Raising Search Costs To Deter Window Shopping Can Increase Profits and Welfare[*]

Greg Taylor[†]

ABSTRACT: Consumers tend to browse products in which they are interested and firms often choose to invest resources in selling to those consumers. A consequence, as I show, is it is optimal for a firm to increase the cost of browsing (even though this drives away potential customers) because doing so allows it to target its sales efforts at those consumers most likely to buy. Despite representing pure waste, this can increase welfare by facilitating the efficient allocation of sales or marketing resources. For a similar reason, consumers often benefit from search costs in aggregate, and prefer them to other means of screening, such as increases in price.

1   Introduction

This article offers a new perspective on the effects of consumer search costs. Recent work has focused on the deliberate raising of search costs (often called "obfuscation" in the literature) to deter search and soften price competition. I study an alternative motivation for creating synthetic search costs, namely to screen for the most valuable consumers and direct sales efforts at them.[1] The article makes two main contributions. Firstly, I show that when a firm can invest in selling to consumers it will wish to impose wasteful search costs to avoid dissipating sales resources on people unlikely to buy. Secondly, and more importantly, I examine the welfare implications of these search frictions. Even when they represent pure waste, an increase in search costs can be welfare-improving. Indeed, transparency or consumer protection policies that eliminate search frictions can reduce both consumer surplus and total welfare. The reason is that a firm invests more in sales assistance when it attracts only serious customers and these investments facilitate the creation of surplus through trade. These results reverse the dominant narrative in the literature, which has mostly neglected the issue of welfare or else shown high search costs to be unambiguously harmful.

Consumers often inspect a product before they buy it. They may do this by a variety of means including visiting a shop, making an inquiry by telephone or email, or reading

[1]Marketers use the term "qualification" for the process of identifying consumers who are the best targets for sales efforts.

product descriptions on a firm's website. By observing such browsing behavior, a firm can generate a list of potential customers. In many industries it is common to expend resources to convert these leads into realized sales. For example, a store might employ sales staff to approach browsing consumers, send email reminders of products consumers left unbought in their virtual shopping baskets, or provide product demonstrations and free samples. This is particularly important in industries with complicated products, those with experience goods, those where bespoke products are customized during the sales process, and when consumers tend to forget or be distracted and must be prompted with a reminder to buy.[2]

Targeting sales efforts at browsers is appealing because the decision to browse is a valuable signal of interest in the product. But there are various reasons why this signal may not be especially informative. For instance, many consumers enjoy browsing (or "window shopping") as a leisure activity without having any intention of purchasing. Some do intend to purchase, but enter a store unlikely to sell a product that meets their needs. Similarly, consumers are increasingly "showrooming"—inspecting goods in a physical store with the express intention of purchasing (from a different seller) online. Such behaviors mean the pool of browsers contains many individuals who, because they are unlikely to buy, are of little value to the seller. This presents a dilemma for firms that can commit resources to converting browsers into buyers, but do not know ex ante who the serious potential customers are. One possible solution to this dilemma is to increase the cost of browsing. There are various ways that this could be achieved:

- Many firms require that consumers proactively register or contact the seller to learn details about a product such as its price, availability, or shipping arrangements. This process can be made more or less onerous.

- Firms can introduce material or administrative barriers to browsing. For example, *Car and Driver* magazine advises its readers, "If it's a rare or particularly expensive car you want to test, visit the dealership first to 'qualify' and expect to return later to drive. 'Obviously, testing a Cobalt is no problem,' explains Chevrolet sales manager Eric Gilliam, 'but testing a Corvette Z06 or a Cadillac CTS-V is.'"[3]

- Browsing can be made unwelcoming or intimidating. Mid-range clothing retailers typically occupy stores with large windows, broad open doors, and prominently marked prices, whereas high-end fashion boutiques often have obscured windows and narrow doorways intimidatingly staffed by a doorman.

---

[2]For example, some industry estimates suggest that less than 3% of online consumers buy on their first visit to a website, giving rise to so-called retargeted advertising—an important and ubiquitous form of online targeted advertising in which ads are specifically used to prompt past visitors to return.

[3]*Car and Driver*, "How to Get That New-Car Test Drive" available at http://www.caranddriver.com/features/how-to-get-that-new-car-test-drive-feature-leave-your-ego-at-home-page-4, retrieved 26 January 2015.

- A representative of fashion retailer Vera Wang remarked "A lot of high school and college students were coming here and weren't serious about buying a wedding dress so that's why we started the [\$500] fitting fee."[4] More generally, some sellers, such as art exhibitors, directly levy a fee for browsing. The analysis below implies that this would be optimal even if the resulting revenue is immediately burned because it allows the artists to spend more time talking to people likely to buy their work.

If it is made more costly to browse then consumers that only have a passing interest are likely to be put-off. There is thus a selection effect: the act of browsing becomes a stronger indicator of interest and the firm is able to target its sales efforts more effectively if search costs are higher. The trade-off is that high search costs exclude some consumers who would otherwise have purchased. The fact that search costs do not exclude consumers uniformly, but rather selectively exclude those with the lowest interest in the product guarantees that some strictly positive search cost is always optimal—a result I establish in various settings. Niche or luxury firms (which target a small but valuable fraction of the consumer population) find screening via search costs to be especially attractive. Firms persist in setting positive search costs under competition, although the presence of a competitor screens-out some disinterested consumers and reduces the need for screening via search costs.

As well as increasing profits, a perhaps surprising result is that search costs can enhance consumer welfare—even accounting for the direct cost imposed on consumers. This is because firms are, consequently, prepared to spend more on assisting consumers. Search costs therefore promote firm investment in activities that facilitate surplus-creating transactions. For example, purveyors of luxury goods often take measures to deter all but the most exclusive clientèle, but lavish attention on those browsers they do attract. By the same token, search costs can increase overall social welfare. Indeed, if the search cost is a fee levied by the firm (rather than pure waste) then the socially optimal search cost is unambiguously positive.

Contribution to the literature

Regulators (*e.g.*, the UK's Office of Fair Trading and the US Federal Trade Commission) and scholars (*e.g.*, Ellison and Ellison, 2009) have raised concerns about strategies to create synthetic search frictions. A key puzzle is this: although it is natural to think that a firm would benefit if the cost of searching its rivals' products increased, why should a firm want to increase the cost of searching its own? Answers include (1) that a time-consuming initial search can wear a consumer out (Ellison and Wolitzky, 2012; Carlin and Ederer,

---

[4]See http://www.marieclaire.co.uk/news/fashion/541731/vera-wang-scraps-appointment-fee-in-new-bridal-store.html, accessed 3rd May 2016. This fee has since been abolished, but the popular press is replete with similar examples of retailers—from shoe stores to sellers of ski equipment or dietary supplements—experimenting with such fees.

2012) or exhaust a consumer's time budget (Hämäläinen, 2016) and therefore deter further search; (2) that obfuscation can cause consumers to overestimate the costs of future search (Ellison and Wolitzky, 2012); (3) that diverting high-search-cost consumers to a rival reduces its demand elasticity and thereby softens competition (Wilson, 2010); (4) that costly return may make a consumer reluctant to leave a firm in the first place (Armstrong and Zhou, 2016); (5) that obfuscation may confuse consumers and make price comparisons harder (Chioveanu and Zhou, 2013; Gu and Wenzel, 2014; Piccione and Spiegler, 2012); or (6) that it can help a firm to commit not to discount its products (Ellison, 2005).

The above papers study an environment in which search costs soften competition and increase profits by reducing consumers' propensity (or ability) to search and compare prices across firms. As noted, I focus on a different reason—namely to better allocate sales resources—and I show that this persists even under monopoly. Thus, the basic mechanism studied in this article is fundamentally different to those present in the papers above. Crudely, search costs in this article can be thought of as forming a part of the firm's informational strategy, rather than its competitive strategy.

Several other papers consider the relationship between search costs and the firm's information problem. In Petrikaite (2015) a firm that sells multiple products sets a different search cost for each, inducing consumers to search low-cost products first. Consumers stop searching earlier when they realize higher valuations so stopping behavior reveals information about willingness to pay and thereby facilitates price discrimination.[5] Petrikaite can be thought of as examining the screening role of search costs within a store, whereas this article considers their role in screening consumers before they enter the store. More substantively, the mechanisms and intuition in the two papers are different: Petrikaite (2015) relies on multi-product retailing and price discrimination (which play no substantive part in my analysis), whereas this article's results are driven by costly sales and ex ante private information (not considered by Petrikaite, 2015).

Shin (2005) considers a model where selling to consumers involves an exogenous cost. Thus, in common with this article, the firm has an interest in identifying the most valuable consumers. In Shin (2005), though, search costs are fixed and firms instead use cheap talk advertising to develop an image as a high-price firm in order to drive away low-value consumers. Thus, Shin's model is one of signalling through advertising, whereas the present article is concerned with the use of search costs to screen consumers. Significantly, I allow the firm to vary its investment in consumer assistance and show that endogenous changes in this investment have important new implications for the effects of search costs. Bar-Isaac, Caruana, and Cuñat (2010) consider the incentive of firms to provide information to consumers that will help them determine their value for a product. A firm

---

[5]In a similar vein, Salop (1977) considers a monopolist that randomizes its prices. High-search-cost consumers sample few prices and pay a high price, whereas low-cost consumers sample many and are rewarded with a lower purchase price. Thus, randomization facilitates price discrimination.

may want to partially inform consumers to reduce the dispersion of consumers' willingness to pay and facilitate surplus extraction, and to help it commit to investing in quality. In contrast, my focus is on the use of search costs to help a firm determine *which* consumers it should "inform." The commitment mechanism addressed by Bar-Isaac, Caruana, and Cuñat (2010) does not play an important role in the results below.

Thus, whereas some papers share elements in common with the model below, the basic mechanism that I consider (involving the creation of search costs, private information about demand, and costly sales effort) is distinct. The first contribution made by this article is, then, to show that consumer browsing and costly sales jointly imply a novel incentive to increase search costs, and to thereby provide a new explanation for various empirically-observed practices.

More importantly, this article also makes a normative contribution by providing a novel welfare analysis of search costs. The papers cited above either do not consider questions of welfare, or show increases in search costs to be harmful for consumers. The dominant message from the literature is, therefore, that increasing search costs is an anti-competitive or exploitative practice. An important contribution of this article is to provide a fresh look at welfare and to show how search costs, despite being wasteful, can help to align consumers' and firms' interests and thereby increase consumer and total surplus. This analysis underlines the importance for policy of considering the strategic purpose of search costs and the indirect consequences that their elimination may have.

The literature on obfuscation grows out of a well-established body of work on consumer search. These articles generally take the cost of search to be a given parameter rather than a strategic choice made by the firm. A fairly common result (see, *e.g.*, Diamond, 1971; Reinganum, 1979; Stahl, 1989) is that (higher) search costs unambiguously harm consumers.[6] Less directly, as one increases the proportion of uninformed consumers (which can be thought of as a reduced form way of modelling higher search costs) in models such as Stahl (1989) and Varian (1980), the result is also typically an increase in price that harms consumers. As a consequence, a reading of this literature would tend to lead one broadly to the conclusion that (i) higher search frictions are bad for consumers and for overall welfare; and (ii) that consumers' interests are often opposed to those of firms (for whom search frictions are a frequently a source of profitability). This article presents a new perspective on both points.

