Word-of-Mouth Recommendations and Discounts in Consumer Search

Nathan Larson
Department of Economics
American University
nlarson@american.edu

July, 2016
Preliminary

Abstract

In a standard differentiated-product sequential search setting (where learning about prices and one's taste for different products is costly), we ask what happens when when consumers already know how much they like some of the products on offer. But because consumers will tend to go straight to a good known product rather than search, the usual hold-up problem is exacerbated, and prices rise. Under broad conditions, the second effect dominates: giving consumers tips about good products depresses utility and increases profits.

If consumers learn about their "known" products from friends and family, firms may pass on discounts through those channels in order to draw more directed searchers. However, firms face a commitment problem, as the more directed consumers they attract, the greater the temptation is to hold up consumers when they arrive. The cost of limited commitment means that firms cannot always discount when they would like to. But when they do, "retail" prices rise, and consumers are generally worse-off than without discounting.

1 Introduction

Consider a couple, Alice and Ralph, who wish to find an architect to design a new home. Suppose there are many architects with varied skill sets and artistic sensibilities, and some will be better-suited to the couple than others – we summarize the couple's utility from hiring any particular architect by a *match value*. If Alice and Ralph do not know their match value with a particular architect or how much he would charge for their project, they can learn both by interviewing him (at some cost). However, there may be a few architects whose work they have already seen, and judged, in the homes of their friends and family – for these architects, Alice and Ralph already know their match value, although they still may have to interview them to get price quotes.

If Alice and Ralph have no friends and family, Wolinsky's (1986) differentiated-product search model gives a reasonable description of how they might proceed: sequentially interview architects at random until finding someone they would like to hire. This paper explores what they should do if they *do* have prior information about how well they are matched with a few of the architects (but are uninformed about the rest). We focus on the case in which the couple must still incur a search cost to get a legally binding price quote from these "known" architects. (The idea is that Alice and Ralph may well have some sense of what their friends and family may have paid, but these past prices do not bind the architect when quoting on a new project.) This assumption applies naturally to markets for skilled services (such as home contracting, consulting, and so forth) but is also relevant in any goods market where firms do not commit to long-term prices. Our main focus is on whether access to "free information" from friends and family improves consumer welfare. In partial equilibrium, the answer is rather clearly yes – since Alice and Ralph are free to ignore the additional information, it can only make them weakly better off. However under broad conditions, this result reverses in general equilibrium – the more information from friends and family that consumers have access to, the worse-off they all are. The fault for this is rooted in the Diamond paradox – when a consumer can direct her search toward better matches first, it becomes more attractive to hold her up with a higher than expected price quote.

Next suppose Alice and Ralph have seen the work of one architect – the one who designed their friend Ed's home. The architect might try to drum up business by extending a "friends and family" discount to anyone that Ed refers. Examples like this in which a consumer learns from the same source both about her taste for a product, and about associated discount offers, are commonplace. Motivated by such examples, we extend our model to allow firms to offer discounts to the consumers who know about the firm's product. For simplicity, we focus on the case in which a consumer is familiar with only one product *ex ante*.[1]

Crucially, we assume that a firm cannot commit to a firm price *ex ante* – the assumption is that whatever the nominal price offer may be, a firm always has scope to tack on unannounced fees at the time of purchase. Instead, firms are permitted to offer discounts off of the retail price that a consumer discovers upon arrival. Because of the hold-up problem, a firm would generally like to commit to lower prices in order to attract more directed

[1] If consumers observed their taste and discount offers for several products before searching, then in effect we would have a situation with two stages of price competition, *ex ante* (to attract directed consumers with discounts) and *ex post* (to keep any arrivals, directed or undirected, from leaving without a purchase). We make no claims about results in this case; it is possible that stronger *ex ante* competition might mitigate the anti-competitive flavor of many of our conclusions, but analysis of this case appears challenging.

consumers. However, with limited commitment, a firm may not be able to make a low price stick. If it expects a consumer mix tilted toward high values (because its discount offer has succeeded in luring consumers who know they like its product), then it will be tempted to hike its retail price *ex post*, partially wiping out the promised discount. This price hike implies a sacrifice on profits from *undirected* searchers who do not receive a discount, as the new retail price will be higher than optimal relative to their demand. If a firm anticipates that discounting *ex ante* will tempt it to distort retail prices too much *ex post*, it may foreswear discounts even though it would like to attract more directed consumers.

When discounting is used in equilibrium, consumers will often be worse-off than they would be if discounts were forbidden. Generally speaking, consumers lucky enough to know about products they like well will benefit, while less lucky consumers will suffer under higher retail prices; expected utility tends to be dominated by the second effect. Discounting generally fattens equilibrium profits due the upward drift in the retail price level. And discounting tends to depress social welfare, as the price wedge between known products and unknown ones distorts consumers' search strategies away from the first-best.

One warning, in the interest of avoiding disappointment later: despite the motivating discussion above, word-of-mouth and recommendations from prior purchasers will play no active role in the model. We will simply take it as a primitive of the model that consumers know their taste (and possibly discount offers) for some products; word-of-mouth and recommendations will be no more than a convenient back-story to explain why this assumption is reasonable. Studying the role of word-of-mouth product information in a richer, less perfunctory way would certainly be worthwhile, but it is not our focus here.

Below, in Section 2, we develop our version of the standard sequential search model and characterize equilibrium when consumers are not familiar with any products prior to search. Section 3 analyzes the model when consumers do have information about some products *ex ante*; conclusions about the effect of this information on consumer welfare are presented in Section 4. Section 5 introduces discounts into the model and presents results, and Section 6 concludes. One final caveat: the work presented below is quite preliminary, and the current state of the paper is rough and incomplete. In particular, references and a discussion of related literature are conspicuously absent. Apologies, and caveat lector!

2 Model

The basic setup is similar to Wolinsky (1986).[2] There are an unlimited number of firms and consumers, each modeled as a continuum of size one. Firms sell differentiated products, and each consumer wishes to buy one unit of some product. Consumer i's valuation, or

[2] See also Anderson and Renault (1999) and Larson (2013), among others.

match value, v_{ij} for firm j's product is drawn randomly from the distribution $F(v)$, i.i.d. across firms. We assume throughout the paper that F has a log-concave density $f(v)$. For a set of m of these firms, consumer i observes the realization of the match value *ex ante*. We call these i's known firms and (without loss of generality) label them 1, 2, ..., m, in decreasing order of match value: $v_{i1} \geq v_{i2} \geq ... \geq v_{im}$. Let $G(v)$ be the distribution of v_{i1}, the best match among i's known firms. If the known firms are a random sample from the market, then we have $G(v) = F(v)^m$. However, we also wish to allow for the possibility that the known firms are not a random sample – for example, if i learns about these firms from friends with similar taste to her own, then their distribution may be distorted away from $F(v)$ by positive selection.

The remaining firms are initially unknown to consumer i, but by paying search cost s she can visit any unknown firm j and learn her match value v_{ij} and firm j's price p_j. Alternatively, by paying search cost $s_p \leq s$, she can visit any known firm and learn its price. At any stage, consumer i may continue searching by picking some previously unvisited firm (known or unknown) to visit. Alternatively, she may choose to purchase at any previously visited firm j, leaving the market with utility $v_{ij} - p_j$. Each consumer chooses her search strategy so as to maximize her expected utility, given her realized known firm match values and her expectation about the distribution of prices. The assumption that $s_p \leq s$ reflects the idea that visiting a known firm may involve less time and effort for the consumer than visiting an unknown firm because she must spend less time investigating the product. A consumer's outside option utility if she makes no purchase is u_0. If candidate equilibrium prices are sufficiently high that consumers would rather take their outside option than engage in search, then no search equilibrium exists, and the market fails. As there is nothing of interest to say in this case, we will generally assume that u_0 is sufficiently small – negative if necessary – that an equilibrium with search exists.

Each firm produces its good on demand at zero marginal cost.[3] From a firm's point of view, the flow of visiting consumers is exogenous. Of these visitors, in expectation λ_d will be *directed consumers* who already know their match values for the firm's product. The remaining number of expected visitors, λ_u, are *undirected consumers* who do not learn their match values until they arrive. The relative numbers of directed and undirected visitors will depend on consumers' search behavior, and a firm cannot discern which consumers are directed and which are not.[4] As we shall see, search behavior depends principally on consumers' expectations about prices (and the search cost parameters), and all equilibria will have the same price at all firms. Anticipating this, a firm's total expected sales

[3] This is just for simplicity of notation; one could think of the marginal cost as being folded into the match value so that v_{ij} represents the net surplus on a sale from j to i.

[4] If a consumer knows about firm j via a connection with one of j's past clients, presumably she could verifiably disclose this fact, proving her directed status. But as we will see, it won't be in her interest to do so.

when it charges p may be summarized by the demand function $D(p\ ;\ p^e)$, where p^e is the price level consumers expect. A firm chooses its price to maximize the expected profit $\pi(p) = pD(p\ ;\ p^e)$. An equilibrium is a profile of prices for firms and search strategies for consumers such that a consumer's search strategy maximizes expected utility given (correct) expectations about prices, and each firm's price maximizes its expected profit given consumers' search strategies.

