

A search cost model of obfuscation

Glenn Ellison*

and

Alexander Wolitzky**

This article develops models in which obfuscation is individually rational for oligopolistic firms. Firms sell a homogeneous good to rational consumers who incur search costs to learn prices. Search costs are endogenized by allowing obfuscation—firms have an unobservable action that increases the time needed to learn their price. One model involves search costs convex in shopping time. We show that slight convexity can dramatically alter the equilibrium price distribution. A second model examines an informational linkage between current and future search costs: consumers are uncertain about a component of search costs. Here, a signal-jamming mechanism can lead to equilibrium obfuscation.

1. Introduction

■ Anyone who has shopped for a mattress, tried to compare the full sets of fees charged by multiple banks or mortgage lenders, or gotten quotes from contractors for a home renovation will find it easy to question the universality of the classic economic argument that firms will disclose all relevant information.¹ Ellison and Ellison (2009) describe practices in which firms intentionally make shopping complicated, difficult, or confusing as “obfuscation” and provide empirical evidence from online shopping. It is easy to think of reasons why it would be collectively rational for firms to practice obfuscation: equilibrium prices are increasing in consumer search costs in many search models, and price discrimination arguments can also be given.² Arguments based on collective rationality, however, bring up a natural critique: why collude on obfuscation rather than just colluding directly on price? In this article, we discuss a search-based model in which it is individually rational for firms to raise consumer search costs.

Diamond (1971) first formalized the connection between search costs and price levels, noting that even an ε search cost could increase prices from the competitive level to the monopoly level because consumers will have no incentive to search if they expect all firms to charge monopoly

*MIT, Microsoft Research, and NBER; gellison@mit.edu.

**Stanford University and Microsoft Research; alexander.wolitzky@gmail.com.

This article was supported by NSF grant SES-0550897 and the Toulouse Network for Information Technology. The second author also received support through a NSF Graduate Research Fellowship.

¹ See Grossman (1981) and Milgrom (1981).

² Diamond (1971) and many subsequent articles connect search costs and equilibrium price levels. Ellison (2005) shows that the joint adoption of add-on pricing strategies can increase prices in a competitive price discrimination model.

prices. Several subsequent articles developed two other important insights: there is a more natural search problem when price dispersion is present, and price dispersion will exist in equilibrium when consumers are differentially informed.³ Our model closely follows that of Stahl (1989), who considers a continuum of consumers shopping for a homogenous good offered by N firms. A fraction μ of the consumers have no search costs and learn all firms' prices. The other $1 - \mu$ pay a search cost of s every time they obtain a price quote. Consumers have identical downward sloping demands $D(p)$. Stahl shows that this produces an elegant, tractable model. All consumers with positive search costs search exactly once. Firms choose prices from a nonatomic distribution on an interval $[\underline{p}, \bar{p}]$ following mixed strategies like those in Varian (1980) and Rosenthal (1980). The model's comparative statics clearly bring out the collective incentive to increase search costs: prices and firm profits increase as the search cost s increases.

Section 2 of our article introduces our model and derives some preliminary results that are common to the different versions we eventually consider. We model obfuscation in a very simple way: consumers are assumed to have a disutility that depends on the total time spent shopping, and each firm is allowed to choose the length of time that is required to learn its price. This time is not observable to consumers until after they have visited the firm. The results here are primarily that a number of the basic features of Stahl's model carry over to our environment: in equilibrium, firms make positive profits and choose prices from a nonatomic distribution with support $[\underline{p}, \bar{p}]$; and consumers search until the expected gain from taking another draw from the price distribution exceeds the expected search costs.