Several other papers are topically related, but address a different set of issues. A literature on targeted advertising takes as given a technology for learning about consumers' types and then examines how firms can use this information in their marketing (see *e.g.*, Athey and Gans, 2010; Bergemann and Bonatti, 2011; Johnson, 2013). The current article,

---

[6]Armstrong and Zhou (2011), Janssen, Moraga-González, and Wildenbeest (2005), Moraga-González, Sándor, and Wildenbeest (forthcoming), and Zhou (2014) show that prices can decrease in search costs for some configurations of the market.

in contrast, is concerned with how a firm can use search costs to generate such information to begin with, and what the welfare consequences of it doing so are. Piccolo, Tedeschi, and Ursino (2015) show that prohibiting deceptive advertising may harm consumers by increasing apparent quality differentiation and weakening price competition. Similarly, Rhodes and Wilson (2016) show how deceptive advertising can reduce prices by lowering consumers' quality expectations—with a potentially positive net benefit to consumers.

Lastly, this article is related to the broader literature on information economics. Mussa and Rosen (1978) provide a classic analysis of monopoly screening through quality menus. Unlike Mussa and Rosen (1978), the firm in my model must invest in selling to consumers before knowing whether they will buy or not and uses a wasteful cost as a screening device, so the mechanism is different. In Mussa and Rosen (1978) screening can be welfare-improving when it leads to the serving of consumers who would otherwise be excluded; in my model, somewhat ironically, screening improves welfare by excluding consumers who would otherwise be served.

2   Baseline Model Description

A single firm sells a product to a unit mass of consumers. Normalize its marginal cost to zero. The firm has three decision variables: its price, $p \in \mathbb{R}_+$; the cost of browsing its products, $s \in \mathbb{R}_+$; and a sales intensity, $\mu \in \mathbb{R}_+$. If the total mass of browsing consumers is m then the firm's cost is $C(\mu m)$, which is twice differentiable. Let $C'(0) = C(0) = 0$, and $C''(\cdot) > 0$.

Each consumer is "interested" in the firm's product with idiosyncratic probability π, which is private information and is distributed in the population according to the distribution function $G \colon [0,1] \to [0,1]$. Let g denote the corresponding density function, which is positive and finite everywhere on the support $[0,1]$. Consumers begin by observing their π and the firm's choice of s, but must browse the firm (at cost s) to learn the price and whether or not the the firm sells a product of interest. If the firm's product is of interest then a consumer demands $D(p,\mu)$ units of the good, with $D_1(p,\mu) < 0$,[7] and $\int_0^\infty D(p,\mu)\,dp < \infty$.[8] I impose the standard requirement that D satisfies the monotone hazard rate property (is log-concave) with respect to p—this ensures that any μ is paired with a unique monopoly price. I also assume that $D_2(p,\mu) \geq 0$ (with strict inequality when $\mu = 0$) and $D_{22}(p,\mu) \leq 0$: demand is increasing and concave in the intensity of sales efforts. In the event that the product is not of interest, demand is zero.

The game proceeds as follows:

---

[7]D is assumed to be differentiable in both arguments over the relevant range; $f_i(\cdot,\ldots,\cdot)$ means the derivative of function f with respect to its i^{th} argument.

[8]With the exception of Section 5, it would be formally equivalent to suppose that consumers had unit demand and willingness to pay $v \sim F(v,\mu)$, which they learn upon browsing. We could then write $E[D(p,\mu)] = 1 - F(p,\mu)$.

1. The firm sets $s \in \mathbb{R}_+$, which is immediately publicly observed.

2. Consumers observe π, and choose whether to browse the firm. The firm quotes a price, $p \in \mathbb{R}_+$, and decides the intensity, $\mu \in \mathbb{R}_+$, with which it pursues sales.

3. Demand and pay-offs are realized.

I look for sub-game perfect Nash equilibria of this game.

Conditional on finding an item of interest, a consumer's benefit from browsing is the surplus obtained from consuming the product:

$$\sigma(p, \mu) \equiv \int_p^\infty D(z, \mu)\, dz.$$

The ex ante expected utility from browsing is thus given by $\pi\sigma(p, \mu) - s$.

Given consumers' decision whether to browse or not, the firm's profit is

$$\Pi(s, \mu, p) = E(\pi|\text{browse})mD(p, \mu)p - C(\mu m). \tag{1}$$

For a given s, let $p^*(s)$ and $\mu^*(s)$ denote the equilibrium p and μ, and define $\sigma^*(s) \equiv \sigma(p^*(s), \mu^*(s))$. When $\Pi(s)$ is written only as a function of s, it is to be understood that this is the function obtained by evaluating $\Pi(s, \mu^*(s), p^*(s))$.

3   ANALYSIS

Optimal sales intensity: search costs screen interest

In this subsection I begin the process of backward induction—solving for the optimal strategies at $t = 2$—and thereby address the question, what can the firm learn about a consumer from the decision to browse?

Consumer browsing decisions: A consumer's net expected utility from browsing is $\pi\sigma^*(s) - s$ and consumers therefore browse if and only if $\pi \geq s/\sigma^*(s) \equiv \bar{\pi}(s)$. Thus, consumers' browsing decisions follow a cut-off rule such that any consumer with $\pi > \bar{\pi}(s)$ chooses to browse. The mass of browsing consumers is $m(s) \equiv 1 - G(\bar{\pi}(s))$. For brevity, I will usually suppress the argument and write $m \equiv 1 - G(\bar{\pi})$.

Optimal price and sales intensity: The firm chooses p and μ to maximize (1). Differentiating (1) with respect to p reveals that the profit-maximizing price is the monopoly price solving

$$p^*(\mu) = -\frac{D(p^*, \mu)}{D_1(p^*, \mu)}. \tag{2}$$

Taking $p = p^*$ as given, (1) becomes

$$\Pi(s, \mu, p^*(\mu)) = E(\pi | \pi > \bar{\pi}) m D(p^*(\mu), \mu) p^*(\mu) - C(\mu m).$$

Using the envelope theorem, the optimal sales policy is characterized by the first-order condition

$$E(\pi | \pi > \bar{\pi}) D_2(p^*, \mu^*) p^* - C'(m\mu^*) = 0. \tag{3}$$

I make the following assumption about the firm's optimisation problem (a discussion follows shortly):

Assumption 1. *For any $\bar{\pi}$ there is a unique profit-maximizing pair (p^*, μ^*).*

Taken together, μ^*, p^*, and $\bar{\pi}$ constitute the equilibrium of the sub-game beginning at $t = 2$ for any given s. Letting $r = D(p^*, \mu^*) p^*$, we can define

$$R(s) \equiv E(\pi | \pi > \bar{\pi}) r = \frac{r}{1 - G(\bar{\pi})} \int_{\bar{\pi}}^1 \pi g(\pi) \, d\pi.$$

$R(s)$ is the revenue value to the firm of having the opportunity to trade with one browsing consumer.

To build intuition about the effects of increasing search costs, it is useful to consider how such an increase affects this revenue along with consumer behavior. Suppose that the following assumption (requiring that consumers experience a net benefit from investments in sales) holds:

Assumption 2. $d\sigma(p^*(\mu), \mu)/d\mu > 0$ *over the relevant range.*

The following result then demonstrates that the firm earns more revenue per browsing consumer and invests more in sales when s is larger.

Lemma 1. *Suppose Assumptions 1 and 2 hold. Then $dR(s)/ds > 0$, $dm/ds < 0$, and $d\mu^*/ds \geq 0$ (with strict inequality if $D_2(p^*, \mu^*) > 0$).*

Proofs are in the appendix. The main work of the proof of Lemma 1 is in establishing that an equilibrium can never be such that a higher s induces more consumers to browse the firm—say, because they expect high s firms to offer a significantly higher sales intensity. When s is positive, serious customers (with a high π) still have much to gain from visiting the store, whereas low-π consumers are especially loath to window shop when doing so is costly. The fact that a consumer is prepared to incur a larger browsing cost therefore conveys a stronger (positive) signal about their interest in the product. As the average level of interest among browsing consumers increases, sales efforts are less often wasted and the firm invests more in them.

Remark on Assumptions 1 and 2: Both of Assumptions 1 and 2 are satisfied by any demand function that can be written in the multiplicatively separable form $D(p, \mu) = h(\mu)D(p)$.[9] One natural interpretation of such demand is that consumers are only able to buy with a probability, h, that depends upon the firm's investment in sales (because, for example, some consumers forget and must be prompted with a reminder to buy). More generally, this form reflects consumers with an intrinsic willingness to pay, $D(p)$, that is scaled by the firm's sales efforts.

Both Assumption 1 and Assumption 2 are also satisfied by various non-separable demand functions. Consider, for example, the linear form $D(p, \mu) = 1 - p/\mu$ or exponential demand $(D(p, \mu) = e^{-p/\mu})$. Both are non-separable and give rise to a p^* that is strictly increasing in μ. Nevertheless, both demand functions satisfy Assumptions 1 and 2.[10]

In general, Assumption 1 will also hold provided that costs are sufficiently convex to make the left-hand side of (3) decreasing in μ. Assumption 2 holds if a small increase in sales effort does not lead to too large an increase in price. This will be true if the right-hand side of (2) does not increase too quickly in μ.

To avoid ambiguity, I will explicitly note when either of Assumptions 1 and 2 are used for a result.

Optimal search cost policy

In Lemma 1 we saw that there can be an informational value to search costs through their ability to screen consumer interest, but that this is coupled with a reduction in the number of consumers visiting the firm (and hence its potential demand). It is therefore interesting to ask: (when) would a firm like to impose a cost for browsing its products? Proposition 1 now establishes that, for any $D(p, \mu)$ satisfying Assumption 1, the firm's profit is maximized by a strictly positive search cost.

Proposition 1. *Given Assumption 1 the firm optimally sets a strictly positive search cost, $s^* > 0$.*

The key to Proposition 1 is the fact that search costs do not exclude consumers uniformly, but rather selectively screen out those with the lowest interest in the product. Once the pool of consumers has been selected for higher levels of interest, the firm can achieve the same result with lower expenditure on marketing and sales resources because these resources will be less often wasted on disinterested consumers. To better understand the trade-off faced by the firm, frame the firm's first-stage problem as one of choosing

---

[9]Under multiplicative separability p^* uniquely solves $p^* = -D(p^*)/D'(p^*)$, and $h''(\mu) \leq 0$ ensures the left-hand side of (3) is decreasing in μ. Thus, Assumption 1 is satisfied. Moreover, $\sigma = h(\mu) \int_{p^*}^{\infty} D(p) \, dp$ is increasing in μ so Assumption 2 is also satisfied.

