise when firms cannot commit to their buy-later prices. In this section we discuss an alternative rationale for exploding offers. In the buy-now discount model, the discount τ induces a returning cost, but it also directly affects firms’ revenue since firms gain revenue from the high prices their returning consumers must pay. In other situations it is plausible that firms must offer the same price to all consumers, but can affect the cost of returning with various non-price means.²¹ The case of exploding offers, where aggressive sellers force new visitors to buy “now or never”, is represented by the extreme case of a very large returning cost. In these situations when firms must offer the same price to all consumers, the only reason to increase the returning cost is the “strategic” motive to deter a consumers from searching further.

Suppose that each firm chooses a returning cost τ_i and a uniform price p_i . Suppose firms can freely choose and $\tau_i \geq 0$ without incurring additional costs. Since the strategic motive to make it hard to return is paramount in this setting, suppose in this section that the match utility u comes from a general distribution distribution F rather than the special case of the uniform distribution. (Recall that the strategic motive is absent

²¹For example, online sellers can ask customers to log on to their accounts or input information again; firms can ask consumers to queue or make another appointment if they want to come back.

with a uniform distribution.) To calculate the symmetric equilibrium strategy (p, τ) , suppose that firm i chooses (p_i, τ_i) while its rival follows the equilibrium strategy (p, τ) . Firm i 's profit is just p_i multiplied by its total demand. Thus, a firm's returning cost τ_i is chosen simply to maximize its total demand given its own price and its rival's price and returning cost.

From Lemma 2 we can deduce: (i) if u has an increasing density, firms will choose a returning cost as high as possible, and in equilibrium there is no returning demand, and (ii) if u has a decreasing density, firms will choose the lowest possible returning cost $\tau_i = 0$. In case (i) the price p is the equilibrium price in a search model without the possibility of consumer recall, while in case (ii) p is the equilibrium price in a search model with costless recall. Thus, when the density of the match utility is increasing the equilibrium involves each firm forcing fresh consumers to buy "now or never", while when the density decreases firms allow consumers to return freely if they do not wish to buy immediately.²²

In the particular case of the uniform distribution, where $f \equiv 1$, given a particular price p firms are indifferent between all levels of the returning cost, and altering τ_i merely shifts returning demand to fresh demand in a one-for-one manner. In this case there are multiple equilibria. Consider for simplicity the case of costless search, so that $s = 0$. Then when consumers can freely return to their initial firm (so that $\tau = 0$), one can check from (10) with $\tau = s = 0$ that the equilibrium price is $p = \sqrt{2} - 1 \approx 0.41$. By contrast, if new consumers are forced to buy "now or never" ($\tau \geq V(p)$), one can check from (10) with $s = 0$ and $\tau = V(p)$ that the equilibrium price solves $p^3 + 2p - 1 = 0$ so that $p \approx 0.45$. Either of these strategies make up a symmetric equilibrium (and there are interior symmetric equilibria as well, where firms make it costly to return but not prohibitively so). Clearly, consumer surplus is higher in the free-return equilibrium (since consumers have the option of buying "now-or-never" even when they are not forced to). It turns out that industry profits are also a little lower with high-pressure selling: the price is higher than with free return, but significantly fewer consumers then buy either product. In this example then, a laissez-faire market may exhibit high-pressure sales techniques being employed in equilibrium, while both consumers and firms would be better off if such practices were prohibited.²³

Table 2 summarizes the outcomes with $s = 0$ and a uniform distribution for match utilities in the three situations we have now encountered. When consumers can freely return to their initial firm, the equilibrium price (buy-now or buy-later) is 0.41. With

²²If the density is non-monotonic, it is possible that firms will adopt an intermediate returning cost, which is positive but still induces some consumers to return.