2.1 Equilibrium with no help from friends

If consumers do not have access to any known firms, the model reduces to a version of the differentiated-products search model covered in Wolinsky (1986), Anderson and Renault (1999), Larson (2013), and others. We will sketch the analysis of this case in broad strokes; for more detail, refer to those earlier papers. Anticipating an equilibrium with a common price level across firms, let p^e be consumers' expectation of that price level. Suppose u' is the best utility offer (net of price) that a consumer has seen so far. If the consumer were to visit one additional firm j and then quit searching and take the best offer in hand, her expected utility would be

$$u' + \Pr\left(v_{ij} - p^e > u'\right) E\left(v_{ij} - p^e - u' \mid v_{ij} - p^e > u'\right) - s,$$

where the middle term is the expected improvement over her current best offer, and the search cost s is the cost of that expected improvement. If the net improvement

$$\Pr\left(v_{ij} - p^e > u'\right) E\left(v_{ij} - p^e - u' \mid v_{ij} - p^e > u'\right) - s$$

is positive, an additional search is worthwhile. Since this expected net improvement is decreasing in the current best offer u', there will be some threshold $\bar{u}$ for the current best offer such that the consumer is indifferent to one additional search:

$$\Pr(v_{ij} - p^e > \bar{u}) E(v_{ij} - p^e - \bar{u} \mid v_{ij} - p^e > \bar{u}) = s$$

By following this line of logic, one can show formally that a consumer's optimal search strategy is to use a threshold rule: purchase at the first firm that offers utility exceeding $\bar{u}$. Again, under the assumption that firms charge a common price, the utility threshold $\bar{u}$ corresponds to a threshold level of the match value: $\bar{v} = \bar{u} + p^e$. Then the condition for this threshold may be rewritten as $\Pr(v_{ij} > \bar{v}) E(v_{ij} - \bar{v} \mid v_{ij} > \bar{v}) = s$, or equivalently, in either of the following three forms:[5]

$$\begin{aligned}(1-F(\bar{v}))\,E(v-\bar{v}\mid v>\bar{v}) &= s\text{, or} \\ \int_{\bar{v}}^{\infty} 1-F(v)\;dv &= s\text{, or} \\ (1-F(\bar{v}))\,E(r(v)\;\mid v>\bar{v}) &= s\end{aligned} \tag{1}$$

where $r(v)=\frac{1-F(v)}{f(v)}$ is the inverse hazard rate of the match value distribution. These relationships emphasize that (given a common price level), the intensity of search will depend on the search cost and the right tail of the match value distribution, but not on the level of prices. The expectation in the expressions above is what is known in the risk literature as the mean excess loss function. In our context, a more descriptive name would be the *excess consumer surplus function*, $XCS(v)=E(v-\bar{v}\mid v>\bar{v})=E(r(v)\;\mid v>\bar{v})$, as it identifies the amount by which an average purchaser's surplus exceeds that of a marginal purchaser.

After each visit, a consumer's search ends with probability $1-F(\bar{v})$, and continues on with probability $F(\bar{v})$. Consequently, she incurs total expected search costs of $\frac{s}{1-F(\bar{v})}$, while her expected match value when she ultimately buys is $E(v\;|v>\bar{v})$. Thus her *ex ante* expected utility from search is $U=E(v\;|v>\bar{v})-p^e-\frac{s}{1-F(\bar{v})}$. Noting that $E(v\;|v>\bar{v})=\bar{v}+E(v-\bar{v}\mid v>\bar{v})$ and using the definition of the threshold $\bar{v}$, this reduces simply to $U=\bar{v}-p^e$. If $\bar{v}-p^e<0$, consumers will refuse to search at all; we will assume this case does not occur, as there is little of interest to say about it.

Next, turn to a firm's pricing problem. If consumers are using cutoff utility $\bar{u}=\bar{v}-p^e$, given the expectation that all firms set price p^e, then a firm that charges price p will sell to a visitor whenever the match value v_{ij} satisfies $v_{ij}-p\geq u$, or equivalently, $v_{ij}\geq\bar{v}-p^e+p$. This yields a profit per visitor of

$$\pi_j(p)=p\left(1-F\left(\bar{v}-p^e+p\right)\right)$$

Log-concavity of $F(v)$ (which follows from log-concavity of $f(v)$) ensures that this profit function is quasiconcave, in which case the optimal price is given by the first order condition $1-F(\bar{v}-p^e+p)=pf(\bar{v}-p^e+p)$. In equilibrium, consumers must hold correct expectations about the price level; substituting $p^e=p$, we have the equilibrium condition

$$p=\frac{1-F(\bar{v})}{f(\bar{v})}=r(\bar{v}) \tag{2}$$

Together, equations (1) and (2) characterize a search equilibrium in the model with no help

[5] For the second, integrate by parts.

from friends.

Some of the comparative statics of the search equilibrium are sensible and straightforward. From (1) it is clear that $\bar{v}$ is decreasing in s, so lower search costs induce consumers to search longer and hold out for better matches. If the match value distribution has a decreasing inverse hazard rate (as is the case for commonly used thin-tailed distributions), then the equilibrium price is decreasing in $\bar{v}$, and so increasing in s: as search costs decline and consumers search longer, prices become more competitive. These conclusions carry into equilibrium consumer utility $U = \bar{v} - r(\bar{v})$: if $F(v)$ has a decreasing inverse hazard rate, then a decline in search costs unambiguously improves consumer welfare (since match quality improves and the price declines).

3 Help from Friends

Now we turn to the main model: consumer i observes m match values *ex ante*, $v_{i1} \geq v_{i2} \geq ... \geq v_{im}$, with the best of these matches distributed according to $G(v)$. First consider consumer i's decision problem when she anticipates facing price p^e at every firm. She is free to search through unknown firms just as above, but at any time she may visit one of the known firms at search cost s_p to get a firm price offer. The first point to notice is that, on the equilibrium path, a consumer never visits any known firm other than the best one, since the utility $v_{i1} - p^e$ dominates $v_{ij} - p^e$ for $j \in \{2, ..., m\}$. In any conceivable situation in which she buys at known firm $j > 1$, she would have been better off visiting firm 1 and buying there instead. Next, observe that the consumer will never visit firm 1 if she does not plan to stop searching and buy there. She knows (barring a deviation by firm 1 from p^e) exactly what surplus to expect at firm 1; if the value of continued search exceeds this surplus, she should not visit firm 1 in the first place.

Given these observations, consumer i's initial situation is as if she were situated in the undirected search model with a current best offer in hand of $\hat{u} = v_{i1} - p^e - s_p$. The only slight difference from the undirected model is the presence of s_p. In the undirected model, a current best offer involves a firm that has already been visited, so the search cost is sunk – taking such an offer involves no additional search cost. In contrast, here, while the consumer correctly anticipates the surplus $v_{i1} - p^e$ at firm 1, she cannot claim it without visiting to pin down a price offer. At this point, we can invoke the standard results from the undirected model to characterize a consumer's optimal search strategy. Let $\bar{u}$ be the threshold utility for undirected search, as above. If the consumer's anticipated utility $\hat{u}$ from visiting her best known firm and purchasing there exceeds $\bar{u}$, she should go there first intending to purchase. Otherwise, she should ignore her known firms and proceed with undirected search.

On the equilibrium path, this is exactly what will happen. However, we must also consider how the consumer should proceed if firm 1 deviates from the price she expected, so that the utility from buying at firm 1 is not what she anticipated when she decided to visit. We claim the following result.

Remark 1 *Suppose a consumer visits her favorite known firm (firm 1) anticipating price p^e and is quoted price p_1. A sufficient condition for her to stop searching and purchase from firm 1 is $p_1 \leq p^e + s_p$.*

This result follows from the logic of the Diamond paradox. Informally, the consumer would not have visited firm 1 unless she expected to do better there, even after paying the search cost s_p, than at her alternative options. But this means that once she arrives at firm 1 and that search cost is sunk, firm 1 could hold her up for a price hike of at least s_p without losing her business. To be a bit more formal about this, having arrived at firm 1, the consumer's utility from accepting its offer is $v_{i1} - p_1$. The consumer's best alternative to purchasing immediately is either to leave to take her second-favorite known offer at firm 2, or to leave and engage in undirected search, whichever is better. Her anticipated utility from leaving for firm 2 is $v_{i2} - p^e - s_p$, so staying at firm 1 is better as long as $p_1 - p^e \leq s_p + (v_{i1} - v_{i2})$. Since she prefers firm 1's product to firm 2's, $p_1 - p^e \leq s_p$ is sufficient to rule out leaving for firm 2.

Alternatively, her expected utility from leaving for undirected search is $\bar{u}$. She would not have visited firm 1 first unless $v_{i1} - p^e - s_p \geq \bar{u}$; that is, unless her expected surplus at firm 1 (including the search cost s_p) were better than undirected search. Given this, once s_p is sunk, $p_1 \leq p^e + s_p$ is a sufficient condition for her utility $v_{i1} - p_1$ from purchasing at firm 1 to exceed $\bar{u}$. The result follows from these two points.

If all of a firm's visitors were directed consumers, this logic would preclude an equilibrium with search: no matter what price consumers anticipated, firms would always have an incentive to hold up visitors by deviating to a slightly higher price. It will be the fact that some of a firm's visitors are engaged in undirected search – and so cannot be held up quite so easily – that puts the brakes on this problem and allows an equilibrium with search to exist.