[10]$D(p, \mu) = 1 - p/\mu$ yields $p^* = \mu/2$, $\mu^* = \kappa^{-1}(E(\pi|\pi > \bar{\pi})/4)/[1 - G(\bar{\pi})]$, and $\sigma = \int_{\mu/2}^{\mu} 1 - \frac{p}{\mu} \, dp = \frac{\mu}{8}$ where $\kappa^{-1}(\cdot)$ is the inverse of the firm's marginal cost function. Exponential D implies $p^* = \mu$, $\mu^* = \kappa^{-1}(E(\pi|\pi > \bar{\pi})/e)/[1 - G(\bar{\pi})]$ and $\sigma = \int_{\mu}^{\infty} e^{-p/\mu} \, dp = \frac{\mu}{e}$.

how many consumers to exclude. The proof of Proposition 1 establishes that the firm would like to exclude the marginal (type-$\overline{\pi}$) consumer if $\mu^* C'(\mu^* m) - \overline{\pi} D(p^*, \mu^*) p^* > 0$. The left-hand side of this condition is the difference between the sales resources saved on the marginal excluded consumer and the revenue forgone from that same consumer. We can substitute for $C'(\mu^* m)$ from (3) and the condition becomes

$$\mu^* E(\pi | \pi > \overline{\pi}) D_2(p^*, \mu^*) p^* - \overline{\pi} D(p^*, \mu^*) p^* > 0. \tag{4}$$

Thus, $d\Pi / d\overline{\pi}$ is positive if

$$\frac{D_2(p^*, \mu^*) \mu^*}{D(p^*, \mu^*)} > \frac{\overline{\pi}}{E(\pi | \pi > \overline{\pi})}. \tag{5}$$

The left-hand side of (5) is the sales-elasticity of demand (a measure of the efficacy of the firm's sales technology). The lower is the sales-elasticity of demand, the more the firm would prefer to provide a low level of sales assistance to many consumers rather than a high level of assistance to few; that is, avoid concentrating sales assistance.[11] Opposing this consideration is the fact that excluding consumers is a way to ensure sales efforts are not wasted on disinterested browsers, which augers for a *greater* concentration of sales efforts. Indeed, the greater is the disparity between the marginal consumer's interest and that of the average browser (*i.e.*, the smaller is the right-hand side of (5)), the larger is the gain by pursing a more exclusive strategy in which a higher level of sales assistance is provided to fewer, more interested consumers. The firm balances these competing forces to determine how many consumers to exclude.[12]

The sales-elasticity of demand will tend to be larger when sales efforts are a very important element of the consumption process (*i.e.*, when demand is initially low, but increases quickly in μ). This, in turn, makes excluding the marginal consumer more attractive. In the most sales-intensive cases, where the elasticity is close to 1, the firm optimally implements a search cost sufficient to exclude all but the most interested of consumers. It is no accident, then, that the examples of search-cost raising strategies mentioned in the introduction—be they couture wedding dresses, shoes, or high-end cars— are all products that rely intensively on sales assistance. A concentration of consumers around $\pi \approx 1$ will tend to drag the value of $E(\pi | \pi > \overline{\pi})$ up, making the right hand side of (5) smaller. Thus, niche or luxury firms targeting a select group of consumers who are much more likely to buy than average will have the strongest incentive to screen browsers.

The following Corollary provides an example of how these considerations translate

---

[11]Heuristically, suppose that every consumer has the same π. A sales elasticity less than 1 implies that halving m and doubling μ would reduce aggregate demand but leave costs unchanged, so profits fall.

[12]Existence of equilibrium is guaranteed because the left-hand side of (5) is less than 1 (by concavity of D) and the right-hand side is continuous, equal to zero when $\overline{\pi} = 0$, and equal to 1 when $\overline{\pi} = 1$. A sufficient condition for uniqueness is that the left-hand side crosses the right-hand side of (5) from above once, when viewed as functions of $\overline{\pi}$. This condition is satisfied by various examples presented below.

into an equilibrium.[13]

Corollary 1. *Suppose that $C(\mu) = \mu^k$, $G(\pi) = \pi$ (uniformly distributed interest), and $D(p, \mu) = \mu^a(1 - p)$, with $k > 1$ and $0 < a < 1$. Then, in the unique equilibrium, $\overline{\pi} = a/(2 - a)$, $p^* = 1/2$,*

$$s^* = \frac{a}{8(2-a)} \left[\frac{(2-2a)^{k-1}(2-a)^{2-k}k}{a/4} \right]^{\frac{a}{a-k}}, \quad and \quad \mu^* = \left[\frac{(2-2a)^{k-1}(2-a)^{2-k}k}{a/4} \right]^{\frac{1}{a-k}}.$$

Welfare consequences of search frictions

Creating unnecessary search frictions cannot be efficient from a first-best perspective because they impose a wasteful cost upon consumers that could be avoided were a social planner able to prevent some consumers from browsing. More relevant from a policy perspective is to consider the second-best case in which a regulator or social planner can influence the value of s (say, by prohibiting common obfuscatory practices)—with consumers and firms responding optimally to this prescription.

Even by this relaxed standard, equilibrium search costs will typically not be socially optimal because the firm does not fully internalize the consequences of search costs for consumers. Aggregate net consumer surplus is

$$\text{CS} = \int_{\overline{\pi}}^1 \left[\pi\sigma(p^*(\mu^*), \mu^*) - s \right] g(\pi)\, d\pi. \tag{6}$$

Thus, the effect of a marginal increase in s upon consumer surplus is

$$\frac{d\text{CS}}{ds} = \int_{\overline{\pi}}^1 \left(\pi \underbrace{\frac{d\sigma}{d\mu^*} \frac{\partial\mu^*}{\partial s}}_{\text{(iii)}} - \underbrace{1}_{\text{(ii)}} \right) g(\pi)\, d\pi - \underbrace{(\overline{\pi}\sigma - s)}_{\text{(i)}=0} g(\overline{\pi}) \frac{\partial\overline{\pi}}{\partial s}. \tag{7}$$

A higher s (i) excludes some consumers, and (ii) imposes a wasteful cost on those consumers who choose to browse. However, these harms are offset to the extent that (iii) excluding low-interest consumers provides the firm with a greater incentive to invest in sales—creating gains from trade for consumers that show up in the form of a higher σ. Overall welfare, W, is given by the sum of consumer surplus and firm profit.

An immediate question is whether the investment effects described above can be strong enough that consumers would benefit from the introduction of a positive search cost. From (3), the size of the firm's investment response depends upon how quickly $E(\pi|\pi > \overline{\pi})$ increases with $\overline{\pi}$. If there is a large cluster of "window shoppers" (consumers with π close to zero) who would be excluded by a small search cost then such a cost will produce

---

[13]The proof is straightforward and is omitted: $\overline{\pi}$ is found immediately from (5) and p^* from (2). Given $(\overline{\pi}, p^*)$, μ^* can be calculated from (3). This equilibrium μ and p imply a σ^*, and the value of s^* is found by solving $\overline{\pi} \equiv s/\sigma^*$ for s.

a significant increase in the average consumer's interest-level and thus prompt a large increase in μ. This intuition is formalized in the next result.

Proposition 2. *Suppose assumptions 1 and 2 hold. Then, fixing $E(\pi)$, consumer surplus and total welfare are maximized by a strictly positive search cost (s) if $g(0)$ is sufficiently large.*

For example, the conditions in Proposition 2 will be satisfied by a u-shaped distribution with enough mass in the tails.[14] A u-shaped distribution for π corresponds to cases where most consumers have little ex ante uncertainty about their *realized* level of interest (*i.e.*, they know for sure whether they are a $\pi \approx 1$ serious customer or a $\pi \approx 0$ window shopper). Intuitively, when consumers know themselves to be window shoppers they indirectly exert an externality on serious customers by reducing the profitability of sales efforts and thereby decreasing the firm's incentive to invest in them. A positive search cost helps to mitigate this externality.

From (3) it is clear that the firm's investment response will also depend upon the properties of its demand and cost functions. Building upon the proof of Proposition 2, the following corollary details the implications for consumer surplus.

Corollary 2. *If demand is multiplicatively separable ($D(p,\mu) = h(\mu)D(p)$) then, starting from $s = 0$, a small increase in search costs causes consumer surplus to rise if*

$$\left[\frac{E(\pi)C'(\mu_0)}{C''(\mu_0) - h''(\mu_0)D(p^*)p^*E(\pi)} \right] \frac{h'(\mu_0)}{h(\mu_0)} g(0) > 1, \qquad (8)$$

where $\mu_0 \equiv \mu^*(0)$ is the unique solution to $C'(\mu_0) = h'(\mu_0)E(\pi)D(p^*)p^*$.

If welfare-creating trade is heavily dependent upon sales assistance then it becomes crucial that sales are targeted as effectively as possible, which can be ensured by excluding those consumers least likely to be interested. Thus, from (8), if $h(\mu)$ is initially low but increases quickly then a positive s is more likely to be beneficial. More generally, we see that (8) is more likely to be satisfied if C is not too convex and h not too concave: under these conditions a marginal increase in consumers' average interest (thanks to screening) is met with a large increase in investment—sufficient to offset the welfare cost of wasteful expenditure on s. A condition very similar to (8) obtains for total welfare. It is omitted from the statement of the proposition because it generates little additional intuition, but details can be found in the proof.

Proposition 2 and Corollary 2 provide conditions that guarantee the optimal search cost is positive. But for a regulator considering a prohibition on cost-raising strategies,

---

[14]More formally, take, for example, $\pi \sim \text{Beta}(\epsilon, \epsilon)$. Then G is u-shaped with $E(\pi) = 1/2$ and $g(0)$ arbitrarily large as $\epsilon \to 0$. One can check that the right-hand side of (5) approaches $\overline{\pi}$ in the $\epsilon \to 0$ limit. Thus, uniqueness of equilibrium follows for any D such that the right-hand side of (5) has slope less than 1 (*e.g.*, $D(p,\mu) = D(p)\mu^a$, with $a < 1$).

a more relevant question is whether such an intervention does better or worse than the equilibrium s^* that would be chosen by the firm.

Proposition 3. *(i) The equilibrium s^* yields higher consumer surplus (and total welfare) than would a search cost of zero if*

$$\frac{\sigma(s^*)}{\sigma(0)} > \frac{E(\pi)}{[1 - G(\bar{\pi})][E(\pi|\pi > \bar{\pi}) - \bar{\pi}]}. \tag{9}$$

(ii) For example, in the specification of Corollary 1, expression (9) becomes

$$\left[\frac{\left(\frac{2-2a}{2-a}\right)^{k-1}}{\frac{2}{2-a}}\right]^{\frac{a}{a-k}} > \frac{(2-a)^2}{4(1-a)^2}. \tag{10}$$

The left-hand side of (9) is consumers' percentage gain from trade due to the firm's enhanced incentive to invest. If this is large enough, then consumers experience a net gain from equilibrium search costs. Thus, (10) is satisfied when the demand parameter a is sufficiently large and the cost parameter k is sufficiently small (*i.e.*, when demand is not too concave in μ and costs are not too convex): from the discussion following Proposition 2, we know that this is when the ability to screen consumers translates into a large increase in μ.