²³The use of exploding offers could be prohibited by mandating a "cooling off period", so that consumers have the right to return a product in some specified time after agreeing to purchase. (They could then return a product if they subsequently find a preferred option.) Many jurisdictions impose cooling off periods for some products, especially those sold in the home.

costless search and return no consumer will ever choose to buy from the first firm immediately, and so fresh demand is zero. In equilibrium, firms set the same price and so each firm’s total demand is split equally between returning and incoming consumers. When firms offer a buy-now discount, section 3.1 shows that the buy-now discount is about 12% of the buy-later price, and as a result there is significant fresh demand and less returning demand. Finally, when exploding offers are used, there is no returning demand at all.

	p	\hat{p}	fresh	returning	incoming	excluded
no discount	0.41	0.41	0%	41%	41%	17%
with discount	0.45	0.51	29%	11%	37%	23%
exploding offer	0.45	n/a	40%	0%	33%	27%

Table 2: The impact on prices and demand of buy-now discounts and exploding offers

4 Other Forms of Search-Based Discrimination

In this section we consider alternative methods of price discrimination in this market. In some situations, firms may be able to distinguish more finely between consumers, and can identify fresh from incoming consumers. (For instance, an online firm may be able to track whether a consumer has previously paid it a visit *and* whether she has already visited the rival seller.) What are equilibrium prices when firms can distinguish all three types of consumers? We continue to assume a uniform distribution for match utilities, that the search cost lies in the range $s \in [0, \frac{1}{8}]$, and that firms can commit to their buy-later prices when a consumer visits for the first time. Let p , \hat{p} , and P be the corresponding prices for fresh consumers, returning consumers, and incoming consumers, respectively. Then the pattern of demand for those consumers who first visit firm i is just as depicted in Figure 1, except that firm j ’s prices are changed from its buy-now prices (p, p_j) to its incoming prices (P, P_j) . As the following proposition shows, in this setting there continues to be an equilibrium with active returning purchase and a buy-now discount, but now the discount applies only to a firm’s *fresh* consumers.

Proposition 3 *When firms can distinguish all three groups of consumers, there is a unique equilibrium with $\tau \in [0, \frac{1}{8} - s]$ and $P = \hat{p} = p + \tau \leq \frac{1}{2} \leq z_\tau$.*

A little surprisingly, in equilibrium firms charge exactly the same price to returning consumers and incoming consumers.²⁴ Thus, in equilibrium a firm discriminates equally against all consumers who have investigated the rival seller, regardless of whether the consumer first sampled it or the rival seller. It is clear that firms have a motive to

²⁴This result also holds for general distributions, and so it is not driven by the assumption of a uniform distribution.

discriminate against consumers who have investigated the rival, since by revealed preference such consumers do not much care for the rival’s product. But it is less clear why a firm discriminates *equally* against all such consumers. After all, a returning consumer has also revealed that she does not much care for the firm’s product either (for otherwise she would have purchased immediately).

Whenever a consumer samples the second product, that consumer makes her choice under complete information and without search frictions: she knows both of her match utilities and both prices, and she needs to incur no further search costs to buy either product. There are two opposing factors affecting the first seller’s choice of buy-later price. First, the first seller offers a lower match utility on average than its rival, and hence its buy-later price should be lower than the rival’s (incoming) price for these consumers. Second, raising its buy-later price reduces a consumer’s incentive to search further, which therefore allows the firm compete for fewer consumers in a disadvantageous situation. For reasons which are not yet fully clear, these two forces cancel out exactly, and the equilibrium buy-later and incoming prices coincide.




Figure 3: Prices with full discrimination

Figure 3 describes the relative positions of the fresh demand price p (the lower solid curve), the returning/incoming demand price $\hat{p} = P$ (the upper solid curve), and the uniform price p_0 (the dashed curve, already depicted on Figure 2A). Relative to the uniform price, the price for returning and incoming consumers becomes higher, but the price for fresh consumers may be higher or lower, depending on the magnitude of the search cost. Note also that, in contrast to the case in Figure 2A, here the non-discounted price depends monotonically on the search cost and never exceeds the monopoly level of 0.5. However, numerical simulations suggest that the impacts of price discrimination on profit, consumer surplus, and total welfare are similar to our main model.

The consequence of relaxing the ability to commit to the buy-later price is also similar to our main model. If firms cannot commit to the buy-later price at all, and if

there is even a small exogenous returning cost, then in equilibrium consumers expect a prohibitively high buy-later price and never return to the first sampled seller. Then each firm acts as a monopolist over its incoming consumers and charges them the monopoly price $\frac{1}{2}$. This further implies that a consumer samples the second seller if and only if the first product's net surplus is below $V(\frac{1}{2})$, and so the price offered to fresh consumers is $\frac{1}{2}[1 - V(\frac{1}{2})] \approx 0.44$ if $s = 0$.