The "obfuscation" in our model is intended to potentially capture a number of real-world phenomena. In the online shopping application, for example, the firm may be choosing both the number of screens that a consumer must click through before she learns the final price, including recommended upgrades, shipping costs, taxes, service fees, etc., as well as the time that it takes each screen to load. In face-to-face retail, the firm may be choosing to tell its salespeople for how long they should talk to customers for before quoting them the final price. In other applications, the firm will not choose waiting times directly, but may instead choose how clearly to convey information about prices, which then maps into waiting time. For example, in the bank application, the firm may be choosing the complexity of its fee structure, which determines how long it would take a consumer to read through the full list of fees for overdrafts, low balances, ATM use, wire transfers, etc. and estimate what he or she will end up paying each month. The time cost of learning the firm's price can also be interpreted as the time required to learn the product's quality and thereby learn a quality-adjusted price. For example, in the case of mattress shopping the price of each mattress at a store, e.g., "Sealy Posturepedic Ruby," may be readily observable, but time will be required to inquire about product attributes and learn which name corresponds to the mattress the consumer had seen at another store and/or to make mental adjustments to account for differences in the attributes of different stores' offerings.⁴

Section 3 analyzes our simplest model of obfuscation. On the firm side, obfuscation is assumed to be costless. On the consumer side, we assume that consumers have a strictly convex disutility $g(t)$ for the time t they spend shopping. We view this as a small departure from the traditional assumption in a realistic direction. We think it is realistic because disutility would be convex in a standard time-allocation model with decreasing returns to leisure.⁵ It is small because for all of our main results $g'(t)$ need only be ε greater than $g'(0)$ even in the $t \rightarrow \infty$ limit, in which case no amount of obfuscation could ever have more than an ε effect on consumers' future

³ Classic articles include Butters (1977), Salop and Stiglitz (1977), Varian (1980), Rosenthal (1980), and Burdett and Judd (1983). See Baye and Morgan (2001) for a recent model applied to online markets and Baye, Morgan, and Scholten (2006) for a nice survey including recent developments.

⁴ See Hendricks, Sorensen, and Wiseman (forthcoming) for a model emphasizing that it may be costly for consumers both to learn product attributes and their own preferences.

⁵ Another interpretation of the convexity of $g(\cdot)$, suggested to us by a referee, is that consumers fall victim to the sunk-cost fallacy, and therefore wish to avoid paying a large total search cost even if they have already paid a portion of it.

search costs. Yet, it is a departure that can greatly alter the equilibrium set. Holding obfuscation levels fixed, our model is much like Stahl's (1989)—the firms' pricing strategies will coincide with those of Stahl's model, with the search cost parameter set equal to the incremental cost of a second search. The fact that obfuscation is endogenous, however, dramatically affects what is possible in equilibrium. Specifically, equilibria in which the upper bound of the price distribution is strictly less than the monopoly price become impossible because a firm can simultaneously make small deviations in two dimensions: increase its price to slightly above \bar{p} and also slightly increase its obfuscation level. Hence, in all equilibria of our model the upper bound of the price distribution is the monopoly price.⁶ Such an upper bound on the equilibrium distribution is only possible if equilibrium search costs are above some lower bound. Therefore, there is a lower bound on the level of equilibrium obfuscation. The lower bound can be zero, but can also be substantial. Obfuscation hurts consumers in two ways: consumers incur higher search costs and pay higher prices. We also derive tractable comparative statics results. For example, in equilibrium obfuscation adjusts so as to offset changes in the exogenous component of consumer search costs.⁷

Section 4 introduces the possibility that obfuscating may be costly for firms. This makes obfuscation levels more deterministic, because in equilibrium each firm must be choosing the minimum level of obfuscation consistent with the equilibrium level of consumer search. It also allows us to discuss cross section relationships between prices and obfuscation. For example, with costly obfuscation firms with the lowest markups will not obfuscate at all, whereas firms with the highest markups do the most obfuscation. We also note that the combination of convex search costs and costly obfuscation can produce models with more complex patterns of search and obfuscation in which some costly searchers visit multiple stores and obfuscation strategies are nonmonotone in price.

Section 5 considers an alternate model of obfuscation. We return to the traditional assumption that consumers have a linear disutility of search time, dropping the strict convexity assumption used in Sections 3 and 4. Instead, we depart from Stahl's model in another direction we find realistic: we assume that there is common uncertainty about how much time is required to learn a firm's price in the absence of obfuscation. A key feature of this model is that consumers' expectations about future search costs increase in the amount of time it takes them to learn the price of the first firm they visit. For example, one could think of this as a model in which consumers are not born knowing how long it takes to get a price quote from a home improvement contractor and in which consumers who spend a long time discussing a project with the first contractor they contact assume that the process of getting a bid from another contractor will also be time-consuming. A natural consequence of such an effect is that obfuscation can occur for signal-jamming reasons. Some predictions of the signal-jamming model are similar to the convex costs model: firms have a strict incentive to obfuscate, and the equilibrium price distributions are a selection from the set of equilibrium distributions of Stahl's model. But the mechanism behind the obfuscation is different and this leads to some interesting differences in predictions. One is an "excess obfuscation problem": obfuscation is almost always above what is necessary to deter search. This leads to lower prices, making firms worse off, but also makes consumers worse off due to high obfuscation. Another is that the selection among the equilibria of Stahl's model is different, largely due to the excess obfuscation problem.