As a concrete example, Figure 1 shows how varying the mass of excluded consumers (*i.e.*, changing s) influences consumer surplus and welfare for particular values of a and k.[15] In this example the equilibrium s is high enough to exclude 82% of consumers, which might be seen as cause for concern. However, it is clear that consumer surplus and total welfare are both higher in equilibrium than would be the case if firms were forced to set $s = 0$.[16]

In Figure 1 we see that, while reducing search costs from their equilibrium level to zero would be harmful, a sufficiently small reduction in s (and hence $\bar{\pi}$) would increase both consumer surplus and total welfare. In other words, the firm sets s too high. This is a more general observation that is not restricted to the particular example plotted:

Proposition 4. *Suppose demand is multiplicatively separable: $D(p, \mu) = h(\mu)D(p)$. In equilibrium, a small reduction in search cost, s, would increase both consumer surplus and total welfare.*

Proposition 4 implies that policy makers are somewhat justified in seeking to protect consumers from unnecessary search costs. However, the above analysis also establishes that the legitimate role search costs play in the efficient allocation of consumer assistance should be a dimension of the policy debate. Indeed, a substantial reduction in the level of

---

[15]The vertical scale is omitted because it can be arbitrarily re-scaled via a parallel shift in $D(p)$.

[16]Note that Figure 1 also provides an example that satisfies the condition in Corollary 2.
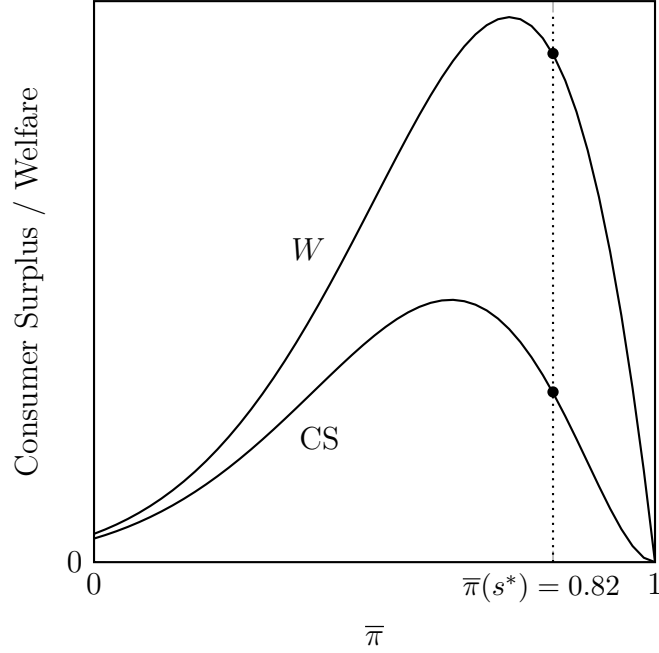
Figure 1: Consumer surplus and welfare as a function of the mass of excluded consumers when $a = 0.9$ and $k = 1$ in the specification of Corollary 1.

search costs will often dilute the pool of serious consumers (and thus the firm's incentive to invest) to such an extent that consumers are made worse-off. Proposition 2 and Corollary 2, in particular, tell us that—especially in contexts where consumers exhibit a propensity to window shop and where sales efforts form an important component of the consumption process—it will often not be optimal to completely eliminate search frictions.

Moreover, the informational burden for a regulator wishing to "fine-tune" s is significant (requiring knowledge of costs, demand, and the sales technology, as well as the ability to measure s and $\overline{\pi}$). In practice, policy makers usually rely on blunter instruments (such as blanket bans on particular cost-raising techniques or mandated transparency). Proposition 3 shows that, with such a policy toolbox, well-meaning interventions may well reduce consumer and total surplus.

The preceding discussion is framed in light of a model in which the firm is free to determine s. But the result that search costs can be welfare-enhancing would continue to hold were those costs the product of some exogenous process rather than the choice of the firm.

4   EXTENSIONS

Search cost menus

With enough flexibility in its sales instruments, the firm might be able to do better by offering a menu to consumers consisting of pairs (μ, s). I investigate such a possibility

when $D(p, \mu) = \mu D(p)$. Suppose that the firm can vary the intensity of sales on a consumer-by-consumer basis, where a larger μ is interpreted as a consumer receiving more intensive sales attention. Costs are additively separable across consumers and convex within each consumer, with $nC(\mu)$ being the cost of providing sales intensity μ to each of a mass n of consumers. Using the revelation principle, we can think of the firm as asking consumers to report a $\widehat{\pi}$ and then allocating to them the corresponding $(\mu(\widehat{\pi}), s(\widehat{\pi}))$.

Lemma 2. *The profit-maximizing menu has the form*

$$\mu(\pi) = \kappa^{-1}[\pi D(p^*)p^*]; \quad s(\pi) = \int_0^\pi \frac{zD(p^*)p^*}{C''\left(\kappa^{-1}[zD(p^*)p^*]\right)} \, dz \int_{p^*}^\infty D(p) \, dp$$

where $\kappa^{-1}(\cdot)$ is the inverse of the firm's marginal cost function.

This menu is such that consumers with a larger π incur greater search costs (to credibly signal their interest) and enjoy higher sales intensity in return. It is clear that the firm benefits from the ability to offer such a menu. The following result establishes that consumers can also be made better-off by this scheme.

Corollary 3. *Suppose $C(\mu) = \mu^2$ and $G(\pi) = \pi^\beta$. The firm's optimal menu of (positive) search costs yields higher consumer surplus and total welfare than does $s = 0$ if $(1 - 2\beta - \beta^2) > 0$; that is if $\beta \in [0, 0.41]$.*

Intuitively, if β is small then most consumers are relatively uninterested "window shoppers" and a firm will not find it worthwhile to invest much in sales if it is forced to provide the same service to everyone. Allowing the firm to offer a menu of search costs ensures that the most interested consumers are identified and receive the attention necessary to realize gains from trade, increasing consumer surplus and welfare.

Browsing fees

Instead of being pure waste, s might be a transfer from consumers to the firm. Indeed, this was the approach taken by Vera Wang with the fitting fee mentioned in the introduction. Such fees are also often used by the organizers of one-off events like art exhibitions or trade fairs, where putting sellers in contact with (serious) buyers during a narrow window of time is especially important.[17] A significant constraint on the more widespread use of such fees is likely to be psychological: consumers often react adversely to what they perceive as arbitrary or exploitative terms of trade (as a case in point, Vera Wang quickly abandoned its fitting fee after a media backlash).

---

[17]Other examples of fees can be found—especially if the interpretation of the model is expanded beyond the domain of consumer search. For instance, many journals charge a submission fee that should disproportionately dissuade authors with little confidence in their manuscript from submitting and free-up editorial resources for more promising manuscripts.

If $\Pi(s)$ is the firm's profit under a wasteful search cost then profit when s is a browsing fee is $\Pi(s) + ms$. Because ms does not depend upon μ or p, the firm's problem with respect to these variables is unchanged from the baseline model and is characterized by (2) and (3). Equilibrium search costs are then determined by solving $\Pi'(s) + m + (dm/ds)s = 0$, where $\Pi'(s)$ is as in the baseline model.[18] We know (by Proposition 1) that $\Pi'(0) > 0$. The following corollary is immediate:

Corollary 4. *Suppose s is a transfer from consumers to the firm. Given Assumption 1, the firm optimally sets a positive search cost, $s^* > 0$.*

Compared to the wasteful case, a browsing fee can result in either a higher or a lower equilibrium s: higher because the firm wishes to increase the browsing fee (which it gets to keep) and lower because driving away consumers reduces the number of times the browsing fee is paid. This is a standard monopolist's marginal versus infra-marginal trade-off, reflecting the fact that the search cost now amounts to a price.[19]

For a given s, consumer surplus (given by (6)) is the same regardless of whether s is waste or a transfer to the firm. Thus, the conditions described in Proposition 2 and Corollary 2 for consumer surplus to be increasing at $s = 0$ also hold verbatim in the model with browsing fees. The same is not true for results on total welfare. When search costs are pure waste they show up as a cost in the expression for total welfare. But under a browsing fee s is merely a transfer that does not directly detract from welfare; the only (social) cost from an increase in the browsing fee is the trade foregone with excluded consumers. Because the least interested consumers yield approximately zero trade surplus in expectation, they can be excluded without detriment to total surplus. This implies that welfare is always increasing in s at $s = 0$. Formally,

Proposition 5. *Suppose Assumption 1 and Assumption 2 hold. Then,*

1. *total welfare is maximized by a strictly positive browsing fee, s;*

2. *for $g(0)$ sufficiently large, consumer surplus is maximized by a strictly positive browsing fee, s.*

As in the model with wasteful search costs, the gains from improved allocation of sales resources can be large enough that the *equilibrium* browsing fee, s^*, yields higher consumer surplus and total welfare than would forcing $s = 0$. For example, if we take the

---

[18]An expression for $\Pi'(s)$ is given in (19).

[19]The supplementary appendix provides an example. When sales are very costly, all but the most interested consumers are excluded under a wasteful search cost. It would then be better for a firm charging a browsing fee to exclude fewer consumers in order to increase the number of times the fee is paid. Conversely, when serving consumers is cheap the firm favors a strategy of attracting many consumers (with a low search cost). But it is exactly when s is low that browsing fee income is increasing in s.
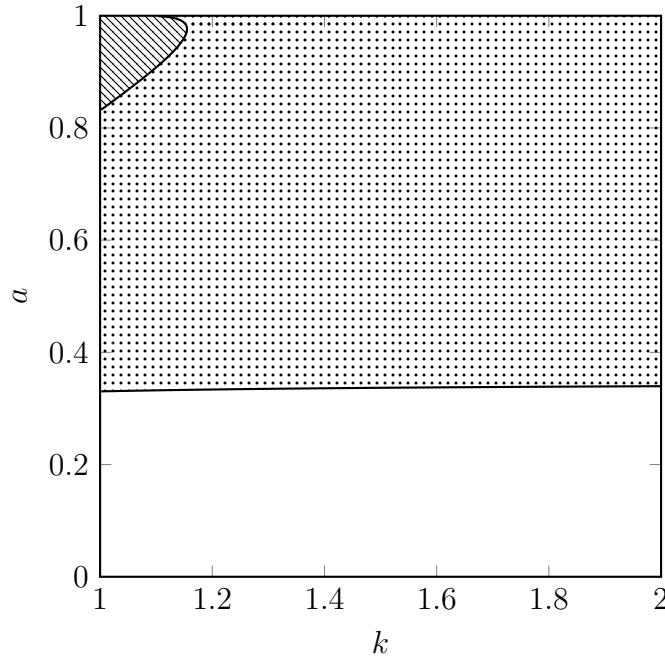
Figure 2: Using the specification of Corollary 1: compared to $s = 0$, the equilibrium browsing fee causes consumer surplus to increase in the hatched region, and causes total welfare to increase in the dotted and hatched regions.