Other forms of price discrimination could also be considered, depending on which consumer groups can be separately identified by firms. We have seen in Proposition 3 that when firms can separately identify all three groups, they do not set distinct prices to incoming and returning customers even though they could do so. It follows that in the alternative setting where firms *cannot* distinguish these two groups of customers the equilibrium prices are as described in Proposition 3 and Figure 3. (It is not clear, however, how relevant this particular information structure is for real markets.)

A perhaps more relevant situation is where firms can identify incoming consumers, but cannot distinguish fresh from returning customers. This implies that a firm sets price p for those consumers who sample it first (i.e., for both fresh and returning consumers), and price P for the incoming consumers who have sampled the rival first. Since fresh consumers and returning consumers pay the same price, consumers no longer face a returning cost. This case might also apply when a firm is somehow prevented from offering different prices to fresh and returning customers.²⁵

This variant essentially replicates the model in Armstrong, Vickers, and Zhou (2009), specialized to duopoly. (In the earlier paper, all consumers visited a prominent firm first, while in the current paper consumers have a random search order. But the assumption that firms can observe the search order makes the two situations the same.) Armstrong *et al.* (2009, Proposition 1) showed that prices were ordered as

$$p \leq p_0 \leq P ,$$

where p_0 is the equilibrium uniform price in (11). In other words, all consumers who buy the first sampled product will pay less than before, while all consumers who buy the second sampled product will pay more than before. This is because the incoming consumers for each firm must be those consumers who were relatively unsatisfied with the rival's product, and so the firm holds greater market power and can afford to set a high price. Separating these consumers makes the remaining demand more price sensitive, and so the associated price becomes lower. For comparison with our previous figures, Figure 4 below depicts how these three prices vary with the search cost, where

²⁵For instance, consider a consumer buying an item online. If she first investigates firm i , that firm offers her a price and match utility knowing that the consumer has not yet visited the rival seller. The consumer can keep firm i 's screen (and its price offer) open while she goes on to investigate the rival. She can then return to i if dissatisfied with j , and pay its original price.

the middle dashed curve is p_0 (already shown on Figures 2A and 3 above), and the lower and upper solid curves are p and P , respectively. Unlike previous cases in this paper, the three prices coincide when $s = 0$. When search and recall is costless, consumers will always sample both products, and so whether or not consumers have first seen the rival product reveals no useful information with which to set discriminatory prices.




Figure 4: Prices when incoming consumers can be identified

The welfare impacts of price discrimination over incoming consumers are again similar to our main model: relative to the uniform-pricing case, price discrimination lowers both consumer surplus and total welfare, and (at least in our duopoly setting) it boosts industry profits if and only if the search cost is relatively small.

5 Discussion and Extensions

This paper has explored how firms could use their knowledge of consumer search behaviour to refine their pricing decisions, and how this practice could affect market performance. Our main model studied the case in which a firm knows whether a consumer is visiting it for the first time or whether the consumer is returning after a previous visit. We saw that firms have an incentive to offer a lower price on a first visit than a second visit, so that new consumers are offered a “buy-now” discount. The ability to offer such discounts acts to raise all prices in the market, which lowers both consumer surplus and total welfare and may even harm firms themselves when the search cost is relatively high. We also explored firms’s incentives to use exploding offers, so that a new visitor to a firm is forced to make her purchase decision immediately, before she has the chance to investigate the rival offering.

Several extensions to this analysis are worthwhile, including the following:

The impact of prominence: In many circumstances, one seller is more prominent than others, with the result that more consumers sample it first. (Recall that De los Santos (2008) showed how one seller in the online book market attracted a greatly disproportionate share of the initial searches by consumers.) We now discuss how prominence could affect our results in the main model. For instance, will a prominent firm offer a greater or smaller buy-now discount than its rival, or will a prominent firm have a greater incentive to use exploding offers than its rival? (In the variants presented in section 4, firms compete for consumers with different search orders separately, so prominence does not affect prices at all.)