Before taking up a firm's pricing problem, we summarize a few results about consumers' optimal search strategy. Using $\bar{u} = \bar{v} - p^e$, the condition under which a consumer purchases at her best known firm may be rewritten as $v_{i1} \geq \bar{v} + s_p$. Thus a fraction $1 - G(\bar{v} + s_p)$ of consumers will visit (and purchase at, on the equilibrium path) their best known firms; the remaining fraction $G(\bar{v} + s_p)$ will engage in undirected search among their unknown firms. As noted earlier, a consumer engaged in undirected search continues searching after each visit with probability $F(\bar{v})$, so the total expected number of firms she visits is $\frac{1}{1-F(\bar{v})}$. Together, these imply that the ratio of directed to undirected visitors at any firm is given by $\lambda_d/\lambda_u = \frac{1-G(\bar{v}+s_p)}{G(\bar{v}+s_p)}(1 - F(\bar{v}))$.

Now consider the profit maximization problem of a firm that anticipates that a fraction λ_d of its visitors have come because it is their best known option, while λ_u have arrived due to undirected search. For the time being, we will focus on necessary conditions for a symmetric equilibrium in which all firms charge the same price, deferring the (more technical) necessary conditions for such an equilibrium to exist. If consumers anticipate price p^e everywhere, then a firm's directed visitors will purchase inelastically at any price $p \leq p^e + s_p$, while its undirected buy with probability $1 - F(\bar{v} - p^e + p)$. Consequently, for prices in the interval $p \in [0, p^e + s_p]$, the firm's profit function is

$$\begin{aligned} \pi(p) &= p(\lambda_d + \lambda_u(1 - F(\bar{v} - p^e + p))) \\ &= p(1 - \lambda_u F(\bar{v} - p^e + p)) \end{aligned}$$

This is not a complete characterization of profits: if the firm charges a price $p > p^e + s_p$, it will begin to lose some of its directed visitors, so different expressions for demand and profits will apply. We do not need those expressions yet, as for now we will focus on the necessary conditions for a symmetric equilibrium in which each firm chooses $p = p^e$. When a firm prices in the interval $p \in [0, p^e + s_p]$, its marginal profit is

$$\pi'(p) = \lambda_d + \lambda_u(1 - F(\bar{v} - p^e + p) - pf(\bar{v} - p^e + p)) \qquad \text{for } p < p^e + s_p$$

The undirected visitor term exhibits the standard trade-off from a price increase: higher profit on inframarginal sales offset by the loss of marginal consumers. However, the directed visitor term shows no downside from a marginal price hike, so the presence of directed consumers will tend to encourage firms to price higher. In a symmetric equilibrium, setting $p = p^e$ must be optimal for the firm, so a necessary condition for equilibrium is the first order condition $\pi'(p^e) = 0$. This is equivalent to

$$p = \frac{\lambda_d/\lambda_u + 1 - F(\bar{v})}{f(\bar{v})},$$

or substituting for the ratio of directed to undirected visitors,

$$p = \frac{1}{G(\bar{v} + s_p)} \frac{1 - F(\bar{v})}{f(\bar{v})} \tag{3}$$

If we let p_u be the equilibrium price level in the pure undirected search model of the previous section, then

$$p = \frac{p_u}{G(\bar{v} + s_p)}$$

so the presence of directed consumers unambiguously pushes the equilibrium price upward.

Furthermore, a reduction in the cost of a price quote at a known firm, s_p, also pushes up prices, since firms anticipate facing relatively more of the price-inelastic directed visitors. Conversely, a reduction in the search cost s at unknown firms improves the utility $\bar{v}$ from undirected search and pushes prices down through two channels: p_u falls, just as earlier, and the fall in $G(\bar{v}+s_p)$ as consumers shift away from directed search.

4 Consumer Welfare: Results

We assume for the moment that a symmetric equilibrium exists with price level characterized by (3). This section considers the conditions under which learning about a few products from friends makes consumers better or worse off than they would have been without this information. Given the optimal search strategy, a consumer's *ex ante* expected utility is

$$U = E\left(\max\left(v_1 - s_p, \bar{v}\right) - p\right)$$

This may be written $U = \bar{v} - p + \Pr\left(v_1 > \bar{v} + s_p\right) E\left(v_1 - s_p - \bar{v} \mid v_1 > \bar{v} + s_p\right)$, where the third term represents the additional consumer surplus generated in the event that a consumer's favorite known firm is more attractive than undirected search. Clearly consumers are better off with this extra surplus than they would have been if directed search were not an option; the question is whether this extra surplus will outweigh the price hike that directed search permits firms to impose. After integrating by parts, consumer utility may be written

$$\begin{aligned} U &= \bar{v} + \int_{\bar{v}+s_p} 1 - G(v)\ dv - p \\ & \bar{v} + \int_{\bar{v}+s_p} 1 - G(v)\ dv - \frac{1}{G(\bar{v}+s_p)} \frac{1-F(\bar{v})}{f(\bar{v})} \end{aligned}$$

Since consumer utility would be $U_0 = \bar{v} - p_u$ in the absence of directed search, the net utility improvement when directed search is possible is $U - U_0 = B(\bar{v}+s_p) - C(\bar{v}, s_p)$, where $B(\bar{v}+s_p) = \int_{\bar{v}+s_p} 1 - G(v)\ dv$ is additional surplus from directed search and $C(\bar{v}, s_p) = p - p_u = \frac{1-G(\bar{v}+s_p)}{G(\bar{v}+s_p)} \frac{1-F(\bar{v})}{f(\bar{v})}$ is the equilibrium price rise. If we define $XCS_G(v)$ to be the excess consumer surplus function under distribution $G(v)$, then this net utility improvement may also be expressed as $U - U_0 = (1 - G(\bar{v}+s_p))\left(XCS_G(\bar{v}+s_p) - \frac{r(\bar{v})}{G(\bar{v}+s_p)}\right)$.

A useful benchmark is the case in which each consumer knows her match value at one randomly chosen firm.

Proposition 1 *Suppose each consumer knows her match value at one randomly chosen firm; that is, $G(v) = F(v)$. In equilibrium, consumers are worse off than they would be without this information.*

Proof. In this case, $B(\bar{v}+s_p) = \int_{\bar{v}+s_p} 1 - F(v)\ dv$ and $C(\bar{v}, s_p) = \frac{1-F(\bar{v}+s_p)}{F(\bar{v}+s_p)} r(\bar{v})$. The benefit term may be written as $B(\bar{v}+s_p) = (1 - F(\bar{v}+s_p)) E(h(v) \mid v > \bar{v}+s_p)$, so

$$\begin{aligned} U - U_0 &= (1 - F(\bar{v}+s_p)) \left(E(r(v_1) \mid v_1 > \bar{v}+s_p) - \frac{r(\bar{v})}{F(\bar{v}+s_p)} \right) \\ &< (1 - F(\bar{v}+s_p)) (E(r(v_1) \mid v_1 > \bar{v}+s_p) - r(\bar{v})) \end{aligned}$$

But log-concavity of $f(v)$ implies that the inverse hazard rate $r(v) = \frac{1-F(v)}{f(v)}$ is weakly decreasing, so the second term in parenthesis is negative. Thus $U - U_0 < 0$. ■

So the benefit of having one match value "in hand" before embarking on search is swamped by the price rise firms extract in anticipation of facing visiting consumers who are more easily held up. Note this result does not depend on the cost of getting a price quote; in particular, it still holds if visiting the known firm is free ($s_p = 0$). While reducing s_p makes visiting the known firm more attractive to consumers, it is precisely the fact that this tilts the balance of demand toward directed consumers that moves prices against them.

A consumer knows her value at n randomly chosen firms

If information about match values is thought of as coming from one's friends' choices, observing one randomly chosen match may sound pessimistic. A consumer with many friends may have the luxury of observing many friends' choices and opting for the firm that is her best match from this larger set. Alternatively, one might think that friends have similar taste, so a sample of friends' choices improves on $F(v)$ by screening out firms that would be bad matches for a consumer. We discuss these possibilities in turn.

To study the first case, suppose that each consumer sees a random sample of n firms with match values drawn from $F(v)$, so the match value at the favorite known firm has distribution $G(v) = F(v)^n$. We write $p_n = \frac{1}{F(\bar{v}+s_p)^n} r(\bar{v})$ for the equilibrium price and U_n for consumer utility.

Proposition 2 *In the model where consumers know n randomly chosen match values, consumers are worse off with larger samples: U_n is decreasing in n. Consequently, for any n, consumers are worse off than they would be under pure undirected search.*

Proof. Write $B_n(\bar{v}+s_p) = \int_{\bar{v}+s_p} 1 - F(v)^n\ dv$, defining $B_0(\bar{v}+s_p) = 0$, and $p_0 = p_u$. Note that $U_n = \bar{v} + B_n(\bar{v}+s_p) - p_n$ and so $U_{n+1} - U_n = \Delta B_n - \Delta p_n$, where $\Delta B_n = B_{n+1}(\bar{v}+s_p) - B_n(\bar{v}+s_p)$ and $\Delta p_n = p_{n+1} - p_n$. The previous proposition established that $\Delta B_0 < \Delta p_0$, so it suffices to show that ΔB_n and Δp_n are decreasing and increasing in n, respectively. For the former, we have $\Delta B_n = \int_{\bar{v}+s_p} F(v)^n (1 - F(v))\ dv$. Since $F(v) \leq 1$, the integrand is decreasing in n. For the latter, we have $\Delta p_n = \frac{1-F(\bar{v}+s_p)}{F(\bar{v}+s_p)^n} r(\bar{v})$ which is increasing in n. ■

The thrust of the result is essentially that the benefits of a larger sample are concave in n, while the equilibrium price has a positive convex relationship with n. Consequently, if the first known firm hurts consumer welfare on net, knowing about additional firms will make consumers progressively worse and worse off.