specification of Corollary 1 then p^* and μ^* are found using (2) and (3) thus:

$$p^* = \frac{1}{2} \text{ and } \mu^* = \left[\frac{k[1 - G(\bar{\pi})]^{k-1}}{aD(p^*)p^*E(\pi|\pi > \bar{\pi})} \right]^{\frac{1}{a-k}} = \left[\frac{8k(1 - \bar{\pi})^{k-1}}{a(1 + \bar{\pi})} \right]^{\frac{1}{a-k}}. \tag{11}$$

In the first stage, the firm sets s to induce the $\bar{\pi}$ that maximizes $\Pi + ms$. From the proof of Proposition 1 we know that $d\Pi/d\bar{\pi} = g(\bar{\pi})[\mu^*C'(m\mu^*) - \bar{\pi}r]$. Moreover, we can write $ms = m\bar{\pi}\sigma$. Thus, the firm targets a $\bar{\pi}$ such that

$$g(\bar{\pi}) \left[\mu^*C'(m\mu^*) - \bar{\pi}r \right] + m\sigma - \bar{\pi}g(\bar{\pi})\sigma + m\bar{\pi}\frac{\partial \sigma}{\partial \mu}\frac{\partial \mu^*}{\partial \bar{\pi}} = 0;$$

that is,

$$\left[\mu^*k \left((1 - \bar{\pi})\mu^* \right)^{k-1} - \frac{1}{4}\bar{\pi}(\mu^*)^a \right] + \frac{1}{8} \left[(1 - 2\bar{\pi})(\mu^*)^a + (1 - \bar{\pi})\bar{\pi}a(\mu^*)^{a-1}\frac{\partial \mu^*}{\partial \bar{\pi}} \right] = 0.$$

Substituting in the value of μ^* from (11) yields an equation that, for any (k, a), can be solved for the optimal $\bar{\pi}$ (the optimal s^* is then found from $\bar{\pi} = s/\sigma$). Figure 2 compares the implied consumer surplus and total welfare under this equilibrium to that which would arise if we constrained $\bar{\pi} = s = 0$ (and let μ^* adjust accordingly). As in the wasteful search costs case, permitting equilibrium search costs is beneficial if a is not too small and (particularly in the case of consumer surplus) k is not too large.

5   OLIGOPOLY

An obvious question concerns whether the incentive to obfuscate is robust to the intro-
duction of competition. To examine this issue, extend the model to incorporate a (finite)
number, n, of firms. Each consumer is interested in a particular homogeneous product for
which he has demand $D(p)$, and which is sold by firm i with probability π_i.[20] A consumer
must incur a cost, s_i, to learn whether firm i sells his product of choice or not. If the
consumer does incur this cost and finds that i indeed sells the product he desires then
an opportunity to trade arises with probability $\mu_i \in [0, 1]$. Thus, a firm's choice of sales
intensity, μ, is interpreted as effort that increases the chance that a trade can take place
and demand has the form $D(p, \mu) = \mu D(p)$.

Let $g(\boldsymbol{\pi})\colon [0, 1]^n \to \mathbb{R}_+$ be the joint density of $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_n)$. Suppose that g
is symmetric (in the sense that switching two of its arguments leaves the value of g
unchanged) and positive on $[0, 1]^n$. Each firm can unilaterally determine its s_i, μ_i, p_i.

The timing of the modified game is as follows:

1. The firms set s_i, p_i, and μ_i.

2. Consumers observe $\mathbf{s} = (s_1, \ldots, s_n)$ and $\boldsymbol{\pi}$.

3. Consumers choose whether to browse a firm (and which firm to browse). Consumers'
   interest and the outcome of the firm's sales efforts are realized, and consumers learn
   the firm's price.

4. Consumers may (but need not) browse an additional firm, in which case steps 3–4
   are repeated.

5. Consumers decide whether (and where) to purchase (recall is free). Pay-offs are
   realized.

In equilibrium, (i) (s_i, μ_i, p_i) maximize Π_i given $(\mathbf{s}_{-i}, \boldsymbol{\mu}_{-i}, \mathbf{p}_{-i})$ and consumer search
behavior; and (ii) consumers search optimally given $(\mathbf{s}, \boldsymbol{\mu}, \mathbf{p})$.

The impossibility of frictionless equilibria

If $s_1 = \cdots = s_n = 0$ then it is (weakly) dominant for a consumer to search every firm
before making a purchase decision. I describe such a market outcome as "frictionless":

Definition 1 (Frictionless Equilibrium)**.** *An equilibrium is frictionless if $s_1 = \ldots = s_n = 0$
and all consumers browse every firm before making a purchase decision.*

---

[20]For example, a consumer wishing to purchase a copy of Tirole's *The Theory of Industrial Organization*
would have a high π for a campus academic bookshop and a low π for a smaller store that caters mostly
to a non-academic clientèle. The opposite would be true for a consumer looking for a mass-market
paperback. A consumer who is "just browsing" would have a low π for both stores.

Suppose that there exists a frictionless equilibrium. The model then has a similar structure to Varian (1980): from i's perspective, consumers who are not willing or able to purchase from any other firm (because, for example, they are not interested in any other firm's products) are like Varian's "uninformed" consumers over whom i has a de facto monopoly. Consumers with multiple firms in their consideration set are "shoppers" over whom the firms compete à la Bertrand. Firms have a simultaneous incentive to undercut (to attract shoppers) and to charge high prices (to exploit loyal consumers); the result is an equilibrium in mixed strategies.

Suppose that firms $2, \ldots, n$ set prices according to a putative symmetric equilibrium CDF, $F(p)$, and have $E(\mu_2) = \ldots = E(\mu_n) \equiv E(\mu)$ (the validity of this supposition to be verified later). We now study 1's best response. Write ψ_z for the event that, having visited every firm, a consumer has z firms other than firm 1 in his consideration set. Let $\rho_z = \Pr(\psi_z)$.[21]

The per-consumer revenue for 1 from a generic choice of p_1, conditional on $\mathbf{s} = \mathbf{0}$ and the consumer being able to buy from 1, are

$$R(p_1) = p_1 D(p_1) \sum_{z=0}^{n-1} \{\rho_z [1 - F(p_1)]^z\} E(\pi_1 | \psi_z). \tag{12}$$

Standard arguments imply that any symmetric equilibrium must have $\max \operatorname{supp} F(p)$ be equal to the monopoly price, p^*.[22] Observe that $R(p^*) = p^* D(p^*) \rho_0 E(\pi_1 | \psi_0)$ is simply the monopoly revenue that can be extracted from a captive consumer who has no alternative offer. Setting $R(p_1) = R(p^*)$ (for the indifference necessary to induce mixing) yields an equation that implicitly defines a symmetric equilibrium distribution, $F(\cdot)$. Because the right-hand side of this equation is constant, inspection of (12) reveals that the equilibrium $F(\cdot)$ is unique and well-defined (*i.e.*, is increasing). For example, with $n = 2$ the equilibrium price distribution is

$$F(p) = \frac{D(p)p\left[E(\pi_1|\psi_0)\rho_0 + E(\pi_1|\psi_1)\rho_1\right] - D(p^*)p^* E(\pi_1|\psi_0)\rho_0}{D(p)pE(\pi_1|\psi_1)\rho_1}. \tag{13}$$

---

[21]Thus,

$$\rho_z = \binom{n-1}{z} \int_0^1 \cdots \int_0^1 \left\{ \prod_{j=2}^{z+1} [E(\mu)\pi_j] \prod_{j=z+2}^{n-1} [1 - E(\mu)\pi_j] \, g(\boldsymbol{\pi}) \right\} d\pi_n \ldots d\pi_1$$

This expression for ρ_z exploits the symmetry of g: We calculate the probability that firms $\{2, \ldots, z+1\}$ are in the consideration set and then multiply by the binomial coefficient $\binom{n-1}{z}$ to account for other permutations of consideration set containing z firms.

[22]$\operatorname{supp} F(\cdot)$ means the support of F. If $\max \operatorname{supp} F(p) < p^*$ then there must be some firm that has the highest price with probability 1 when setting $p_i = \max \operatorname{supp} F(p)$ (ties would require a shared mass point, which some firm would undercut). Because such a firm sells only to loyal consumers, it prefers to deviate to p^*. Similarly, if $\max \operatorname{supp} F(p) > p^*$, a firm could increase both its per-consumer profit and its demand by shifting mass from $\max \operatorname{supp} F(p)$ to p^*.

If the π_i's are independently distributed then (13) further simplifies to

$$F(p) = \frac{D(p)p - [1 - E(\mu)E(\pi)]p^*D(p^*)}{E(\mu)E(\pi)pD(p)}. \tag{14}$$

Characterisation of the equilibrium support is completed by setting $F(\underline{p}) = 0$ and solving for $\underline{p}$.

Because it is indifferent, firm 1's profit is the same as when it sets $p_1 = p^*$: $\Pi_1 = R(p^*)\mu_1 - C(\mu_1)$. Solving $\partial\Pi_1/\partial\mu_1 = 0$ and setting $\mu_1 = E(\mu)$ yields the symmetric value of μ necessary to complete the characterisation of a (putative) frictionless equilibrium. There might also be frictionless equilibria that are asymmetric. The following result establishes that, once firms are allowed to set $s_i > 0$, there can be no frictionless equilibrium, symmetric or otherwise.

Proposition 6. *There is no frictionless equilibrium of the game with $n > 1$ firms.*

Intuitively, frictionless price competition is such that firms' revenues are proportional to the mass of captive consumers (who are not willing or able to purchase elsewhere). But a firm can still attract these consumers even if it deviates to a small but positive s. Moderate increases in s therefore have only a secondary effect on a firm's revenues. On the other hand, increasing s significantly decreases firms' sales costs by excluding consumers who will buy from from rivals or who have a low interest in the deviator's product.

It should be emphasized that, while search frictions in my model ultimately relax price competition in the manner already studied elsewhere in the literature, the relaxation of competition is not the driving force behind the incentive to introduce them. Indeed, although the proof of Proposition 6 identifies a profitable deviation with higher prices, the deviating firm could have sold to exactly the same set of consumers at the same high price without increasing s above zero. The benefits of positive search costs come, instead, through the better allocation of sales resources. The unilateral incentive to raise costs identified here is, then, independent from, and complementary to the competition-reducing motive studied elsewhere in the literature.