In the non-monetary returning cost model presented in section 3.4, each firm just chooses its returning cost to maximize the sum of its fresh demand and returning demand (recall that a firm's incoming demand is not affected by its returning cost). Given prices, this decision is independent of how many consumers sample the firm first. Hence, at least for monotonic utility density functions (in which cases the choice of returning cost is independent of prices), prominence does not alter firms' incentive to choose their returning cost. In particular, if the density for match utility is increasing, a firm has an incentive to make exploding offers regardless of how prominent it is.

Turning to the buy-now discount model, consider first the extreme case where *all* consumers first visit one of the firms, say firm A . Thus, firm A has only fresh and returning demand, with respective prices p and \hat{p} , while firm B only has incoming demand with price P . It is then clear that the equilibrium choices for these three prices are exactly as described in Proposition 3 when full discrimination is possible. (When calculating equilibrium prices in the fully discriminating case, the share of consumers who first sample one firm makes no difference to the analysis, since the segments of consumers who first sample A and consumers who first sample B are treated entirely separately by the firms.) Hence, when one firm is prominent, its buy-now price is the lower solid line on Figure 3 and its buy-later price is the higher solid line, which is also the (single) price offered by the less prominent firm. Since each of the prominent firm's prices are (weakly) lower than the its rival's price, consumers have a strict incentive to investigate the prominent firm first, even if they have a choice over their search order.²⁶

However, this extreme case cannot shed light on the less prominent firm's incentive to offer buy-now discount (since it has no opportunity to do so when its only demand comes from incoming consumers). In a less extreme setting where $1 + \theta$ (with $\theta < 1$) consumers sample the prominent firm first, one can show that the prominent firm's buy-now discount decreases with θ while the less prominent firm's buy-now discount increases with θ . That is, the less prominent firm will actually offer a deeper buy-now

²⁶In particular, the buy-now discount model with *ex ante* symmetric firms also has two *asymmetric* equilibria (in addition the symmetric equilibrium discussed in section 3) in which all consumers sample a certain firm first and this firm then provides better deals.

discount than its rival. For example, if $s = 0$, the prominent firm's buy-now discount decreases from 0.06 to 0.053 as θ varies from zero to one, while the less prominent firm's buy-now discount increases from 0.06 to 0.063.²⁷

More ornate schemes: In our buy-now discount model, sellers may be able to extract more surplus from buyers by offering them an additional option—namely, buyers can pay a deposit d for the option to return and buy at a specified price q .²⁸ With this new option, more consumers may opt to search on, and among the consumers who do search on, those having relatively high valuations of the first product will buy the deposit contract while others having relatively low valuations will not since they rarely come back. In the uniform distribution case, one can show numerically that: (i) starting from the equilibrium in our main model, each firm has a strict incentive to introduce a deposit contract; (ii) the purchase price in the deposit contract is even lower than the buy-now price (i.e., $q < p$) but the consumers who buy the deposit contract and eventually come back pay more than fresh consumers (i.e., $d + q > p$); (iii) with the new instrument firms earn less in equilibrium. This extension could be extended further, so that firms offer a *menu* of deposit contracts (a bigger deposit would grant the right to come back and buy at a lower price).

Using search behaviour to infer search cost: When consumers have heterogenous search costs, which plausibly will be private information initially, consumer search behaviour may be informative of search costs. The consumers who buy the first sampled product immediately have (on average) relatively high search costs, which may induce firms to charge them a high price when they are distinguishable. Meanwhile, those consumers who keep searching have (on average) relatively low search costs, which may induce firms to charge them a low price.

To illustrate in a stark model, consider the situation in which firms can distinguish all three groups of consumer (fresh, returning and incoming) and so charge each group a distinct price. Suppose that products are homogeneous, and each consumer is willing to pay up to v for a unit. Suppose there are two possible costs of search: some consumers have costless search (i.e., $s = 0$), while the others have extremely high costs of searching the second firm, and so will immediately buy from the first firm they visit provided that that firm's price is below v . Then the equilibrium in this market involves each firm setting the buy-now price equal to monopoly price v , and the incoming and returning

²⁷This is because, as we pointed out in footnote 12, in equilibrium a firm's buy-now discount decreases with its rival's buy-now price. When a firm becomes more prominent, it will charge a lower buy-now price as it has fewer incoming consumers (over whom it has more monopoly power), which in turn induces its rival to offer a greater buy-now discount.