As an example, consider the case in which the taste distribution $F(v)$ is Type I extreme value with mean μ and variance σ^2, a commonly used assumption in the discrete choice literature. Suppose the valuation distribution $G(v)$ at a consumer's favorite known firm is also Type I extreme value with the same variance and mean μ_G. Then Proposition 2 implies that if $\mu_G \geq \mu$, consumers are worse off with access to this favorite known firm than they would be without it, and furthermore, improvements in μ_G leave them worse off. This follows from the fact that $G(v)$ is equivalent to the distribution of the largest of $n = \exp\left((\mu_G - \mu)\,\pi/\sqrt{6\sigma^2}\right)$ draws from $F(v)$.[6]

A consumer knows her value at one firm (with poor matches screened out)

Next consider the case in which a consumer is more likely to be well-matched with her friends' firms than with a random firm. As motivation, suppose each consumer has one friend whose taste is identical to her own. In the background, we imagine that this friend found his firm through a search process like those we have been discussing here, using some cut-off match value $\hat{v}$. (This cut-off might be equal to $\bar{v}$, but it need not be – this friend could have had a different search cost.) In this case, the friend's match value will be a draw from a truncated version of $F(v)$. With this motivation, we suppose that each consumer observes a single known firm with match value drawn from $G(v) = F_{\hat{v}}(v) = \frac{F(v)-F(\hat{v})}{1-F(\hat{v})}$ on $[\hat{v}, \infty)$. We assume $\hat{v} < \bar{v} + s_p$, so that the friend's firm is sometimes not worth a visit. If this were not the case, there would be no equilibrium with search, as the Diamond paradox would apply in full force: all consumers would prefer to start at their known firms, and anticipating this, firms would face no restraint on raising prices relative to consumers' expectations.

Proposition 3 *If each consumer knows her value at one firm drawn from $G(v) = F_{\hat{v}}(v)$, so that matches worse than $\hat{v}$ are screened out, then consumers are worse off than they would have been under pure undirected search. Furthermore, an increase in $\hat{v}$ (improving the distribution of the known firm) makes consumers worse off.*

Proof. As earlier, the benefit term from directed search may be written

$$B(\bar{v} + s_p) = (1 - G(\bar{v} + s_p))\, E\left(\frac{1 - G(v_1)}{g(v_1)} \mid v_1 > \bar{v} + s_p\right).$$

[6] The proof of the proposition is written for integral n (in keeping with the motivation in the text), but it is not hard to confirm that the mathematical result holds for any positive real n.

But note that truncating a distribution to the left does not change its inverse hazard rate: $\frac{1-G(v)}{g(v)} = \frac{1-F(v)}{f(v)}$ for $v \geq \hat{v}$. Thus the net benefit from directed search is

$$U - U_0 = B(\bar{v} + s_p) - C(\bar{v}, s_p) = (1 - G(\bar{v} + s_p)) \left(E(r(v_1) \mid v_1 > \bar{v} + s_p) - \frac{r(\bar{v})}{G(\bar{v} + s_p)} \right)$$

where $r(v) = \frac{1-F(v)}{f(v)}$ as earlier. Aside from the replacement of $F(\bar{v} + s_p)$ with $G(\bar{v} + s_p)$, the second term in parentheses is identical to its counterpart in the proof of Proposition 1; it is negative by the same argument used there.

For the second claim, it suffices to show that $\frac{d}{d\hat{v}}(U - U_0)$ is negative. Taking this derivative, we have

$$\begin{aligned} \frac{d}{d\hat{v}}(U - U_0) &= f(\hat{v}) \frac{1 - F(\bar{v} + s_p)}{(1 - F(\hat{v}))^2} E(r(v_1) \mid v_1 > \bar{v} + s_p) - f(\hat{v}) \frac{1 - F(\bar{v} + s_p)}{(F(\bar{v} + s_p) - F(\hat{v}))^2} r(\bar{v}) \\ &= f(\hat{v}) \left(\frac{B(\bar{v} + s_p)}{1 - F(\hat{v})} - \frac{C(\bar{v}, s_p)}{F(\bar{v} + s_p) - F(\hat{v})} \right) \end{aligned}$$

We have already established that $B(\bar{v} + s_p) \leq C(\bar{v}, s_p)$. Clearly $1 - F(\hat{v}) \geq F(\bar{v} + s_p) - F(\hat{v})$, so the term in parentheses inflates the cost term by a larger factor than the benefit term, so the claim follows. ■

The logic echoes the case in which the known firm is a random draw: the more that consumers take up the option of going to a known firm, the more the shift in demand composition toward directed consumers lets firms extract higher prices, swamping the benefits from that option. Because a better known option (higher $\hat{v}$) implies a larger shift in the composition of demand, it does not make consumers better off. In particular, as the worst-case threshold for the known firm rises toward $\hat{v} = \bar{v} + s_p$, the fraction of consumers engaged in undirected search (proportional to $F(\bar{v} + s_p) - F(\hat{v})$) vanishes, removing the last trace of discipline on firms' pricing decisions.

Corollary 1 *Suppose each consumer knows her value at n firms drawn from the truncated distribution $F_{\hat{v}}(v)$. Then an increase in the sample size n makes consumers worse off.*

This is a straightforward extension of the arguments in Propositions 2 and 3.

Do consumers benefit if it becomes easier to get a price quote at a known firm?

Next we consider the welfare impact of changes in s_p, the cost of obtaining a firm price quote at a firm where the consumer already knows her valuation. Clearly, a reduction in this cost should induce a shift in demand composition toward directed visitors; given the results so far, it should not be surprising that this shift may induce a price change that leaves consumers worse off. We consider a scenario that nests all of the cases discussed above: suppose each consumer has n known firms with valuations drawn from the truncated distribution $F_{\hat{v}}$ (so that $G(v) = F_{\hat{v}}(v)^n$).

Proposition 4 *Suppose consumers have n known firms with valuations drawn from $F_{\hat{v}}$. Then consumer utility is increasing in the cost s_p of obtaining a price quote from a known firm.*

Proof. In this case, we start from $U = \bar{v} + B(\bar{v} + s_p) - p$. It suffices to show that $\frac{d}{ds_p}(B(\bar{v} + s_p) - p)$ is positive. We have $\frac{dB(\bar{v}+s_p)}{ds_p} = -(1 - G(\bar{v} + s_p))$ and $\frac{dp}{ds_p} = \frac{g(\bar{v}+s_p)}{G(\bar{v}+s_p)^2} r(\bar{v}) = \frac{g(\bar{v}+s_p)}{G(\bar{v}+s_p)} p$, so it will suffice to show $\left.\frac{G(v)(1-G(v))}{g(v)}\right|_{v=\bar{v}+s_p} \leq p$. Direct computations give

$$\begin{aligned} \frac{G(v)(1-G(v))}{g(v)} = \frac{F_{\hat{v}}(v)(1 - F_{\hat{v}}(v)^n)}{n f_{\hat{v}}(v)} &= \frac{1 - F_{\hat{v}}(v)}{f_{\hat{v}}(v)} F_{\hat{v}}(v) \left(\frac{1}{n} \sum_{k=0}^{n-1} F_{\hat{v}}(v) \right) \\ &\leq \frac{1 - F_{\hat{v}}(v)}{f_{\hat{v}}(v)} = \frac{1 - F(v)}{f(v)} \end{aligned}$$

So then because $r(v)$ is a decreasing function, we have $\left.\frac{G(v)(1-G(v))}{g(v)}\right|_{v=\bar{v}+s_p} \leq r(\bar{v} + s_p) \leq p = \frac{1}{G(\bar{v}+s_p)} r(\bar{v})$. ∎

4.1 When does having recommended firms make consumers worse off: general conditions

We owe the simplicity of the analysis so far to the assumption that a consumer's valuations for products she knows about *ex ante* are distributed similarly to her valuations for unknown products. While this assumption is fairly natural, one could imagine circumstances under which the favorite known firm distribution G is only tenuously connected to the raw distribution of values F. (For example, we might wish to allow for the firms a consumer sees to be shaped by arbitrary forms of selection.) In this section, we consider general conditions under which access to known firms helps or harms consumers' equilibrium utility. The latter are easier to come by; for the former we focus on examples that illustrate the features G must have in order for consumers to benefit.

Let $r_G(v) = \frac{1-G(v)}{g(v)}$ be the inverse hazard rate for the favorite known firm. As earlier, the net gain or loss to consumers from access to known firms is $U - U_0 = B(\bar{v} + s_p) - C(\bar{v}, s_p)$, which may be written as

$$\begin{aligned} U - U_0 &= (1 - G(\bar{v} + s_p)) \left(E(r_G(v_1) \mid v_1 > \bar{v} + s_p) - \frac{1}{G(\bar{v} + s_p)} r(\bar{v}) \right) \\ &= (1 - G(\bar{v} + s_p)) \left(XCS_G(\bar{v} + s_p) - \frac{1}{G(\bar{v} + s_p)} p_u \right) \end{aligned}$$

where the excess consumer surplus function is defined analogously for G as it was for F. First we highlight a simple ranking of consumer and producer surplus that applies, in any context, not just search, whenever demand is logconcave.