Search costs and equilibrium browsing behavior

Proposition 6 rules-out frictionless equilibria. More intuition for the effect of oligopoly search costs can be obtained by looking at what happens when all firms set $s > 0$. In a similar spirit to Diamond (1971), it is then consistent with equilibrium for every firm to set the monopoly price, p^*. Indeed, suppose that a consumer finds that they have the opportunity to purchase from a firm they visit, which sells their desired good at the monopoly price. Such a consumer does not expect to gain from further search because any additional search is expected to also yield the monopoly price. On the other hand,
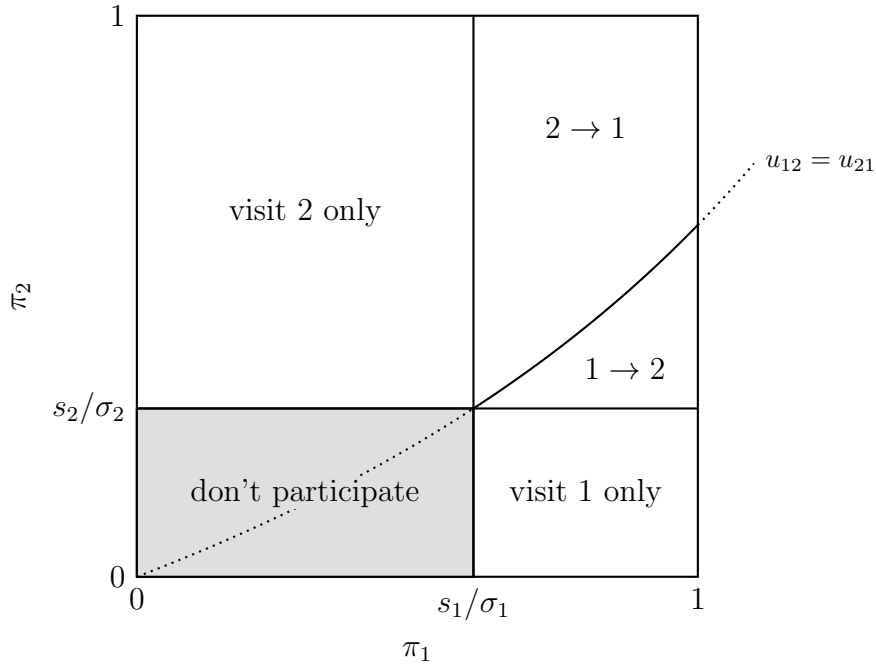
Figure 3: Browsing behaviour as a function of consumers' interest in firms 1 and 2.

consumers who are either uninterested in a product they view or are unable to buy from the firm that sells it may wish to search again.

If $n = 2$ the optimal consumer browsing behavior is as depicted in Figure 3 (in a slight abuse of notation, I write $i \to j$ to denote browsing i first, followed by j.). A consumer obtains expected utility $\pi_1 \sigma_1 - s_1$ from visiting 1 and is, therefore, willing to visit 1 only if $\pi_1 > s_1/\sigma_1$. A symmetric condition holds for 2. Consumers who are, in principle, prepared to visit both firms will start with the i that maximizes

$$u_{ij} \equiv \pi_i \sigma_i - s_i + [1 - \pi_i \mu_i][\pi_j \sigma_j - s_j] \tag{15}$$

and visit the second firm only if they do not buy from the first. The impact of competition is therefore to introduce a new margin for the effects of search costs: as well as excluding consumers with low interest (those with $\pi_1 < s_1/\sigma_1$), a positive s_i also drives some consumers to visit a rival first.

The firms' problem with respect to choosing μ is substantially unchanged—the equilibrium μ satisfies the analog of (3) (along with the equilibrium condition that consumers' expectations are correct). What is new here is that each firm computes $E(\pi|\text{visit})$ not via the simple threshold rule, $E(\pi|\pi > \bar{\pi})$, but rather by integrating over the relevant areas in Figure 3 (taking its rival's strategy as given).

If we impose some more structure on the model then it is possible to use the browsing behavior depicted in Figure 3 to study equilibrium search costs. To that end, I take interest to be uniformly distributed, costs to be quadratic, and demand to be linear.

Example 1. *Suppose that $n = 2$, $g(\pi_1, \pi_2) = 1$, $C(\mu) = \mu^2/2$, and $D(p) = (1 - p)$.*

1. *There is an equilibrium in which $s_1 = s_2 = 0.097$. In this equilibrium, consumers with $\pi_1, \pi_2 < 0.78$ are excluded from the market.*

2. *Compared to the symmetric frictionless outcome (derived above), the equilibrium yields lower consumer surplus but higher total welfare.*

3. *If firm 2 were to exit the market then firm 1's equilibrium search cost would increase to $s_1 = 0.1$. Consumers with $\pi_1 < 0.8$ would be then excluded.*

The calculations for Example 1 can be found in the supplementary appendix. In spite of the presence of competition, firms have a substantial equilibrium incentive to increase search costs—enough to deter 78% of potential visitors.

Much like Diamond (1971), introducing a positive (but potentially very small) search cost results in a discontinuous increase in price and therefore a discontinuous decrease in the surplus generated by each trade. Thus, the model is constructed in a fashion that would appear to strongly prejudice it against any finding that search costs might be welfare-improving. This makes part 2 of Example 1 all the more striking.

Part 3 of Example 1 says that a duopolist implements a lower search costs and excludes fewer consumers than does a monopolist. While this may not seem too surprising, the intuition for the result is interesting. As well as the obvious fact that firms are essentially competing to attract the most valuable consumers (those for whom both π_1 and π_2 are large), there is a second and more novel effect of competition. There are some consumers for whom π_1 is small and π_2 is large. As a monopolist, firm 1 would increase s to exclude such consumers and avoid wasting sales resources on them. However, under a duopoly these consumers will naturally tend to visit firm 2 first and, thanks to their high π_2, never make it to firm 1. Thus, firm 2 acts like a sponge that absorbs many of the consumers firm 1 would most wish to exclude, thereby reducing the need for screening.

We can also look at the welfare effects of increasing s from zero. The intuition underlying results like Proposition 2 continues to hold under oligopoly: the presence of many window shoppers will have the effect of reducing firms' investment incentives, to the detriment of welfare. An easy way to see this is to suppose that a mass ω of consumers are window shoppers who have $\boldsymbol{\pi} = \mathbf{0}$, while the remaining $1 - \omega$ have $\boldsymbol{\pi}$ atomlessly distributed according to $g(\boldsymbol{\pi})$. The following result then obtains:

Proposition 7. *For $n > 1$ and ω sufficiently large, there exists an $\epsilon > 0$ such that $s_1 = \ldots = s_n = \epsilon$ yields higher consumer surplus and total welfare than does the frictionless equilibrium derived above.*

The proof of Proposition 7 can be found in the supplementary appendix.

6 Summary and Final Remarks

A perennial problem faced by retailers is the time and resources wasted on trying to sell to consumers who are "just looking" or are otherwise unlikely to buy. In this article I have shown that a firm benefits from using search costs to screen its visitors—allowing it to target its sales efforts at more serious customers. This effect holds despite the fact that excluding browsers reduces a firm's potential demand. The reason is that costly search only deters marginal consumers who are the least likely to have an interest in the product in question. The firms that are likely to benefit most from screening through search costs are those that sell a product that is sales-intensive; those that attract many "window shoppers"; and those for which the most valuable customers are in a small minority.

Costly search imposes a direct cost upon consumers, but also carries efficiency benefits by facilitating the better use of sales or marketing resources and thereby compelling the firm to invest more in customer assistance. This effect can be large enough to offset the direct costs so that both welfare and consumer surplus can increase when firms are permitted to raise search costs. The same reasoning also implies that naturally occurring search costs (rather than synthetic costs created by the firm) can be welfare-improving. If search costs take the form of a transfer from consumers to the firm (rather than being pure waste) then welfare is always maximized by a strictly positive search cost. In the supplementary appendix, I allow for screening using prices as well as search costs. Nevertheless, firms rely solely on search costs for screening because, unlike a distortion in price, they do not reduce profits from inframarginal consumers. For a similar reason, consumers are better-off being screened by search costs than when price increases are used to achieve the same degree of screening.

The extant literature has shown that search costs can harm consumers through a relaxation of competition or the facilitation of surplus extraction. In the analysis above raising search costs can also lead to an increase in price but nevertheless leave consumers better-off. Thus, this article suggests that search costs can have positive effects that should be weighed against their harms. In particular, the strategic role played by search costs is important in judging their likely welfare consequences. In markets in which costly investments in sales are important, and where window shopping is pervasive, costly search is more likely to play an important and beneficial role in the deployment of sales resources. In markets where these conditions are not satisfied, regulators should view obfuscatory practices with greater suspicion as they are then more likely to be driven by perfidious strategic motives.

A Proofs

Proof of Lemma 1. Let $z(\mu) \equiv E(\pi | \pi > \overline{\pi}) D_2(p^*(\mu), \mu) p^*(\mu) - C'(m\mu)$. Applying the implicit function theorem to totally differentiate (3) yields

$$\frac{dz(\mu^*)}{d\mu^*} d\mu^* + D_2(p^*(\mu^*), \mu^*) p^*(\mu^*) \frac{\partial E(\pi | \pi > \overline{\pi})}{\partial \overline{\pi}} \frac{d\overline{\pi}}{ds} ds = 0$$

$$\implies \frac{d\mu^*}{ds} = -\frac{D_2[p^*(\mu^*), \mu^*] p^*(\mu^*) \frac{\partial E(\pi | \pi > \overline{\pi})}{\partial \overline{\pi}} \frac{d\overline{\pi}}{ds}}{\frac{dz(\mu^*)}{d\mu^*}}, \quad (16)$$

That the denominator is negative in the neighborhood of μ^* is implied by Assumption 1 (in particular, for $\mu = 0$ the left hand side of (3) is positive so it must be decreasing at its unique zero). Thus, $d\mu^*/ds$ has the same sign as $d\overline{\pi}/ds$ (firms invest more in sales when they attract more interested consumers). Assumption 2 then implies that $d\mu^*/ds \geq 0$ and $d\overline{\pi}/ds \geq 0$. To see that the latter inequality must be strict, suppose that $d\overline{\pi}/ds = 0$. By Assumption 1, $d\overline{\pi}/ds = 0$ implies that $d\mu^*/ds = dp^*/ds = 0$. However it must then be that $d\sigma/\partial s = 0$. We would then have

$$\frac{d\overline{\pi}}{ds} = \frac{\sigma - s\frac{d\sigma}{ds}}{\sigma^2} = 1/\sigma > 0, \quad (17)$$

a contradiction. That $dm/ds < 0$ follows immediately from $d\overline{\pi}/ds > 0$. Lastly, $d\overline{\pi}/ds > 0$ implies that $dR(s)/ds > 0$ must hold by a revealed preference argument because $R(s) = E(\pi | \pi > \overline{\pi}) p^* D(p^*, \mu^*)$. ∎

Proof of Proposition 1. The effect of an increase in s on profits is

$$\frac{d\Pi(s, \mu^*, p^*)}{ds} = \Pi_1(s, \mu^*, p^*) + \Pi_2(s, \mu^*, p^*) \frac{d\mu^*}{ds} + \Pi_3(s, \mu^*, p^*) \frac{dp^*}{ds}.$$

Noting that $\Pi_2(s, \mu^*, p^*) = \Pi_3(s, \mu^*, p^*) = 0$ by definition, this effect is

$$\frac{d\Pi}{ds} = \Pi_1(s, \mu^*, p^*) = \frac{d\overline{\pi}}{ds} \left\{ \frac{\partial E(\pi | \pi > \overline{\pi})}{\partial \overline{\pi}} rm + \frac{\partial m}{\partial \overline{\pi}} \left[E(\pi | \pi > \overline{\pi}) r - \mu^* C'(m\mu^*) \right] \right\}. \quad (18)$$