²⁸For example, some business schools demand a deposit from applicants who want to keep the admission offer for a longer time. It is also a common practice in many business-to-business transactions.

price equal to marginal cost (which is zero).²⁹ Thus, firms will attempt to exploit the inert consumers with a high buy-now price, but once a consumer reveals she is willing to shop around, she is offered a competitive deal.

Consumers' incentives to conceal/reveal their search history. Notice that from an *ex ante* perspective, in all our models a consumer will be better off if she can conceal her search history, so that firms are forced to set the non-discriminatory price.³⁰ Thus, if it is costless to pretend to be a new visitor (e.g., by deleting cookies on your computer), all consumers will do this, and the market will operate as a standard search market with uniform prices (as in Wolinsky's framework). But if there are some costs involved in concealing search history, or if some consumers do not think to do so, there will remain an incentive to condition prices on observed search history.

Consumers may also have incentive to (selectively) reveal their search history. For instance, they may want to force the current seller to offer a better deal by providing hard information of a previous price offer (if this is possible).³¹ Investigating how such a possibility could affect price competition and market performance is an interesting topic for future investigation.

APPENDIX

Proof of Lemma 1: We calculate the consumer's optimal strategy by backward induction. Denote by $v(u_i)$ the consumer's expected continuation payoff if she decides to investigate firm j when her match utility from firm i is u_i . For convenience, write $\hat{p}_i = p_i + \tau_i$ for the total expected cost of buying from firm i if a consumer first leaves and then returns to the firm. Thus, a returning customer for firm i expects to pay \hat{p}_i and so she will never do this if $u_i < \hat{p}_i$. Therefore, $v(u_i) = V(p)$ when $u_i < \hat{p}_i$. On the other hand, if $u_i \geq \hat{p}_i$ then she will return to i after sampling j if $u_i - \hat{p}_i > u_j - p_j$, where p_j is firm j 's actual (not expected) price. Since the consumer expects j to charge p , her expected payoff from investigating j is then

$$\int_0^{p+u_i-\hat{p}_i} (u_i - \hat{p}_i) dF(u_j) + \int_{p+u_i-\hat{p}_i}^{u_{\max}} (u_j - p) dF(u_j) - s = V(p) + \int_0^{u_i-\hat{p}_i} F(p+x) dx, \quad (12)$$

²⁹The argument is similar to Armstrong, Vickers and Zhou (2009, section 4).

³⁰This is somewhat related to models of sequential bargaining. If one buyer is negotiating with a sequence of sellers, then the buyer may gain from keeping the order (and the outcome) of negotiations secret from sellers. Noe and Wang (2004) present such a model, and find that when the objects sold are complements for the buyer, then the buyer obtain greater surplus when he randomizes and conceals the order in which he approaches the sellers.

³¹In the model in Daughety and Reinganum (1992), if a consumer makes contact with two sellers, she can force the sellers to compete in Bertrand fashion and force her price down to marginal cost.

where the equality follows after integrating by parts and the definition of V in (1). In sum, the consumer's expected payoff if she decides to investigate firm j is

$$v(u_i) = \begin{cases} V(p) & \text{if } u_i \leq \hat{p}_i \\ V(p) + \int_0^{u_i - \hat{p}_i} F(p+x)dx & \text{if } u_i \geq \hat{p}_i . \end{cases}$$

Note that $v(\cdot)$ is a (weakly) increasing function with slope no greater than 1.

Clearly, the consumer will decide to investigate firm j if and only if

$$u_i - p_i < v(u_i) . \quad (13)$$

Since (13) holds when $u_i = 0$, and the left-hand side has slope 1 while the right-hand side has slope less than 1, there is at most one point, say u^* where the curves in (13) cross so that $u^* - p_i = v(u^*)$, and (13) holds if and only if $u_i < u^*$. If the curves never cross, so that (13) always holds for $u_i \leq u_{\max}$, then the price p_i is so high that the consumer will always continue onto the rival no matter how high the match utility from i is. Otherwise the curves will cross, and u^* is the reservation utility level for continuing onto j .