Lemma 1 *Suppose a firm faces a consumer with valuation $\tilde{v}$ ˜ $F(v)$ for its product and outside option utility $\hat{u}$. Suppose $F(v)$ has a logconcave density, and the firm sets the profit-maximizing price p^*. Let the firm's profit margin on a sale be PS, and let the consumer's expected surplus on a sale in excess of her outside option be XCS. Then $PS \geq XCS$.*

Proof. To emphasize that normalizing marginal cost to zero is inessential, let the firm have marginal cost $c \geq 0$, so that the profit function is $\pi(p) = (p-c)(1-F(\hat{u}+p))$. The firm's first order condition gives $PS = p^* - c = r(\hat{u}+p^*)$, where $r(v) = \frac{1-F(v)}{f(v)}$ as usual. A consumer's surplus from a sale is $E(v \mid v > \hat{u}+p^*) - p^*$, so

$$XCS = XCS(\hat{u}+p^*) = E(v - \hat{u} - p^* \mid v > \hat{u}+p^*) = E(r(v) \mid v > \hat{u}+p^*)$$

Then because $r(v)$ is decreasing, we have $XCS \leq PS$. ■

Intuitively, when consumer tastes are thin-tailed – i.e., when inframarginal consumers do not value a good too much more than marginal ones – a firm can capture the lion's share of surplus from a sale. This is the key fact underlying Proposition 1. The value to a consumer of a free additional search (assuming $s_p = 0$), is $XCS(\bar{v})$ times the chance that search results in a purchase. Because $XCS(\bar{v})$ is less than the current price level p_u, the extra search cannot be worthwhile for consumers unless the chance it is successful exceeds the proportionate increase in the price level that it induces in general equilibrium. But the chance the extra search leads to a purchase (again assuming $s_p = 0$) is $1 - F(\bar{v})$, while the proportionate price rise is the purchase/no purchase ratio $\frac{1-F(\bar{v})}{F(\bar{v})}$, so the extra search cannot be worthwhile.

Proposition 5 *The following are sufficient conditions for consumers to be worse off with favorite known firm v_1 ˜ $G(v)$ than they would have been under pure undirected search.*

(i) $XCS_G(v) \leq XCS(v)$ for all v

(ii) $r_G(v) \leq r(v)$ for all v

(iii) $G(v)$ is a downward shift of $F(v)$; that is, $G(v) = F(v+z)$ for a positive constant z.

(iv) $\frac{G(v)(1-G(v))}{g(v)} \leq \frac{F(v)(1-F(v))}{f(v)}$ for all v and $r_G(v)$ is decreasing.

Proof. (i) We have $XCS_G(\bar{v}+s_p) \leq XCS(\bar{v}+s_p)$ by hypothesis and $XCS(\bar{v}+s_p) \leq r(\bar{v}) = p_u \leq \frac{p_u}{G(\bar{v}+s_p)}$ because $r(v)$ is decreasing.

(ii) The condition $r_G(v) \leq r(v)$ for all v implies $XCS_G(\bar{v}+s_p) \leq XCS(\bar{v}+s_p)$, so the result follows from (i).

(iii) We have $r_G(v) = r(v+z)$; given $r(v)$ decreasing, this implies $r_G(v) \leq r(v)$ for all v, and so the result follows from (ii).

(iv) Since $F(\bar{v}+s_p) \leq 1$, it suffices to show $G(\bar{v}+s_p)\, XCS_G(\bar{v}+s_p) \leq F(\bar{v}+s_p)\, p_u = F(\bar{v}+s_p)\, r(\bar{v})$. Furthermore, given $r(v)$ decreasing, it suffices to show

$$G(\bar{v}+s_p)\, XCS_G(\bar{v}+s_p) \leq F(\bar{v}+s_p)\, r(\bar{v}+s_p) = \frac{F(\bar{v}+s_p)(1-F(\bar{v}+s_p))}{f(\bar{v}+s_p)}.$$

But given $r_G(v)$ decreasing, we have

$$G(\bar{v}+s_p)\, XCS_G(\bar{v}+s_p) \leq G(\bar{v}+s_p)\, r_G(\bar{v}+s_p) = \frac{G(\bar{v}+s_p)(1-G(\bar{v}+s_p))}{g(\bar{v}+s_p)},$$

so the result follows. ■

Parts (i)-(iii) are not entirely satisfying, as they amount to conditions under which a consumer's valuation at his favorite known firm is less auspicious than at a random firm. (Note that condition (ii) implies that $G(v)$ is first-order stochastically dominated by $F(v)$.) Given Proposition 1, the consumer welfare result is not too surprising, and furthermore the case when favorite known firms are superior to a random draw, not inferior, would seem more empirically relevant.

Condition (iv) is a bit more interesting. To flesh out its potential, consider an example in which both $F(v)$ and $G(v)$ are logistic, with mean and variance (μ, σ^2) or (μ_G, σ_G^2) respectively. The logistic distribution has the feature that $\frac{F(v)(1-F(v))}{f(v)} = \sigma$ is constant, so condition (iv) holds whenever $\sigma_G \leq \sigma$. Consequently, if consumers get a lower-variance draw at their favorite known firms than they do from a random firm, then access to a favorite known firm hurts them – regardless of how much better or worse that draw is on average than a random firm.

5 Referral Discounts

We have seen that consumers may not be enthusiastic about acting on referrals to known firms because they expect to be held up when they arrive (just as any arriving consumer, directed or undirected, will be). This means that firms might wish to give consumers assurances in advance that they will not be exploited too much if they show up. We assume that these assurances take the form of a referral discount at known firms: in addition to learning her taste for a known firm's product, a consumer also learns about any discount it is willing to offer her off of its retail price.

In principle, a firm might wish to publicize a discount offer to all consumers, not just those who are familiar with its product. This would involve reaching undirected searchers, perhaps with advertising, at which point they would no longer be considered undirected. While this is a reasonable line of inquiry, we will not pursue it here. Our motivation for focusing on discounts to "connected" consumers is twofold. First, it is natural to imagine that the channel by which a consumer learns about a known product – for example, a referral by a friend – is also the channel through which discount offers are passed on. Second, casual empiricism suggests that "refer a friend" discounts are very common.

We focus on discounts rather than price offers because the latter are not particularly interesting. If a firm is able to commit to a price offer in advance, the hold-up problem is solved (at least vis-à-vis directed searchers), and there is not much further to say. This is not as limiting as it may seem – it is often hard to make a credible commitment to an advance price offer because the firm will alway have an incentive to violate the spirit of that commitment by tacking on ancillary fees and surcharges after the consumer arrives. That said, because rational consumers will anticipate the retail price, a \$20 discount will amount to an *implicit* price offer. The difference is that consumers understand that their take-up of the discount offer will affect the firm's *ex-post* incentives in setting its retail price. Bluntly, if the firm is offering a 90% discount, then they should expect to find a very high retail price, and factor this into their decisions about where to shop. Thus one can think of discounts as incentive-compatible price offers to connected consumers, in the sense that if consumers take the discounted price at face value, their demand response to it affects the firm's retail pricing exactly as much as they were expecting.

We will restrict attention to the special case in which each consumer has exactly one known firm, with value distributed according to $G(v)$. Thus each consumer sees at most one discount offer before searching. If a single consumer were to see multiple discount offers before searching, we would in effect have a hybrid model with a first stage of price competition with differentiated products and a second stage of search, with the added complication that first stage price offers must be incentive-compatible relative to the search demand they induce at the second stage. While this would be worth pursuing, our ambitions at this point will be more modest.

5.1 Setup and Analysis

Let us set up some notation. We will focus on symmetric equilibria; as earlier, let p^e refer to the expected (retail) price that consumers expect to face at any firm. This is the price that any undirected consumer will pay; a directed consumer may pay less if she is offered a discount. As earlier, let $\bar{v}$ be the threshold match value for an undirected searcher; $\bar{v}$ is pinned down by the search cost s and optimal search behavior according to (1), just as before. In addition to observing the match value at her known firm, a consumer sees a discount offer $d \geq 0$. Let $\hat{p}(d)$ be the retail price she anticipates at her known firm, given the discount offer, with $p_d(d) = \hat{p}(d) - d$ the discounted price she expects to pay. We will often drop arguments and write $\hat{p}$ and p_d; the dependency between these prices and the discount will be fleshed out momentarily.

The logic governing a consumer's optimal search strategy is just as earlier, with minor changes to reflect the discount. As before, there is no reason to visit her known firm unless she plans to buy immediately (assuming the price is as she expected). The utility of embarking on undirected search remains $\bar{v} - p^e$, while the utility from going to her known firm with plans to buy immediately is $v_1 - p_d - s_p$. Thus she will visit her known firm if $v_1 \geq \hat{v} = \bar{v} - p^e + p_d + s_p = \bar{v} - p^e + \hat{p} - d + s_p$ and search otherwise. And a revised version of the hold-up problem described in Remark 1 continues to apply. If upon arriving at her known firm she faces a price $p_1 > \hat{p}$ higher than the retail price she expected, as long as the price hike is small ($p_1 < \hat{p} + s_p$) she will still follow through with the purchase rather than walking away (since the search cost s_p is now sunk). Thus, conditional on arriving at her known firm, her demand is inelastic to small price hikes.