Our article is related to a number of others. Ellison and Ellison (2009) provide informal descriptive evidence on obfuscation among a group of e-retailers and present empirical evidence that suggests that at least two mechanisms are involved: consumers appear to have substantially incomplete knowledge of prices, and firms' add-on pricing strategies appear to create an adverse-selection effect that would be expected to increase equilibrium markups. A number of subsequent

⁶ This "slight upward deviation" argument is somewhat reminiscent of the Diamond paradox, but the idea of a "double deviation" in price and obfuscation is new. Arguments of this kind do not eliminate equilibria with prices bounded away from the monopoly price in standard search models like Stahl (1989).

⁷ In this model, and in all of the variants we consider, a monopoly seller would have no incentive to engage in obfuscation. Thus, obfuscation is a consequence of competition.

articles have explored obfuscation mechanisms.⁸ Ellison (2005) discusses add-on pricing in the context of a competitive price discrimination model. He notes that add-on pricing is not individually rational in the base model, but could be made individually rational by adding a subpopulation of irrational consumers who were exploited by the add-on strategy. Gabaix and Laibson (2006) work out an explicit model along these lines.

Spiegler (2006) provides an alternate boundedly-rational approach. In his model, consumers are only capable of evaluating products on one of many dimensions. Firms “obfuscate” by randomizing and making the product more attractive on some dimensions (e.g., making fees lower if some contingency arises) and less attractive on others. He notes that an increase in the competitiveness of the market (more firms) leads to an increase in obfuscation but no change in average prices. This is somewhat similar in spirit to our finding that decreases in exogenous search costs don’t change average prices because they are fully offset by a change in obfuscation. Other comparative statics differ, however, and the meaning of obfuscation and the mechanisms are completely different. Eliaz and Spiegler (2011) address some related topics, e.g., whether firms with higher prices do more or less to inform consumers, in another elegant model with boundedly rational consumers. Their model is more similar to the traditional information revelation literature than to our article in that informing consumers is the costly action.

Carlin (2009) and Wilson (2010) are most closely related to our article. Each also models obfuscation as a strategic decision by firms that increases search costs in a model with optimal consumer search. Carlin’s model differs from ours both in the focus and in the type of search model it uses. The search model is an all-or-nothing model along the lines of Salop and Stiglitz (1977) and Varian (1980).⁹ More importantly, Carlin’s focus is primarily on how obfuscation affects market prices, whereas some of our main motivations are to explore why it is individually (as opposed to collectively) rational to obfuscate and how obfuscation varies in the cross section. Carlin does make obfuscation individually rational and not just collectively rational, but this is done in a fairly straightforward manner so that the article can focus on other things: consumers observe a summary statistic (like the average obfuscation level) before deciding whether to conduct an all-or-nothing search and do not observe any individual firm’s obfuscation level, so an increase in obfuscation by any one firm leads to exactly the same outcome as would a smaller coordinated increase by all firms.¹⁰ In particular, obfuscation by a firm in Carlin’s model affects the search incentives of the entire pool of consumers, while in our model it only affects those consumers that visit the firm; in this regard, Carlin’s model is more similar to Robert and Stahl (1993), who model advertising as informing a fraction of the population about a firm’s price.