Writing out the terms more explicitly:

$$\Pi_1(s, \mu^*, p^*) = \frac{r g(\overline{\pi}) \int_{\overline{\pi}}^1 g(\pi)(\pi - \overline{\pi}) \, d\pi}{1 - G(\overline{\pi})} \frac{d\overline{\pi}}{ds} - \frac{r g(\overline{\pi}) \int_{\overline{\pi}}^1 g(\pi)\pi \, d\pi}{1 - G(\overline{\pi})} \frac{d\overline{\pi}}{ds} + g(\overline{\pi}) \frac{d\overline{\pi}}{ds} \mu^* C'(m\mu^*)$$

$$(19)$$

$$= g(\overline{\pi}) \frac{d\overline{\pi}}{ds} \left[\mu^* C'(\mu^* m) - \overline{\pi} r \right].$$

At $s = 0$ (*i.e.*, $\bar{\pi} = 0$), $\mu = \mu^*(0)$, $p = p^*(0)$, this is

$$\Pi_1(0, \mu^*(0), p^*(0)) = g(\bar{\pi}) \frac{d\bar{\pi}}{ds} \mu^* C'(\mu^*), \tag{20}$$

which has the same sign as $d\bar{\pi}/ds$. This cannot be negative because $s = 0 \implies \bar{\pi} = 0$. Suppose that $d\bar{\pi}/ds = 0$. By Assumption 1, $d\bar{\pi}/ds = 0$ implies that $d\mu^*/ds = dp^*/ds = 0$. However it must then be that $d\sigma^*(s)/ds = 0$. (17) would then hold—a contradiction. It must, therefore, be the case that $d\bar{\pi}/ds > 0$ and hence $\Pi_1[0, \mu^*(0), p^*(0)] > 0$. ∎

Proof of Proposition 2. Consumer surplus can be reformulated as

$$\mathrm{CS} = \int_{\underline{\pi}}^{1} (\sigma\pi - s) g(\pi) \, d\pi = \int_{\bar{\pi}}^{1} \sigma(\pi - \bar{\pi}) g(\pi) \, d\pi.$$

Because μ^* depends on s only via $\bar{\pi}$ (and p^* is pinned-down by μ^*) we can use Assumption 1 to write

$$\frac{dCS}{ds} = \frac{d\bar{\pi}}{ds} \frac{dCS}{d\bar{\pi}} = \frac{d\bar{\pi}}{ds} \left[\frac{d\sigma}{d\mu} \frac{d\mu}{d\bar{\pi}} \int_{\bar{\pi}}^{1} (\pi - \bar{\pi}) g(\pi) \, d\pi - \sigma \int_{\bar{\pi}}^{1} g(\pi) \, d\pi \right] \tag{21}$$
$$= \frac{d\bar{\pi}}{ds} \left[\frac{d\sigma}{d\mu} \frac{d\mu}{d\bar{\pi}} [E(\pi | \pi > \bar{\pi}) - \bar{\pi}] m - \sigma m \right].$$

From Lemma 1, we know that $d\bar{\pi}/ds > 0$ so

$$\frac{dCS}{ds} > 0 \iff \frac{d\sigma}{d\mu} \frac{d\mu}{d\bar{\pi}} [E(\pi | \pi > \bar{\pi}) - \bar{\pi}] - \sigma > 0. \tag{22}$$

An expression for $d\mu^*/d\bar{\pi}$ is found along similar lines to (16):

$$\frac{dz(\mu^*)}{d\mu^*} d\mu^* + D_2(p^*(\mu^*), \mu^*) p^*(\mu^*) \frac{\partial E(\pi | \pi > \bar{\pi})}{\partial \bar{\pi}} d\bar{\pi} = 0$$
$$\implies \frac{d\mu^*}{d\bar{\pi}} = - \frac{D_2(p^*(\mu^*), \mu^*) p^*(\mu^*) \frac{\partial E(\pi | \pi > \bar{\pi})}{\partial \bar{\pi}}}{z'(\mu^*)}. \tag{23}$$

Evaluating this expression yields

$$\frac{d\mu^*}{d\bar{\pi}} = - \frac{D_2(p^*, \mu^*) p^* [E(\pi | \pi > \bar{\pi}) - \bar{\pi}]}{m z'(\mu^*)} g(\bar{\pi})$$

Combining with (22) and letting $s \to 0$ (which implies $\bar{\pi} \to 0$, $m \to 1$) reveals $dCS/ds|_{s=0} > 0$ if

$$- \frac{d\sigma^*(0)}{d\mu} \frac{D_2(p^*, \mu_0) p^* E(\pi)^2}{z'(\mu^*) \sigma^*(0)} > \frac{1}{g(0)}, \tag{24}$$

where $\mu_0 \equiv \mu^*(0)$. By Assumption 1 and Assumption 2, the left-hand side is strictly positive (in particular we saw in the proof of Lemma 1 that $z'(\mu^*) < 0$) and the right-hand

side goes to zero as $g(0) \to \infty$. ∎

Proof of Corollary 2. Condition (8) is found from (24) in three steps. First, evaluate $z'(\mu^*)$. Second, let $D_2(p, \mu) = h'(\mu)D(p)$, $D_{22}(p, \mu) = h''(\mu)D(p)$, $dp^*/d\mu = 0$, $(d\sigma/d\mu)/\sigma = h'(\mu)/h(\mu)$. Thirdly, write $C'(\mu_0) = E(\pi)h'(\mu_0)D(p^*)p^*$ (from (3)).

Remark on obtaining a condition for $dW/ds|_{s=0} > 0$: From (19) we can write $d\Pi/ds = (d\Pi/d\overline{\pi})(d\overline{\pi}/ds)$. Combining with (21) yields

$$\frac{dW}{ds} = \frac{d\overline{\pi}}{ds}\left[\frac{dCS}{d\overline{\pi}} + \frac{d\Pi}{d\overline{\pi}}\right].$$

Using (20) we therefore obtain a condition for $dW/ds|_{s=0} > 0$ by adding a new term, $C'(\mu_0)\,\mu_0/\sigma^*(0)$ to the left-hand side of (24). This new term is positive so the condition for $dW/ds|_{s=0} > 0$ is easier to satisfy than $dCS/ds|_{s=0} > 0$. ∎

Proof of Proposition 3. Writing $\sigma_s \equiv \sigma(p^*, \mu^*(s))$, the equilibrium s yields higher consumer surplus than does $s = 0$ if

$$\int_{\overline{\pi}}^1 [\sigma_{s^*}\pi - s^*]\,g(\pi)\,d\pi - \int_0^1 \sigma_0\pi g(\pi)\,d\pi = \int_{\overline{\pi}}^1 \sigma_{s^*}(\pi - \overline{\pi})g(\pi)\,d\pi - \int_0^1 \sigma_0\pi g(\pi)\,d\pi$$

$$= [1 - G(\overline{\pi})]\left[E(\pi|\pi > \overline{\pi}) - \overline{\pi}\right]\sigma_{s^*} - E(\pi)\sigma_0$$

is positive. From Corollary 1 we know that $\overline{\pi} = a/(2-a)$, implying $E(\pi|\pi > \overline{\pi}) = 1/(2-a)$. Thus, $CS|_{s=s^*} - CS|_{s=0} > 0$ if

$$\frac{\sigma_{s^*}}{\sigma_0} = \frac{[\mu^*(s^*)]^a}{[\mu^*(0)]^a} > \frac{(2-a)^2}{4(1-a)^2}. \tag{25}$$

$\mu^*(s^*)$ is given in Corollary 1, while $\mu^*(0)$ is given by (3) with $E(\pi|\pi > \overline{\pi}) = E(\pi) = 1/2$ and $m = 1$: $[a\mu^{a-1}/8] - k\mu^{k-1} = 0 \implies \mu^*(0) = [8k/a]^{\frac{1}{a-k}}$. Substituting $\mu^*(s^*)$ and $\mu^*(0)$ into (25) and simplifying yields (10). ∎

Proof of Proposition 4. At equilibrium we know (see (4)) that $D(p^*, \mu^*)p^*\overline{\pi}$ must intersect $\mu^* E(\pi|\pi > \overline{\pi})D_2(p^*, \mu^*)p^*$ from below when viewed as a function of $\overline{\pi}$:

$$\frac{d\mu^*}{d\overline{\pi}}\left[E(\pi|\pi > \overline{\pi})h'(\mu^*)D(p^*)p^* + h''(\mu^*)D(p^*)p^*\mu^* E(\pi|\pi > \overline{\pi})\right] +$$

$$\frac{\partial E(\pi|\pi > \overline{\pi})}{\partial \overline{\pi}}\mu^* h'(\mu^*)D(p^*)p^* < \frac{d\mu^*}{d\overline{\pi}}h'(\mu^*)\overline{\pi}D(p^*)p^* + h(\mu^*)D(p^*)p^*.$$

Cancelling the $D(p^*)p^*$ terms and rearranging:

$$\frac{d\mu^*}{d\overline{\pi}}h'(\mu^*)\left[E(\pi|\pi > \overline{\pi}) - \overline{\pi}\right] - h(\mu^*)$$

$$< -\frac{d\mu^*}{d\overline{\pi}}h''(\mu^*)\mu^* E(\pi|\pi > \overline{\pi}) - \frac{\partial E(\pi|\pi > \overline{\pi})}{\partial\overline{\pi}}\mu^* h'(\mu^*). \quad (26)$$

From (22), the left-hand side of (26) must be positive if consumer surplus is to be increasing in s at equilibrium. Thus, the right-hand side must also be positive. Using (23), the requirement that the right-hand side of (26) is positive becomes

$$-\frac{h'(\mu^*)D(p^*)p^*\frac{\partial E(\pi|\pi>\overline{\pi})}{\partial\overline{\pi}}}{mC''(\mu^* m) - h''(\mu^*)D(p^*)p^* E(\pi|\pi > \overline{\pi})}h''(\mu^*)\mu^* E(\pi|\pi > \overline{\pi})$$

$$-\frac{\partial E(\pi|\pi > \overline{\pi})}{\partial\overline{\pi}}\mu^* h'(\mu^*) > 0.$$

This condition fails when $C''(\cdot) = 0$ (cancelling terms then yields $0 > 0$), and the left-hand side is decreasing in $C''(\cdot)$.