The curves cross at the “flat” portion of $v(\cdot)$, i.e., $u^* < \hat{p}_i$, if and only if $\tau_i = \hat{p}_i - p_i \geq V(p)$. When the three prices satisfy this inequality, the consumer will continue onto j if and only if $u_i - p_i < V(p)$. Moreover, the consumer will never return to i if she leaves it, no matter how low is u_j or how high the actual (as opposed to the expected) p_j is. (To check this, observe that her payoff if she returns to i after leaving it is $u_i - \hat{p}_i$, and since she chooses to continue to j we have $u_i - \hat{p}_i < V(p) + p_i - \hat{p}_i$ which is negative given that $\hat{p}_i - p_i \geq V(p)$.) This proves part (ii) of the lemma.

Now suppose $\tau_i = \hat{p}_i - p_i < V(p)$, so that the curves cross on the increasing portion of $v(\cdot)$. In this case, from the expression for the continuation payoff on the left-hand side of (12) the threshold utility u^* satisfies

$$\begin{aligned} u^* - p_i &= v(u^*) \Leftrightarrow \int_{p+u^*-\hat{p}_i}^{u_{\max}} [u_j - (p + u^* - \hat{p}_i)]dF(u_j) - s = \tau_i \\ &\Leftrightarrow V(p + u^* - \hat{p}_i) = \tau_i \\ &\Leftrightarrow z_{\tau_i} = p + u^* - \hat{p}_i \\ &\Leftrightarrow u^* = a_{\tau_i} + p_i - p . \end{aligned}$$

In sum, when $\tau_i < V(p)$ the consumer should continue on to investigate the second product if and only if $u_i - p_i < a_{\tau_i} - p$. And if she does continue to j , she should return to i whenever $u_i - \hat{p}_i > u_j - p_j$ where p_j is firm j 's actual price. This completes the proof of part (i).

Proof of Lemma 2: As discussed in the text, changing τ_i has no impact on firm i 's incoming demand and so the impact on the firm's total demand is the sum of the impacts on its fresh and its returning demand. Suppose that $p_j = p$ on Figure 1, and also that $\tau_i < V(p)$ so that there is a region of returning demand. Then by inspecting Figure 1 the impact of a small increase in τ_i on the sum of fresh and returning demand is proportional to

$$-f(p_i + \tau_i)F(p) - f(a_{\tau_i} + p_i - p)(1 - F(z_{\tau_i}))\frac{da_{\tau_i}}{d\tau_i} - \int_p^{z_{\tau_i}} f(u_j)f(p_i + \tau_i + u_j - p)du_j$$

where the first term reflects the shrinkage of returning demand from the left, the second term reflects the net expansion of fresh demand and the shrinkage of returning demand from the right, and the last integral term represents the impact of shifting the diagonal line in Figure 1 to the right. From (4) and integration by parts, the above expression equals

$$\int_p^{z_{\tau_i}} F(u) f'(u + p_i - p + \tau_i) du .$$

Since the firm's total demand is continuous in τ_i and does not depend on τ_i when $\tau_i \geq V(p)$, all results in the Lemma follow immediately.

Proof of Proposition 1: Given the discussion in the text, the only part of this result which remains to be proved is that there is a pair (p, τ) in the rectangle $[0, \frac{1}{2}] \times [0, \frac{1}{8} - s]$ which solves (9) and (10). Let $0 < p_\tau \leq \frac{1}{2}$ be the p which solves (10) given τ . Now consider (9) with p replaced by p_τ . We show that this has a solution $\tau \in [0, \frac{1}{8} - s]$. First, if $\tau = 0$, the left-hand side becomes $z_0^2 - p_0^2$ which is positive because $z_0 \geq \frac{1}{2}$ for $s \leq \frac{1}{8}$, while $p_0 \leq \frac{1}{2}$. Second, if $\tau = \frac{1}{8} - s$ then $z_\tau = \frac{1}{2}$ and the left-hand side of (9) is negative if

$$\frac{1}{4} - p_\tau^2 \leq 2\tau . \tag{14}$$

When $z_\tau = \frac{1}{2}$, (10) implies $\frac{1}{p_\tau} - p_\tau - \frac{3}{2} = 2\tau$. Hence, (14) is equivalent to

$$p_\tau^3 - p_\tau^2 - \frac{7}{4}p_\tau + 1 = (p_\tau - \frac{1}{2})(p_\tau^2 - \frac{p_\tau}{2} - 2) \geq 0 ,$$

which is true for $p_\tau \in (0, \frac{1}{2}]$.