Wilson (2010) does focus on the question of why obfuscation is individually rational and develops a very nice argument (which is also very different from ours). The primary difference between Wilson’s model and ours is that Wilson assumes that the firm-specific search costs are observable to consumers when they choose which stores to visit. One’s first thought might be that this will make obfuscation impossible, because consumers will always choose to visit firms with the lowest search costs first. Wilson’s clever observation is that while it is true that many or all consumers will visit the low-search-cost firm first, this does not necessarily render obfuscation unappealing. Obfuscation can provide strategic-commitment benefits: by making itself less attractive to the consumers with positive search costs, the obfuscating firm induces its rival to focus more on these consumers and raise prices, which can benefit both firms. Our article differs from his in the assumptions on observability, in the mechanisms that drive obfuscation, and in the details of many results. For example, in his article obfuscating firms tend to charge lower prices, whereas obfuscation is associated with charging high prices in our model.

⁸ Ellison (2006) includes a survey of some of this literature.

⁹ Baye, Morgan, and Scholten (2006) refer to these as clearinghouse models.

¹⁰ White (2009) studies the incentives of a search engine to improve search quality. This also is similar in some ways to the coordinated increase problem.

2. Model and preliminary results

■ In this section we present our model and derive some basic results. Our model is similar to that of Stahl (1989) with two additions: search costs are allowed to be a nonlinear function of the number of searches carried out, and the per-search cost is an endogenous choice of the firms. The results in this section show that some standard results carry over: consumer search strategies can be characterized using cutoff rules, firms earn positive profits in a dispersed price equilibrium, and equilibrium price distributions are atomless.

□ **Model.** Suppose N firms indexed by $i = 1, 2, \dots, N$ are selling undifferentiated goods to a unit mass of consumers. We suppose consumers are of two types: proportion μ are “costless searchers” who automatically learn all firms’ prices and proportion $1 - \mu$ are “costly searchers” who must incur search cost $g(t)$ to spend a total time t searching. Ascertaining firm i ’s price requires time $\tau + t_i$, where $\tau > 0$ is exogenous and t_i is the obfuscation level chosen by firm i . For example, τ might be the amount of time it takes for unavoidable tasks like driving to a store or loading a webpage and reading product descriptions to verify that the item is what the consumer wants, whereas t_i might be the amount of additional time required to find the product given where the firm has placed it or the time it takes to find information that the firm has not posted as prominently, e.g., product attributes, over-credit-limit fees, shipping charges, etc.^{11,12} Therefore, a costly searcher would incur cost $g(\tau + t_1)$ to learn the price of a firm that chooses obfuscation level t_1 if this is her first search and would incur total cost $g(2\tau + t_1 + t_2)$ if she chose to continue her search and also learn the price of a second firm that chooses obfuscation level t_2 . We assume that $g(0) = 0$ and that $g(\cdot)$ is twice continuously differentiable, strictly increasing, and weakly convex. Consumers cannot observe a firm’s obfuscation level before they visit it and learn its price, but do have rational expectations about the distribution of obfuscation levels.

As in Stahl (1989), consumers have downward-sloping demand functions $D(p)$ that satisfy $\int_0^\infty D(x) dx < \infty$ and $R(p) \equiv pD(p)$ is the revenue a firm obtains from selling to consumer with demand $D(p)$ at price p . We assume that $R(p)$ is continuously differentiable with unique maximum p^m and that $R'(p) > 0$ if $p < p^m$. Each firm i out of $N \geq 2$ firms chooses price p_i and obfuscation level t_i . Firms produce at zero marginal cost. However, firm i incurs a fixed obfuscation cost of $c(t_i)$ when it chooses obfuscation level t_i ; we assume that the obfuscation cost function is differentiable with $c(0) = 0$ and $c'(t) \geq 0$ for all t . In some sections we will focus on the case of costless obfuscation, $c(t) = 0$ for all t , which allows for the simplest results.

Observe that if consumers have unit demand up to a choke price p^m then all of our assumption are satisfied with the exception that $R(p)$ would be continuously differentiable everywhere except for p^m (rather than being continuously differentiable everywhere). All of our results and proofs would go through under this alternative assumption with the exception of the statement in Proposition 7 that when obfuscation is costly firms will never use enough obfuscation to drive the highest prices all the way up to p^m . In addition, the second part of Proposition 1, while still true, may only hold vacuously. We discuss these issues more when presenting the propositions in question.