Turning to firms' pricing decisions and equilibrium conditions, the same caveats apply as before. We will focus on necessary first-order conditions for profit maximization and a symmetric equilibrium. In the neighborhood of a symmetric profile of retail price and discount strategies, small price deviations by a firm will fall into the region where directed consumers may be treated as inelastic. Our first-order conditions will rule out these local deviations. At this point we do not have simple conditions that ensure that firms' profit functions are single-peaked (which would guarantee that locally optimal pricing is also globally optimal for a firm).[7] As an advance warning, we will often pursue clean results by taking $G(v) = F(v)$ (known firms follow the same match value distribution as any other firm) and $s_p \to 0$. The latter, in particular, may seem a contradiction; after all, the logic of the hold-up problem that directed consumers face is that s_p is strictly positive (giving a firm room to hike its price slightly once that cost is sunk). However, the logic of the hold-up problem applies for arbitrarily small (but positive) s_p, and because this permits all firms' prices to inch up in tandem, a smaller s_p puts no inherent limit on how high the general price level may rise. Thus our limiting results for $s_p = 0$ should be taken as a clean approximation of market outcomes when s_p is very small.

In a candidate symmetric equilibrium, let d^e be the common discount level. This implies a common threshold $\hat{v}^e = \bar{v} - d^e + s_p$ above which consumers shop at their known firms and below which they search. Following our earlier analysis, this means that each firm receives a volume $D_u = \frac{G(\hat{v}^e)}{1-F(\bar{v})}$ of undirected arrivals. If it charges retail price $\hat{p}$, the fraction of

[7] There is no particular reason to worry that profits will not be quasiconcave, given reasonable distributions for $F(v)$ and $G(v)$. The complication is that demand reflects a mixture of $F(v)$ and $G(v)$ (with the latter truncated to inelastic below a threshold), so a clean condition guaranteeing (say) log-concavity of demand is hard to obtain.

them who will purchase is $1 - F(\bar{v} - p^e + \hat{p})$.

The firm can control how many of its potential directed consumers arrive with the discount d it offers. However, because these arrivals will be *ex post* inelastic, the more of them that it lures into coming, the higher it will be tempted to set its retail price $\hat{p}$, and these consumers anticipate this consequence when evaluating the discount *ex ante*. With this in mind, we start by evaluating the firm's retail price-setting incentives in the event that it has offered a discount d and managed to lure a volume D_d of directed consumers who expect to face retail price $\hat{p}^e$ (which will generally be different from p^e, as it depends on the particular discount that was offered). This will tell us the actual retail price $\hat{p}$ that the firm offers. Then we impose the conditions that $\hat{p}^e = \hat{p}$ (consumers correctly anticipate the *ex post* retail price implied by the discount and their demand response), and that the volume D_d of arrivals is consistent with optimal consumer behavior given correct expectations about the retail price.

Given d, D_u, and $\hat{p}^e$, the firm's total demand at retail prices in a neighborhood around $\hat{p}^e$ is given by

$$D(\hat{p}) = D_d + D_u (1 - F(\bar{v} - p^e + \hat{p})) \tag{4}$$

and so its profit at prices in this neighborhood is given by

$$\pi(\hat{p}) = (\hat{p} - d) D_d + D_u \pi_u(\hat{p}) \tag{5}$$

where

$$\pi_u(\hat{p}) = \hat{p}(1 - F(\bar{v} - p^e + \hat{p})) \tag{6}$$

Notice that discount payments dD_d are effectively sunk at this point. The first-order condition for *ex post* profit maximization is given by

$$\pi'(\hat{p}) = D_d + D_u \pi'_u(\hat{p}) = 0 \quad \text{or} \tag{7}$$

$$D_d = -D_u \pi'_u(\hat{p}) \tag{8}$$

Equation (8) has a natural interpretation. If there were no undirected searchers to worry about, the firm would face no discipline at all on its temptation to hold up the directed ones. But when D_u is positive, hiking up the retail price $\hat{p}$ to exploit directed consumers imposes a cost on profits from undirected searchers: the firm is over-pricing relative to their demand, and overall profits suffer as they begin to take their business elsewhere in greater numbers. It is precisely the presence of these undirected searchers that gives credibility to the discount the firm offers to its directed consumers – they know that the firm cannot *completely* unwind that discount with a higher retail price for fear of losing too much undirected business. By the same token, the larger D_d is, the more the firm will be willing to distort its undirected sales with a higher price – the marginal directed consumer will understand that by virtue of showing up, she increases D_d and has this side effect on the retail price.

Next, consistency of directed consumers' decisions about whether to visit requires, first, that consumers expect a retail price consistent with (8), and second, that D_d is given by the fraction of directed consumers who are willing to visit when they anticipate price $\hat{p}$. Given d and the anticipated $\hat{p}$, only consumers with match values above the threshold $\hat{v}$ given above will visit. Thus we have

$$D_d = 1 - G\left(\hat{v}\right) \tag{9}$$

where $\hat{v}$ satisfies

$$\hat{v} = \bar{v} - p^e + p_d + s_p \ , \tag{10}$$

the implicit discount price p_d satisfies

$$p_d = \hat{p} - d \tag{11}$$

and the retail price $\hat{p}$ satisfies (8).

A firm's initial decision involves what discount d, if any, to offer. But given these relationships, we might just as well treat its strategic variable as directed demand D_d, the marginal directed consumer $\hat{v}$, or the implicit discount price p_d – whichever is convenient – with the other variables and $\hat{p}$ pinned down by the consistency conditions and *ex post* profit maximization. This is true mathematically, but to help motivate this flexibility of the strategic variable more intuitively, consider a sequence of cheap talk dialogue between the firm and potential directed consumers. First the firm promises a discounted price p_d. Consumers respond that, if we were to believe that offer, we would show up in numbers given by (10) and (9), and then you could not help yourself from setting a retail price according to (8). Consequently, to convince us that your promised discount price is really credible, you will have to commit to a discount d given by (11).

With this motivation, let us write the firm's *ex ante* profit as

$$\Pi\left(p_d\right) = \pi_d\left(p_d\right) + D_u \pi_u\left(\hat{p}\left(p_d\right)\right) \tag{12}$$

where

$$\pi_d\left(p_d\right) = p_d D_d = p_d\left(1 - G\left(\hat{v}\left(p_d\right)\right)\right)$$

with $\hat{v}\left(p_d\right)$ given by (10) and $\hat{p}\left(p_d\right)$ determined by (8), writing both $\hat{v}$ and $\hat{p}$ so as to emphasize the dependence on the promised discount price p_d. It should be pointed out that discounts are required to be weakly positive: $d \geq 0$, and so the discount price is constrained by $p_d \leq \hat{p}$. This follows from the fact that the firm cannot distinguish directed and undirected searchers unless the former present their discount offers. If the firm were to try to charge them a premium rather than a discount, they would simply choose to blend in with the crowd. If marginal profit $\Pi'(p_d)$ is strictly positive at $p_d = \hat{p}$, then (under suitable regularity conditions), the firm does best to offer no discount at all.[8] Conversely, if $\Pi'(p_d) < 0$ where $p_d = \hat{p}$, then the firm will discount, and its optimal discount price must satisfy the first-order condition

$$\Pi'(p_d) = \pi_d'(p_d) + D_u \pi_u'(\hat{p}(p_d)) \frac{d\hat{p}}{dp_d} = 0 \tag{13}$$

In order to make headway in understanding (13), it will be helpful to temporarily impose a stronger condition than quasiconcavity on profits.

Condition 1 *The profit function $\pi_u(\hat{p})$ over undirected consumers is strictly concave with $\pi_u'(\hat{p}) \to -\infty$ as $\hat{p} \to \infty$.*

These strong assumptions guarantee that (8) will have an interior solution with the unambiguous comparative statics properties described in the remark below; we will discuss what happens when Condition 1 is not met later.

Remark 2 *An increase in the promised discount price p_d will induce an increase in $\hat{v}$ and a reduction in D_d: fewer directed consumers will patronize the firm. Under Condition 1, this induces a reduction in the retail price $\hat{p}$ and a reduction in the discount d that the firm must offer.*

So under Condition 1, the discount and retail prices move in opposite directions. The first claim follows directly from (10) and (9). Then Condition 1 guarantees that (8) has a solution with the righthand side positive and increasing in $\hat{p}$, so a reduction in D_d must imply a reduction in $\hat{p}$.

As a benchmark it is useful to define p_u^m and p_d^m, the "monopoly" prices that maximize $\pi_u(\hat{p})$ and $\pi_d(p_d)$ respectively (recognizing that these prices depend on market-wide price levels that are suppressed here). The former, p_u^m, corresponds to the *ex post* retail price that the firm would set if it faced no directed consumers. The latter represents the "full commitment" price the firm would like to offer to directed consumers *ex ante*, if it could make a binding offer. We already know from equation (8) that if the firm offers a positive discount it will set $\hat{p} > p_u^m$: that is, it will set a higher than desirable retail price for undirected consumers. Furthermore, because $\pi_u'(\hat{p})$ and $\frac{d\hat{p}}{dp_d}$ are both negative, the *ex ante* first-order equation (13) implies that $\pi_d'(p_d) < 0$, so the firm will promise a discount price $p_d > p_d^m$. In other words, the firm will discount less to directed consumers than it would if

[8] Specifically, quasiconcavity of $\Pi(p_d)$ will suffice.

it could make them a binding price offer. This is because in the absence of commitment, discounting carries a cost – the consequent overpricing to undirected retail consumers.

Although the following conclusion is perhaps a bit tautological in light of these factors, it provides a useful context for thinking about discounting.