Proof of Proposition 3 Suppose firm i deviates to (p_i, \hat{p}_i, P_i) while firm j follows the equilibrium strategy (p, \hat{p}, P) . Let $\tau_i \equiv \hat{p}_i - p_i$ and $\tau = \hat{p} - p$. Then firm i 's total demand can be calculated to be

$$Q_T = \underbrace{1 - (p_i + a_{\tau_i} - P)}_{\text{fresh demand}} + \underbrace{\frac{1}{2}(z_{\tau_i}^2 - P^2)}_{\text{returning demand}} + \underbrace{(a_\tau + p - P)(1 - P_i) - \frac{1}{2}(z_\tau - P)^2}_{\text{incoming demand}} .$$

To understand this expression, note that for a consumer who samples firm i first, her choice is the same as in the buy-now discount case, except that she now expects that firm j will charge her P instead of p if she searches on. This explains the fresh and returning demands above. (Notice that the condition for active returning purchase in equilibrium is now $P < z_\tau$.) The expression for i 's incoming demand can be derived from calculating j 's incoming demand on Figure 1, but permuting the two firms, i.e., replace p_j with i 's actual incoming price P_i , replace τ_i with τ , and replace the “buy now” threshold of $a_{\tau_i} + p_i - p$ with $a_\tau + p - P$.

It follows that firm i 's profit is

$$p_i[1 - (p_i - P + a_{\tau_i})] + \frac{1}{2}(p_i + \tau_i)(z_{\tau_i}^2 - P^2) \\ + P_i \left\{ (a_\tau + p - P)(1 - P_i) - \frac{1}{2}(z_\tau - P)^2 \right\} . \quad (15)$$

This profit is additively separable in the three tariff parameters (p_i, τ_i, P_i) . Similarly to the buy-now discount model, τ_i does not affect the sum of i 's fresh and returning demand, and so τ_i is chosen to maximize $\tau(z_\tau^2 - P^2)$, which has first-order condition

$$z_\tau^2 - P^2 - \frac{2\tau z_\tau}{1 - z_\tau} = 0 . \quad (16)$$

Second, p_i only affects fresh and returning demand, and so the first-order condition for p is

$$1 - (2p - P + z_\tau + \tau) + \frac{1}{2}(z_\tau^2 - P^2) = 0 . \quad (17)$$

Finally, P_i only affects firm i 's incoming demand, and the first-order condition for P is

$$(z_\tau + \tau + p - P)(1 - 2P) - \frac{1}{2}(z_\tau - P)^2 = 0 . \quad (18)$$

As in the buy-now discount model, these first-order conditions are also sufficient.

In the following, we first show that the solution to (16)–(18), if it exists, must satisfy $P = p + \tau$. In a second step, we then show from (16)–(17), that there exists a unique solution with $(P, \tau) \in [0, \frac{1}{2}] \times [0, \frac{1}{8} - s]$.

Let $\theta \equiv P - p$. We can rewrite equations (17) and (18) as

$$1 - 2p - \tau = z_\tau - P - \frac{1}{2}(z_\tau^2 - P^2)$$

and

$$1 - 2P = \frac{(z_\tau - P)^2}{2(z_\tau + \tau - \theta)} ,$$

respectively. Subtracting the second of these from the first yields

$$2\theta - \tau = \frac{z_\tau - P}{2} \left(2 - (z_\tau + P) - \frac{z_\tau - P}{z_\tau + \tau - \theta} \right) .$$

Rewrite (16) as

$$\tau = \frac{1 - z_\tau z_\tau^2 - P^2}{z_\tau},$$

and subtracting it from the above equation yields

$$2(\theta - \tau) = \frac{z_\tau - P}{2} \left(2 - \frac{z_\tau - P}{z_\tau + \tau - \theta} - \frac{z_\tau + P}{z_\tau} \right) = \frac{(z_\tau - P)^2(\tau - \theta)}{2z_\tau(z_\tau + \tau - \theta)}.$$

Hence, the solution if exists must satisfy $\tau = \theta$, i.e., $P = \hat{p} = p + \tau$.