The game proceeds as follows. First, firms simultaneously and noncooperatively choose obfuscation levels and prices. Then, costless searchers automatically learn all firms’ prices and can buy from any firm, while costly searchers search strategically: they draw a new randomly selected firm with each search and may stop searching and buy from any firm they have visited

¹¹ We assume that consumers can return to a previously visited firm at no cost. This could fit Internet search if consumers leave open a browser window. The driving example would not fit this well if consumers cannot purchase over the phone and must drive back to a store to purchase from it.

¹² Costless searchers are not affected by obfuscation. Different reasons might make this appropriate. Costless shoppers could be experts who learn prices in the course of other activities or search via a technology not subject to obfuscation. Alternatively, they could enjoy shopping or receive utility from getting the best price that offsets the time costs.

at any point. We will look for symmetric sequential equilibria of this game; henceforth by “SE” we mean symmetric sequential equilibrium.¹³

□ **Search strategies.** In this section we show that standard results on optimal search strategies carry over to our model. To state this formally, note first that every symmetric strategy profile induces a price distribution $F(p)$. If the price distribution is given by $F(p)$ and a consumer has already spent total time t_0 searching and has observed price p_0 but no lower prices, then the consumer’s expected cost to searching again is $\mathbb{E}_t[g(t_0 + \tau + t) - g(t_0)]$, whereas her expected benefit from searching again and then buying from the lowest-price firm she has observed is

$$\begin{aligned} V(p_0) &\equiv \int_{\underline{p}}^{p_0} \left(\int_x^\infty D(p) dp - \int_{p_0}^\infty D(p) dp \right) f(x) dx \\ &= \int_{\underline{p}}^{p_0} \left(\int_x^{p_0} D(p) dp \right) f(x) dx \\ &= \int_{\underline{p}}^{p_0} \left(\int_{\underline{p}}^p f(x) dx \right) D(p) dp \\ &= \int_{\underline{p}}^{p_0} D(p) F(p) dp, \end{aligned}$$

where \underline{p} is the infimum of the support of $F(p)$.

We begin by showing that optimal consumer search is given by continuing search if $\mathbb{E}_t[g(t_0 + \tau + t) - g(t_0)] < V(p_0)$ and by stopping search if $\mathbb{E}_t[g(t_0 + \tau + t) - g(t_0)] > V(p_0)$.¹⁴

Lemma 1. In any SE, a costly searcher stops searching if $V(p_0) < \mathbb{E}_t[g(t_0 + \tau + t) - g(t_0)]$ or if all N firms have been visited and continues searching if $V(p_0) > \mathbb{E}_t[g(t_0 + \tau + t) - g(t_0)]$.

We present a formal proof in the Appendix. It proceeds by induction on the number of stores remaining using a two-case argument: if the incremental cost of the next search is less than $V(p_0)$, then searching must be optimal because searching exactly once is better than not searching; and if the incremental cost is greater than $V(p_0)$, then not searching must be optimal because convexity of $g(\cdot)$ implies that continuing to search is less appealing than it would be if incremental search costs were constant (i.e., if $g(\cdot)$ were linear), and standard results imply that continuing is not optimal in that case.

Lemma 1 suggests why firms have an incentive to obfuscate. Note that $\mathbb{E}_t[g(t_0 + \tau + t) - g(t_0)]$ is increasing in t_0 , by the convexity of $g(\cdot)$. Hence, by increasing its obfuscation level a firm increases its customers’ future search costs. If $\lim_{t \rightarrow \infty} g'(t) = \infty$, then the argument would be very simple: by using enough obfuscation a firm could completely deter future search by its customers. Our main results, however, hold even if $g'(t)$ is only slightly greater than $g'(0)$ even in the $t \rightarrow \infty$ limit. Here, even a large deviation may not be enough to deter future search depending on how other firms are pricing. But firms can always slightly increase future search costs and our subsequent argument shows that this has a big effect of what can occur in equilibrium.

□ **Price equilibrium.** In this section we recall some standard results for the case where t is a parameter rather than a choice variable and show that properties of these equilibria carry over to our model.

Before doing so, we should note that our model sometimes has equilibria in which the costly searchers are inactive. If exogenous or endogenous search costs are sufficiently high, then costly

¹³ There is no standard definition of sequential equilibrium in games with a continuum of strategies. We use the term for a sequentially rational strategy profile with the restriction on beliefs that a consumer who observes a deviation by one firm continues to believe that other firms use their equilibrium strategies.