Remark 3 *A necessary condition for a firm to use discounting is that $p_d^m < p_u^m$. That is, the monopoly price it would like to offer directed consumers ex ante, before searching, is below the monopoly price it would like to charge undirected consumers ex post, upon arrival.*

Broadly speaking, two factors affect whether the condition $p_d^m < p_u^m$ is met. Because of the hold-up problem in a search equilibrium, firms tend to charge *all* consumers a higher price than they would like to commit to *ex ante*. Indeed, they would like to commit to discounted prices for undirected searchers too, if they were able to reach them. This factor tends to favor discounting. However, the valuations of directed consumers are not necessarily distributed identically to those of undirected consumers; this depends on the distributions $G(v)$ and $F(v)$. If a consumer's known firm is cherry-picked from the overall distribution – perhaps because she learns about its product from friends who share her taste – then the monopoly price for directed consumers could be higher than p_u^m. We will defer any formal analysis of this comparison for the time being.

Next we flesh out the relationship between the retail and discount prices. To identify $d\hat{p}/dp_d$, we differentiate condition (8) implicitly to get

$$\frac{d\hat{p}}{dp_d} = -\frac{1}{D_u \pi_u''(\hat{p})} \frac{dD_d}{dp_d}$$

or using (10) and (9),

$$\frac{d\hat{p}}{dp_d} = \frac{g(\hat{v})}{D_u \pi_u''(\hat{p})}$$

or cast as an elasticity,

$$\varepsilon_{\hat{p},p_d} = \frac{d\hat{p}}{dp_d} \frac{p_d}{\hat{p}} = \frac{1}{\hat{p}\pi_u''(\hat{p})} \frac{D_d}{D_u} \varepsilon_d$$

where ε_d is the *ex ante* demand elasticity of directed consumers. Speaking informally, we may say that the firm has good commitment power when $|d\hat{p}/dp_d|$ is small; in this case it is able to credibly promise a discounted price without distorting its retail price too much. Commitment power is strong when directed demand is a small fraction of total demand (D_d/D_u small) – or precisely when discounting wouldn't be very helpful. Commitment power is also strong when undirected profits are strongly concave – this indicates that the firm will not be tempted to distort retail prices very much *ex post* because the harm done to undirected profits at the margin grows quickly. Finally, commitment power is strong when directed demand is inelastic – this reflects the fact that discounting p_d will elicit a small demand response, and it is the tilt of total demand composition toward directed consumers that induces distortion of the retail price. Of course, this is not particularly helpful either, as it is precisely when directed demand is inelastic that the firm will have little to gain from discounting.

Using these results and (8), we may write the *ex ante* first-order condition as

$$\Pi'(p_d) = \pi'_d(p_d) - \frac{D_d}{D_u}\frac{g(\hat{v})}{\pi''_u(\hat{p})} = 0 \tag{14}$$

5.2 Symmetric Equilibrium When Discounts Are Possible

Equation (14) and the accompanying analysis establish the incentives facing an individual firm, given the pricing of other firms (p^e and d^e) and the consumer search behavior it induces (threshold $\hat{v}^e$ and undirected consumer arrivals D_u). Now we close the model and look for a symmetric equilibrium. As there is quite a bit to keep track of, we begin fairly simply by examining whether there may be an equilibrium *without* discounting.

5.2.1 Will Discounts Be Used?

Consider a candidate strategy profile in which all firms charge p^e and none of them discounts. A deviation to a positive discount corresponds to a deviation to a promised discount price $p_d < p^e$. With this in mind, we consider the sign of $\Pi'(p_d)$, evaluated at $p_d = \hat{p} = p^e$, under a symmetric strategy profile.

The threshold for directed consumers reduces to $\hat{v} = \bar{v} + s_p$, undirected arrivals are $D_u = G(\hat{v}) / (1 - F(\bar{v}))$, and it is easily confirmed with (8) that the optimality of pricing at $\hat{p} = p^e$ must imply that the retail price is given, as earlier, by

$$p^e = \frac{1}{G(\hat{v})}\frac{1 - F(\bar{v})}{f(\bar{v})} \tag{15}$$

The first point to note is that firms' desire to discount below p^e does not necessarily mean that discounting will occur. That is, as discussed earlier, discounting will be attractive in principle, if firms had commitment power, when $p_d^m < p^e$. Suppose this is the case, so that $\pi'_d(p^e)$ is strictly negative. However, an individual firm will only deviate to a positive discount if $\Pi'(p^e)$ is strictly negative, and it may not be since the wedge $\Pi'(p^e) - \pi'_d(p^e)$ between total and directed marginal profits is strictly positive.[9] In effect, this wedge tells us that the marginal "commitment cost" of discounting is positive even evaluated at a discount of zero. Examining (14), we can see that this is true because D_d is strictly positive even in a candidate equilibrium with no discounting – that is, prices have already been distorted

[9]Recall that $\pi''_u(p^e) < 0$.

away from p_u^m by the presence of directed consumer demand, so additional distortions have a first-order (rather than second-order) effect. We summarize this point below.

Remark 4 *Suppose that when discounts are ruled out, the equilibrium retail price level is p^e. Suppose also that any firm would like to commit to offering directed consumers a price below p^e. The latter need not imply that discounts will be used, if they are permitted.*

Using expressions for $\pi_d(p_d)$, $\pi_u(p_u)$, D_d, and D_u, and p^e, a firm's incentive to discount below the common list price may be written

$$\Pi'(p^e) = 1 - G(\hat{v}) - \frac{g(\hat{v})}{f(\bar{v})}\frac{1-F(\bar{v})}{G(\hat{v})} - \frac{g(\hat{v})}{\pi_u''(p^e)}\frac{(1-G(\hat{v}))(1-F(\bar{v}))}{G(\hat{v})} \tag{16}$$

5.2.2 Special Case: Known products are like random products and price quotes are free

In an important special case, the incentive to discount simplifies dramatically. For most of the subsequent analysis, we assume the following.

Known products are not special $G(v) = F(v)$

Free price quotes $s_p = 0$

The first condition is simply that a consumer's match value at her known firm is drawn from the same distribution as her match value at any other firm. Thus we rule the possibility that she draws from a distribution with positive selection relative to $F(v)$, perhaps because she is exposed to the choices of people with tastes like hers. The second condition is that getting a firm price quote from a known firm is costless – here we would remind the reader of our earlier caveat that $s_p = 0$ should be thought of as a clean, limiting case for results when s_p is positive but small. The first assumption will have a substantive impact on our results; the second, less so – its practical significance is that the directed and undirected thresholds for a purchase will collapse to a single threshold $\hat{v} = \bar{v}$.

Under these assumptions, the incentive to discount away from a (candidate symmetric equilibrium) list price reduces to

$$\Pi'(p^e) = -\frac{(1-F(\bar{v}))^2}{F(\bar{v})}\left(1 + \frac{f(\bar{v})}{\pi_u''(p^e)}\right)$$

Thus, an individual firm will (will not) have an incentive to discount away from the common list price if the term in parentheses is positive (negative). Because undirected profits $\pi_u(p^e)$ were assumed concave, we have:

Result 1 *At a candidate symmetric equilibrium with no discounting and list price* p^e*, a firm will have an incentive to offer a discount if* $|\pi_u''(p^e)| > f(\bar{v})$*. It will prefer to offer no discount if* $|\pi_u''(p^e)| < f(\bar{v})$*.*

Notice that $f(\bar{v})$ is the marginal increase in directed arrivals induced by a marginal discount; as discussed earlier, this drives the temptation to hike one's list price. And $|\pi_u''(p^e)|$ reflects, in effect, the strength of the firm's commitment *not* to hike its list price. When the latter outweighs the former, a deviation to discounting will be worthwhile. In this case, any symmetric equilibrium *must* involve discounted offers.

Using expressions for undirected demand and the equilibrium list price p^e, we can reframe this condition in terms of the primitives of the model: the distribution of match values. Note that $\pi_u''(p^e) = -(2f(\bar{v}) + p^e f'(\bar{v}))$, so the condition for discounting to occur may be written as $f(\bar{v}) + p^e f'(\bar{v}) > 0$, where the price is given by $p^e = (1 - F(\bar{v})) / (F(\bar{v}) f(\bar{v}))$. Given the assumptions we have made, either group's demand (directed or undirected) at a price p is given by $1 - F(\bar{v} - p^e + p)$, so $f(\bar{v})$ is the slope of that demand curve, evaluated at the candidate equilibrium list price. Then the condition for discounting can be given an economic interpretation.

Result 2 *At a candidate symmetric equilibrium with no discounting and list price* p^e*, a firm will have an incentive to offer a discount if demand is concave, or if demand is convex but the price elasticity of its slope is less than one. Otherwise a firm will prefer not to discount.*

Concave demand implies $f'(\bar{v}) > 0$ which is sufficient for the condition above. The price elasticity of the slope of demand, evaluated at $p = p^e$, is $\eta = \frac{p^e}{f(\bar{v})} f'(\bar{v})$, so if demand is convex, $|\eta| < 1$ implies $f(\bar{v}) + p^e f'(\bar{v}) > 0$. To summarize, firms will wish to discount if demand (evaluated at the candidate equilibrium price) is not *too* convex. The condition on η does not appear to map naturally into more familiar conditions on demand (such as logconcavity or ρ-concavity of demand or of the density $f(v)$). We will give an example later for which demand and $f(v)$ are both convex but logconcave, but the condition $|\eta| < 1$ is not met.

Examples

Uniform match values, linear demand If match values are uniformly distributed, then demand is (weakly) concave, so any equilibrium must involve discounting.

Max from a uniform sample Suppose that $F(v) = v^a$ for $a > 1$. This family of distributions might be of interest if, a consumer's match value at a firm were to represent the best of a number of uniform draws. For example, if each firm offers n products, with match values for each product distributed $U[0,1]$, then a consumer's best match upon visiting the firm would be distributed as $F(v) = v^n$. In this case, demand is strictly concave, so any equilibrium involves discounting.