The result for total welfare follows from the fact that $W = \text{CS} + \Pi$ and Π is, by definition, stationary at equilibrium. ∎

Proof of Lemma 2. The firm's problem is to choose the menu that maximizes profits

$$\max_{\mu(\pi),s(\pi)}\int_0^1 \left[\pi\mu(\pi)D(p^*)p^* - C(\mu(\pi))\right]g(\pi)\,d\pi$$

subject to consumers' incentive compatibility and participation constraints

$$\pi\mu(\pi)\int_{p^*}^\infty D(p)\,dp - s(\pi) \geq \max\left\{0, \pi\mu(\widehat{\pi})\int_{p^*}^\infty D(p)\,dp - s(\widehat{\pi})\right\} \quad \forall\widehat{\pi}.$$

Begin by finding the $\mu(\pi)$ to maximize profits on a consumer-by-consumer basis. Profit from a type-π consumer is $\pi\mu D(p^*)p^* - C(\mu)$, which is maximized by $\mu(\pi) = \kappa^{-1}[\pi D(p^*)p^*]$. Substituting this into consumers' utility function yields

$$u = \pi\kappa^{-1}[\widehat{\pi}D(p^*)p^*]\int_{p^*}^\infty D(p)\,dp - s(\widehat{\pi})$$

$$\frac{du}{d\widehat{\pi}} = \frac{\pi D(p^*)p^*}{C''\left[\kappa^{-1}(\widehat{\pi}D(p^*)p^*)\right]}\int_{p^*}^\infty D(p)\,dp - s'(\widehat{\pi}).$$

In equilibrium, utility must be maximized when consumers report honestly: $\widehat{\pi} = \pi$. The corresponding first-order condition is

$$\frac{\pi D(p^*)p^*}{C''\left(\kappa^{-1}[\pi D(p^*)p^*]\right)}\int_{p^*}^\infty D(p)\,dp - s'(\pi) = 0,$$

which implies

$$s(\pi) = s(0) + \int_0^\pi \frac{zD(p^*)p^*}{C''\left[\kappa^{-1}(zD(p^*)p^*)\right]} \, dz \int_{p^*}^\infty D(p) \, dp.$$

Assuming that s can never be negative, the relevant boundary condition is such that a $\pi = 0$ consumer's participation constraint is satisfied: $s(0) = 0$. ∎

Proof of Corollary 3. If s is constrained to be zero then the firm's profit is $E(\pi)D(p^*)p^*\mu - C(\mu)$, implying $\mu^* = \kappa^{-1}[E(\pi)D(p^*)p^*]$. The effect on consumer surplus of allowing the firm to offer its optimal menu of search costs is therefore

$$\Delta\text{CS} = \int_0^1 \left[\pi \left\{ \kappa^{-1}[\pi D(p^*)p^*] - \kappa^{-1}[E(\pi)D(p^*)p^*] \right\} \int_{p^*}^\infty D(p) \, dp - s(\pi) \right] g(\pi) \, d\pi.$$

Using $C(\mu) = \mu^2$ and $G(\pi) = \pi^\beta$, we find that

$$\Delta\text{CS} \geq 0 \iff (1 - 2\beta - \beta^2) \underbrace{\frac{\beta D(p^*)p^*}{4(1+\beta)^2(2+\beta)} \int_{p^*}^\infty D(p) \, dp}_{>0} \geq 0.$$

The result is immediate. ∎

Proof of Proposition 5. To prove part 1: Welfare is given by

$$W = \underbrace{\int_{\underline{\pi}}^1 g(\pi)\pi\sigma(p^*(\mu^*), \mu^*) \, d\pi - ms}_{\text{CS}} + \underbrace{ms + E(\pi|\pi > \overline{\pi})mD(p^*, \mu^*)p - C(m\mu^*)}_{\text{Profit}}$$

The two ms terms cancel. When $s = 0$ we have

$$\frac{d}{ds}\left(\int_{\underline{\pi}}^1 g(\pi)\pi\sigma(p^*(\mu^*), \mu^*) \, d\pi \right) = \int_0^1 g(\pi)\pi \frac{d\sigma(p^*(\mu^*), \mu^*)}{d\mu^*} \frac{\partial\mu^*}{\partial s} \, d\pi,$$

which is positive by Assumption 2 and Lemma 1. That

$$\frac{d}{ds}\left(E(\pi|\pi > \overline{\pi})mD(p^*, \mu^*)p - C(m\mu^*) \right)$$

is positive when $s = 0$ was established in Proposition 1.

The proof of part 2 is identical to that for Proposition 2. ∎

Proof of Proposition 6. Suppose that there exists a frictionless equilibrium. Write $F_i(\cdot)$ for the equilibrium price distribution used by firm i, $\overline{p}_i \equiv \max \text{supp } F_i(p)$, and $\underline{p}_i \equiv \min \text{supp } F_i(p)$. Let $\mu_i(p)$ denote that value of μ chosen by firm i when it sets $p_i = p \in \text{supp } F_i$.

Begin by establishing that $\max_i \bar{p}_i = p^*$. It cannot be the case that two firms share a mass point at some p because there would then be a tie with positive probability and one of the firms would wish to slightly undercut the mass point. Suppose $\max_i \bar{p}_i < p^*$. There then exists a j with $\bar{p}_j = \max_i \bar{p}_i$ such that j sells only to loyal consumers (viz. consumers who are either not interested or not able to buy from any other firm). But j could sell to its loyal consumers at p^* and would therefore prefer to increase its price. Likewise, if $\max_i \bar{p}_i > p^*$ then a firm setting $p_j = \max_i \bar{p}_i$ could (strictly) increase its per-consumer profit and (weakly) increase its demand by shifting mass from $\max_i \bar{p}_i$ to p^*. Summarizing so far: we know that there exists at least one firm that sometimes sets $p_i = p^*$, and that this firm attracts only loyal consumers whenever it does so. Call this firm 1. Its profits in the putative equilibrium are

$$p^* D(p^*) E(\pi_1 | \psi_0) \rho_0 \mu_1(p^*) - C(\mu_1(p^*)). \tag{27}$$

Now, suppose that firm 1 deviates to s_1 positive but small. If

$$\pi_1 \sigma_1^e - s_1 \equiv \pi_1 \mu_1^e \int_{p_1^e}^{\infty} D(p)\,dp - s_1 \geq \int_{\underline{p}_j}^{\infty} D(p)\,dp$$

(where an e superscript denotes consumers' expectations following the deviation) then a type π_1 consumer would wish to visit 1 even after the best possible realisation of demand and price at firm j. It could then conceivably be optimal for some consumers to visit 1 before j, despite the fact that a visit to the latter is costless. However, note that this inequality requires $p_1^e < \underline{p}_j$. The belief $p_1^e < \min_{j \neq 1} \underline{p}_j$ can never be consistent with profit-maximizing behavior because 1 could increase its price slightly and still have a lower price than every rival with probability 1. Thus, it must be that consumers first visit a firm other than 1, and that some of these consumers (in expectation) never reach firm 1. More formally, if the mass of consumers visiting firm 1 is m_1 then we must have $m_1 < 1$.

If firm 1 pairs its deviation in s_1 with the actions $p_1 = p^*$ and $\mu_1 = \mu_1(p^*)$ (*i.e.*, it plays the μ it would have played in the putative frictionless equilibrium given price p^*) then its profits will be $p^* D(p^*) E(\pi_1 | \psi_0) \rho_0 \mu_1(p^*) - C(m_1 \mu_1(p^*))$, which is strictly greater than (27). The deviation to $s_1 > 0$ is therefore profitable. ∎

References

Armstrong, M. and Zhou, J. "Paying for Prominence." *Economic Journal*, Vol. 121 (2011), pp. F368–F395.

Armstrong, M. and Zhou, J. "Search Deterrence." *Review of Economic Studies*, Vol. 83 (2016), pp. 26–57.

Athey, S. and Gans, J. S. "The Impact of Targeting Technology on Advertising Markets and Media Competition." *American Economic Review*, Vol. 100 (2010), pp. 608–613.

Bar-Isaac, H., Caruana, G., and Cuñat, V. "Information Gathering and Marketing." *Journal of Economics and Management Strategy*, Vol. 19 (2010), pp. 375–401.

Bergemann, D. and Bonatti, A. "Targeting in Advertising Markets: Implications for Offline vs. Online Media." *RAND Journal of Economics*, Vol. 42 (2011), pp. 417–443.

Carlin, B. I. and Ederer, F. "Search Fatigue." Working Paper No. 17895, National Bureau of Economic Research, 2012.

Chioveanu, I. and Zhou, J. "Price Competition with Consumer Confusion." *Management Science*, Vol. 59 (2013), pp. 2450–2469.

Diamond, P. "A Model of Price Adjustment." *Journal of Economic Theory*, Vol. 3 (1971), pp. 156–168.

Ellison, G. "A Model of Add-On Pricing." *Quarterly Journal of Economics*, Vol. 120 (2005), pp. 585–637.

Ellison, G. and Ellison, S. F. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica*, Vol. 77 (2009), pp. 427–452.

Ellison, G. and Wolitzky, A. "A Search Cost Model of Obfuscation." *RAND Journal of Economics*, Vol. 43 (2012), pp. 417–441.

Gu, Y. and Wenzel, T. "Strategic Obfuscation and Consumer Protection Policy." *Journal of Industrial Economics*, Vol. LXII (2014), pp. 632–660.

Hämäläinen, S. "Obfuscation by Substitutes, Drowning by Numbers." Working Paper, University of Helsinki, 2016.

Janssen, M. C., Moraga-González, J. L., and Wildenbeest, M. R. "Truly Costly Sequential Search and Oligopolistic Pricing." *International Journal of Industrial Organization*, Vol. 23 (2005), pp. 451–466.

Johnson, J. P. "Targeted Advertising and Advertising Avoidance." *RAND Journal of Economics*, Vol. 44 (2013), pp. 128–144.

Moraga-González, J. L., Sándor, Z., and Wildenbeest, M. R. "Prices and Heterogeneous Search Costs." *RAND Journal of Economics* (forthcoming).

Mussa, M. and Rosen, S. "Monopoly and Product Quality." *Journal of Economic Theory*, Vol. 18 (1978), pp. 301–317.

Petrikaite, V. "Consumer Obfuscation by a Multiproduct Firm." Working Paper, Barcelona Graduate School of Economics, 2015.

Piccione, M. and Spiegler, R. "Price Competition Under Limited Comparability." *Quarterly Journal of Economics*, Vol. 127 (2012), pp. 97–135.

Piccolo, S., Tedeschi, P., and Ursino, G. "How Limiting Deceptive Practices Harms Consumers." *RAND Journal of Economics*, Vol. 46 (2015), pp. 611–624.

Reinganum, J. F. "A Simple Model of Equilibrium Price Dispersion." *Journal of Political Economy*, Vol. 87 (1979), pp. 851–858.

Rhodes, A. and Wilson, C. M. "False Advertising." Working Paper Number 72693, Munich Personal RePEc Archive, 2016.

Salop, S. "The Noisy Monopolist: Imperfect Information, Price Dispersion and Price Discrimination." *Review of Economic Studies*, Vol. 44 (1977), pp. 393–406.

Shin, J. "The Role of Selling Costs in Signaling Price Image." *Journal of Marketing Research*, Vol. 42 (2005), pp. 302–312.

Stahl, D. O. "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review*, Vol. 79 (1989), pp. 700–712.

Varian, H. R. "A Model of Sales." *American Economic Review*, Vol. 70 (1980), pp. 651–659.

Wilson, C. M. "Ordered search and equilibrium obfuscation." *International Journal of Industrial Organization*, Vol. 28 (2010), pp. 496–506.

Zhou, J. "Multiproduct Search and the Joint Search Effect." *American Economic Review*, Vol. 104 (2014), pp. 2918–2939.