We now show existence and uniqueness. Let us focus on (16)–(17). First, we have already shown that in equation (16) τ decreases with P . Second, replace p in (17) by $P - \tau$ and then the two equations imply

$$P = 1 - z_\tau + \frac{\tau}{1 - z_\tau},$$

which increases with τ . Hence, the solution if exists must be unique. Substituting this expression for P into (16) yields

$$z_\tau^2 - \left(1 - z_\tau + \frac{\tau}{1 - z_\tau} \right)^2 = \frac{2\tau z_\tau}{1 - z_\tau}.$$

If $\tau = 0$, then $z_\tau = 1 - \sqrt{2s} \geq \frac{1}{2}$ since $s \leq \frac{1}{8}$. So the left-hand side is positive, greater than the right-hand side which is zero. If $\tau = \frac{1}{8} - s$, then $z_\tau = \frac{1}{2}$. So the left-hand side is negative, less than the right-hand side which is positive. Hence, this equation has a solution $\tau \in [0, \frac{1}{8} - s]$. Finally, given that τ belongs to this interval, it follows from (16) that $P \leq \frac{1}{2}$.

References

- ACQUISTI, A., AND H. VARIAN (2005): “Conditioning Prices on Purchase History,” *Marketing Science*, 24(3), 1–15.
- ANDERSON, S., AND R. RENAULT (1999): “Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model,” *RAND Journal of Economics*, 30(4), 719–735.
- ARBATSKAYA, M. (2007): “Ordered Search,” *Rand Journal of Economics*, 38(1), 119–126.
- ARMSTRONG, M., AND J. VICKERS (2010): “Competitive Nonlinear Pricing and Bundling,” *Review of Economic Studies*, 77(1), 30–60.
- ARMSTRONG, M., J. VICKERS, AND J. ZHOU (2009): “Prominence and Consumer Search,” *Rand Journal of Economics*, 40(2), 209–233.

- CHEN, Y. (1997): “Paying Customers to Switch,” *Journal of Economics and Management Strategy*, 6(4), 877–897.
- CHU, W., E. GERSTNER, AND J. HESS (1995): “Costs and Benefits of Hard-Sell,” *Journal of Marketing Research*, 32(1), 97–102.
- DAUGHETY, A., AND J. REINGANUM (1992): “Search Equilibrium with Endogenous Recall,” *Rand Journal of Economics*, 23(2), 184–202.
- DE LOS SANTOS, B. (2008): “Consumer Search on the Internet,” NET Institute Working Paper 08-15.
- DECK, C., AND B. WILSON (2006): “Tracking Customer Search to Price Discriminate,” *Economic Inquiry*, 44(2), 280–295.
- DIAMOND, P. (1971): “A Model of Price Adjustment,” *Journal of Economic Theory*, 3, 156–168.
- ELLISON, G., AND A. WOLITZKY (2008): “A Search Cost Model of Obfuscation,” mimeo, MIT.
- FUDENBERG, D., AND J. TIROLE (2000): “Customer Poaching and Brand Switching,” *Rand Journal of Economics*, 31(4), 634–657.
- GAL-OR, E., M. GAL-OR, AND A. DUKES (2007): “Optimal Information Revelation in Procurement Schemes,” *Rand Journal of Economics*, 38(2), 400–418.
- HART, O., AND J. TIROLE (1988): “Contract Renegotiation and Coasian Dynamics,” *Review of Economic Studies*, 55(1), 509–540.
- JANSSEN, M., AND A. PARAKHONYAK (2008): “Consumer Search with Costly Recall,” mimeo.
- NOE, T., AND J. WANG (2004): “Fooling All the People Some of the Time: A Theory of Endogenous Sequencing in Confidential Negotiations,” *Review of Economic Studies*, 71(3), 855–881.
- REYNOLDS, S., AND J. WOODERS (2009): “Auctions with a Buy Price,” *Economic Theory*, 38(1), 9–39.
- ROTH, A., AND X. XING (1994): “Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions,” *American Economic Review*, 84(4), 992–1044.
- VISHWANATH, T. (1989): “Job Search, Stigma Effects, and Escape Rate from Unemployment,” *Journal of Labor Economics*, 7(4), 487–502.

WOLINSKY, A. (1986): "True Monopolistic Competition as a Result of Imperfect Information," *Quarterly Journal of Economics*, 101(3), 493–511.