¹⁴ This is not a direct corollary of standard results since here the expected search costs faced by consumers depend on the entire history of the obfuscation levels they have encountered.

searchers will not get even a single price quote. We will mostly ignore these equilibria and use the phrase “nontrivial SE” to mean a SE in which the costly searchers all get at least one price quote, which we refer to as “entering.”¹⁵

Proposition 1. (Stahl 1989) Suppose that every firm’s level of obfuscation is fixed exogenously at t . Then the price distribution for any nontrivial SE takes one of two possible forms:

1. If there exists an $r \in (0, p^m)$ for which

$$\int_p^r D(p) \left(1 - \left[\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(\bar{p})}{R(p)} - 1 \right) \right]^{\frac{1}{N-1}} \right) dp = g(2(\tau + t)) - g(\tau + t),$$

then the equilibrium price distribution is $F(p) = 1 - \left[\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(\bar{p})}{R(p)} - 1 \right) \right]^{\frac{1}{N-1}}$, with $\bar{p} = r$, and $R(p) = \left[\frac{1-\mu}{1+(N-1)\mu} \right] R(\bar{p})$.

2. If there does not exist such a value of r , then the equilibrium price distribution is $F(p) = 1 - \left[\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(\bar{p})}{R(p)} - 1 \right) \right]^{\frac{1}{N-1}}$, $\bar{p} = p^m$, and $R(p) = \left[\frac{1-\mu}{1+(N-1)\mu} \right] R(\bar{p})$.

Proof. By Lemma 1, costly searchers search for a second time after observing price p_0 if $V(p_0) > g(2(\tau + t)) - g(\tau + t)$ and do not search for a second time if $V(p_0) < g(2(\tau + t)) - g(\tau + t)$. The result then follows immediately from Stahl’s analysis for $c = g(2(\tau + t)) - g(\tau + t)$. *Q.E.D.*

Note that Proposition 1 can be thought of as showing that two slightly different types of mixed equilibria arise. The first type arises when search costs are small. In these equilibria, the constraint that consumers must be willing to buy from a firm offering price \bar{p} rather than searching again is binding and pins down the upper bound of the support of the price distribution. The upper bound of the support and the distribution of prices vary with the search cost in these equilibria. The second type arises when search costs are larger. In these equilibria, consumers strictly prefer buying from the first firm they visit to getting another price quote. The upper bound of the price distribution is always the monopoly price. The price distribution is independent of the search cost over the range of search costs for which this case applies.

As noted earlier, the second conclusion of Proposition 1 may only hold vacuously in a model with unit demands: if $g(t)$ is linear, then the incentive to search for a second price quote after observing $p = p^m$ is exactly the same as the incentive to obtain an initial price quote, so except for a single value of t no consumers will buy from a firm that sets $p = p^m$ and there cannot be an equilibrium of this form.¹⁶ Equilibria of the second type exist for a broader range of t when search costs are strictly convex because convexity makes second searches more costly than first searches. When we endogenize obfuscation, equilibria of this type will exist under very weak conditions: intuitively firms can choose a level of obfuscation that makes equilibrium search costs fall in the required range.

We now turn our attention to the case where t is endogenous. With an exogenously fixed t and two types of consumers it is well known that every SE price distribution is atomless, that every firm makes positive profits, and that every costly searcher buys from the first firm she visits. The first two results continue to hold generally when t is a choice variable, while the last result requires the additional assumption that obfuscation is costless for firms (which is the case explored in Section 3).

Lemma 2. Every firm makes positive profits in any nontrivial SE.

¹⁵ In a trivial SE only costless searchers enter so firms are Bertrand competitors and price at cost. Obfuscation levels must be high enough so that costly searchers still do not enter. There may also be SE in which costly searchers mix over entering. Janssen, Moraga-Gonzalez, and Wildenbeest (2005) study equilibria in which costly searchers mix in Stahl’s model (without obfuscation).

¹⁶ Stahl (1989, 1996) does not consider unit demands, but the corresponding propositions are more robust to unit demands because of a difference in the model specification. Stahl assumes that costly searchers get their first price quote for free and only pay for subsequent price quotes.