Exponentially distributed match values Demand is convex, and the slope elasticity reduces to $|\eta| = \frac{1}{F(\bar{v})}$, so an equilibrium with no discounting generally exists.

Logistic match values Demand is convex when consumers are sufficiently choosy (s small, hence $\bar{v}$ large). But the slope elasticity reduces to $\eta = \frac{1-2F(\bar{v})}{F(\bar{v})^2} = \frac{(1-F(\bar{v}))^2}{F(\bar{v})^2} - 1 > -1$. So any equilibrium must involve discounting.

Normally distributed match values Again, demand is convex when consumers are sufficiently choosy. But numerical investigation suggests that any equilibrium must involve discounting.

Match values with a decreasing, triangular density Suppose $F(v) = 2v - v^2$, with $f(v) = 2 - 2v$, and $1 - F(v) = (1-v)^2$ for $v \in [0,1]$. Then demand is convex but $\frac{1}{2}$-concave, and its slope is weakly concave. However, if consumers are not very choosy, firms will prefer not to discount. In particular, firms will not discount if $\bar{v} < 1 - \frac{\sqrt{2}}{2} \approx 0.29$ (but will wish to discount otherwise).

5.3 Equilibria with Discounting

This section attempts to characterize symmetric equilibria when discounts are used. The efforts here are preliminary: we will restrict attention to the special case in which $G(v) = F(v)$ and $s_p = 0$, and even so we shall focus mainly on computed examples for specific distributions.

A necessary condition for a symmetric equilibrium in which all firms charge list price $\hat{p}$ and offer the (credible) discount price p_d to directed consumers is that the first-order condition (14) holds at $\hat{p}$ and p_d and that $\hat{p}$ is consistent with (8). The last condition requires the list price to satisfy

$$\hat{p} = \frac{1 - F(\bar{v})}{F(\hat{v}) f(\bar{v})} \tag{17}$$

where the directed consumer threshold now satisfies $\hat{v} = \bar{v} - d = \bar{v} - \hat{p} + p_d$. In this case, it is convenient to treat the marginal directed consumer $\hat{v}$ as the "driving" variable, with p_d and $\hat{p}$ determined by $\bar{v}$ and $\hat{v}$. Notice that $\Pi'(p_d) = \Pi'(\hat{v})$, since the discount price and $\hat{v}$ move one-for-one. Then with some substitution, the first-order condition (14) is given by

$$\Pi'(\hat{v}) = \pi_d'(\hat{v}) - \frac{D_d}{D_u}\frac{f(\hat{v})}{\pi_u''(\hat{p})} = 0 \tag{18}$$

where $D_d = 1 - F(\hat{v})$, $D_u = \frac{F(\hat{v})}{1-F(\bar{v})}$, $\pi'_d(\hat{v}) = D_d - p_d f(\hat{v})$, $\pi''_u(\hat{p}) = -(2f(\bar{v}) + \hat{p}f'(\bar{v}))$, and prices are given by $p_d = \hat{p} - (\bar{v} - \hat{v})$ and (17).

5.3.1 Example: Uniformly Distributed Match Values

To illustrate some general patterns, we consider an example where the distribution of match values is uniform: let $v \sim U[0,1]$, so $F(v) = v$. With a bit of algebra the first-order condition (18) can be shown to have the solutions

$$\hat{v} = \frac{3\bar{v}+1}{8} \pm \frac{1}{8}\sqrt{(9\bar{v}-5)(\bar{v}+3)}$$

for $\bar{v} \geq \frac{5}{9}$; for $\bar{v} < \frac{5}{9}$, $\Pi'(\hat{v})$ is strictly negative on $[0, \bar{v}]$. While both solutions do appear to represent equilibria, we focus on the larger one, as an informal dynamic analysis indicates the equilibrium corresponding to the smaller solution is unstable. For this distribution, the relationship between $\bar{v}$ and the search cost s is given by $s = \frac{1}{2}(1-\bar{v})^2$. In the figures below we plot equilibrium prices, utility, and profits versus s. As a benchmark for comparison, we also plot the corresponding equilibrium outcomes when discounting is forbidden.

Both retail ($\hat{p}$) and discount (p_d) prices are increasing in the search cost; the gap between them grows with s, indicating larger discounts as search grows more costly. Retail prices are always higher with discounting than they would be without, but for s sufficiently large, even consumers who receive a discount pay a higher net price than they would have if discounting were impossible. At $s = \frac{8}{81}$ (corresponding to $\bar{v} = \frac{5}{9}$), both prices diverge. This is the point at which there are no longer enough undirected searchers to keep firms honest with their retail pricing: both discounts and prices spiral out of control. Thus far we have hand-waved over a consumer's decision to engage in search versus opting out of purchase by treating the outside option u_0 as arbitrarily small, so as to rule out the latter. We cannot be so cavalier here: for any fixed u_0, the divergence of prices implies that a search equilibrium fails to exist, and the market fails, as search costs approach and exceed $s = \frac{8}{81}$. This market failure is driven by discounting; if discounting is ruled out, the equilibrium price continues to rise gradually with s.

Discounting drives a wedge between the utility of directed consumers – those with a known product good enough to buy immediately ($v_1 \geq \hat{v}$) – and undirected consumers (everyone else). However, overall *ex ante* expected utility is lower with discounting than without. Furthermore, when search costs are high enough, even the lucky consumers who know about a product that is good enough to buy are worse-off when firms are able to offer them discounts.

In contrast, overall expected profits are greater with discounting than without. As the figure indicates, discounting does permit a firm to operate at a more profitable point on its

Figure 1

demand curve for directed consumers, as compared to undirected consumers. But at least in this example, the difference is small. The bulk of the improvement in overall profits, relative to a no-discounting regime, comes from the general equilibrium upward drift in retail prices that discounts instigate.

6 Concluding Remarks

[That's all, folks.]

References

[1] Anderson, S., Renault, R.: Advertising content. Amer. Econ. Rev. 96, 93-113 (2006)

[2] Anderson, S., Renault, R.: Pricing, product diversity, and search costs: a Bertrand-Chamberlin-Diamond model. RAND J. Econ. 30, 719-35 (1999).

[3] Arbatskaya, M.: Ordered search. RAND J. Econ. 38, 119-127 (2007)

[4] Armstrong, M., Vickers, J., Zhou, J.: Prominence and consumer search. RAND J. Econ. 40, 209-233 (2009)

[5] Bagnoli, M., Bergstrom, T.: Log-concave probability and its applications. Econ. Theory. 26, 445-469 (2005)

[6] Caplin, A., Nalebuff, B.: Aggregation and imperfect competition: on the existence of an equilibrium. Econometrica. 59, 25-59 (1991)

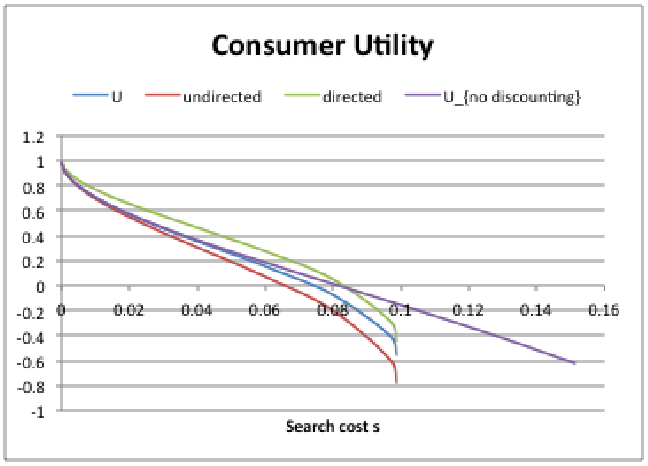

Figure 2

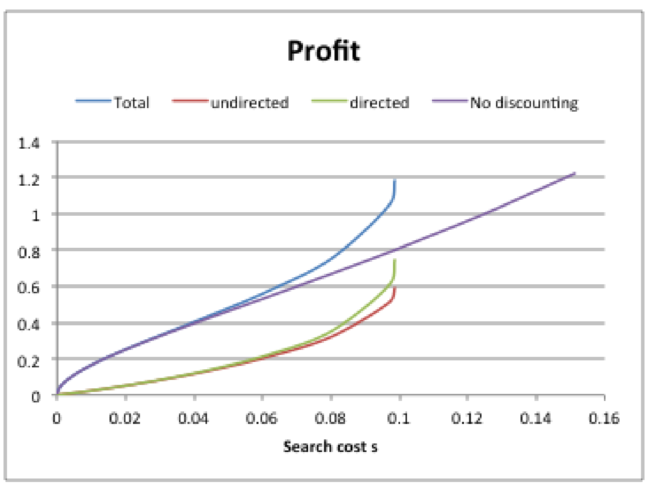

Figure 3

[7] Christou, C., Vettas, N.: On informative advertising and product differentiation. Int. J. Ind. Organ. 26, 92-112 (2008)

[8] Diamond, P.: A model of price adjustment. J. Econ. Theory. 3, 156-68 (1971)

[9] Haan, M., Moraga Gonzalez, J.: Competing for attention in a consumer search model. Econ. J. forthcoming.

[10] Larson, N. Niche products, generic products, and consumer search. Economic Theory. 52, 793-832 (2013).

[11] Wolinsky, A.: True monopolistic competition as a result of imperfect information. Quart. J. Econ. 96, 493-511 (1986)