Lemma 3. If $F(p)$ is a nontrivial SE price distribution, then it is atomless.

Lemma 4. If $c(t) = 0$ for all t , then on the equilibrium path of any SE every costly searcher searches at most once.

We defer the proofs of these lemmas to the Appendix. The only one that is nonstandard is Lemma 4. A quick summary of the argument is that we first show that the firm that sets the highest price must choose sufficient obfuscation to prevent consumers from searching again—Lemma 3 implies that the firm would otherwise earn zero profits, violating Lemma 2—and then note that firms setting lower prices must also prevent consumers from shopping again because otherwise they would do better to emulate the highest-priced firm's obfuscation level. Proposition 8 in Section 4 shows that the added assumption that obfuscation is costless is necessary for this result: it presents an example with costly obfuscation where some costly searchers visit two firms before purchasing.

3. An obfuscation model: costless obfuscation and convex disutility of search

■ In this section we analyze our model under the assumption that obfuscation is costless and consumer disutility for shopping, $g(t)$, is strictly convex. The costless obfuscation assumption seems natural for applications where obfuscation consists of not taking a straightforward action that would help consumers. For example, banks can be thought of as practicing obfuscation when they fail to post complete lists of their account fees in a prominent location, and the online retailers discussed in Ellison and Ellison (2009) practice obfuscation when they do not design product description pages to contain all of the information consumers would like. The convex search costs assumption is a departure from much of the previous literature, but we think it is a departure in a realistic direction.¹⁷ Readers who are skeptical that real-world search costs are highly convex should note that our results only require a slight degree of convexity. For example, they will hold even if $g'(t)$ is uniformly bounded above by $g'(0) + \varepsilon$ for a small positive ε . Note that with such slight convexity any obfuscation that firm i does will have at most an order ε effect on the incremental search costs that consumers will incur if they choose to visit another firm.

Our assumption that obfuscation is costless has some drawbacks. One is that there will be substantial indeterminacy as to the actual obfuscation levels chosen in equilibrium. Another is that equilibria will not be strict and it would be a weakly dominant deviation from the equilibrium to set $t = \infty$. Nonetheless, we begin with the model with costless obfuscation because it is the setting in which our main insights come through most cleanly.

The first subsection below discusses the impact of convex search costs on the possible distributions of equilibrium prices. Our main conclusion is that the possibility of obfuscation sometimes has a dramatic effect on the equilibrium price distribution. The second subsection discusses equilibrium obfuscation levels. Here, we derive lower and upper bounds on the amount of obfuscation that can occur in equilibrium. The lower bound points to a second channel through which consumer welfare can be substantially affected. The third subsection presents comparative statics. For example, it shows that equilibrium obfuscation adjusts to offset changes in the exogenous component of consumer search costs.

□ **Equilibrium price distributions.** We may now state our first main result characterizing price distributions with endogenous search costs: if costly searchers search, $F(p)$ is a SE price distribution of the model with endogenous obfuscation if and only if it is the SE price distribution of Stahl's fixed-search cost model that has the upper bound of its support equal to p^m . In other words,

¹⁷ A notable exception is Stiglitz (1987), which considers both convex and concave search costs in a model where consumers observe the market price distribution, though not which firms charge which prices, and focuses on existence of pure-strategy equilibrium and on when firms' demand curves are kinked.

equilibrium price distributions of the first possible form in the characterization of Proposition 1 cannot arise in our endogenous obfuscation model, and the (unique) price distribution of the second possible form is the only possible nontrivial SE price distribution of our model. This price distribution is a nontrivial SE price distribution if and only if exogenous search costs are low enough that consumers are willing to enter when prices are given by this distribution and no firms obfuscate.

Proposition 2. $F(p)$ is a price distribution for a nontrivial SE only if

$$F(p) = 1 - \left[\left(\frac{1 - \mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right]^{\frac{1}{N-1}} \tag{1}$$

for all $p \in [\underline{p}, p^m]$, where \underline{p} is given by $R(\underline{p}) = \left[\frac{1-\mu}{1+(N-1)\mu} \right] R(p^m)$. In particular, $F(p)$ is independent of τ , the exogenous component of consumer search costs. Nontrivial SE exist if and only if the search costs satisfy $g(\tau) \leq V(\infty)$, where $V(\infty)$ is the value of search when a single price quote is drawn from the distribution $F(p)$ described in equation (1). The set of nontrivial SE equals the set of joint distributions over p and t such that the marginal distribution over p equals $F(p)$, and the marginal distribution over t is such that

$$\mathbb{E}_t[g(\tau + t)] \leq \int_{\underline{p}}^{\infty} D(p)F(p) dp, \tag{2}$$

and such that, for all $p_0 \in [\underline{p}, p^m]$ and all t_0 that are ever chosen by a firm that sets price p_0 ,

$$\mathbb{E}_t[g(2\tau + t + t_0)] - g(\tau + t_0) \geq \int_{\underline{p}}^{p_0} D(p)F(p) dp. \tag{3}$$

The proof is again presented in the Appendix. The existence proof consists of showing that $F(p)$ makes firms indifferent over all prices in its support and that there is some obfuscation level for which equations (2) and (3) are both satisfied. Equation (2) is the “entry” condition necessary for consumers to be willing to get an initial price quote. Equation (3) is the “stopping” condition necessary for consumers to never get a second price quote. It is intuitive that both can be satisfied because the benefit of getting an additional price quote is larger when consumers have not yet gotten any price quotes.

The conclusion that this is the only possible price distribution even when exogenous search costs are sufficiently small so that the equilibrium would look very different without obfuscation is the more striking result. The first step in its proof is to note that the equilibrium price distribution must coincide with the equilibrium of Stahl’s model for some search cost. The dramatic selection among these follows from considering the “double deviation” in which a firm sets its price at ε above the upper bound of the price distribution and does enough extra obfuscation to make consumers willing to pay the extra ε ; this slight upward double deviation in price and obfuscation is reminiscent of the slight upward single deviation in price at the heart of the Diamond paradox. Such deviations will be profitable unless the upper bound of the equilibrium price distribution is the monopoly price. This leaves the distribution of Stahl’s model with $\bar{p} = p^m$ as the only equilibrium price distribution. Note that this argument applies even if a firm’s deviation to $t = \infty$ only slightly increases its customers’ expected future search costs.¹⁸

Our intuition for why making g even slightly convex can have a large effect on the equilibrium set comes from considering this double deviation. In Stahl’s model with small exogenous search costs, firms are indifferent between all prices in the interval $[\underline{p}, \bar{p}]$. When setting price \bar{p} firms know they have zero probability of attracting costless shoppers and profit only from selling to costly shoppers who visit them first. They would make more money off the costly shoppers if

¹⁸ In some applications $g(\cdot)$ could be concave, e.g., if there is learning-by-doing in search. In this case, firms increase future search costs by *decreasing* their search time. Hence, either obfuscation will be sufficient to lead to the equilibrium in Proposition 2 or perhaps the other equilibrium in Proposition 1 might occur with the firms setting the highest prices using zero obfuscation.

they charged them a higher price, but they cannot do so because at price \bar{p} consumers are exactly indifferent between purchasing and getting a second price quote. Having even a slight ability to affect consumers' future search costs makes it profitable to charge slightly above the highest price being charged in equilibrium. This type of unraveling ensures that obfuscation must be sufficient to shift the upper bound of the price distribution all the way up to p^m even when it would be much lower without obfuscation.

One striking fact about Proposition 2 is that, as long as $g(\tau) \leq V(\infty)$, any reduction in the exogenous fixed component of consumer search costs has no effect whatsoever on the equilibrium distribution of prices and profits—any reduction in τ that would lead to lower prices must be offset by changes in the equilibrium level of obfuscation (see Proposition 3). Hence, our model provides a formalization of the observation in Ellison and Ellison (2009) that improvements in search technology need not make search more efficient. Their empirical findings are consistent with the idea that the reduction in search costs online have led to greater equilibrium obfuscation, although probably not with the extreme finding of this section that the response can be so large as to keep the price distribution unchanged.

□ **Equilibrium obfuscation levels.** We now consider equilibrium obfuscation levels. Obfuscation can create substantial search costs which are another important channel of welfare effects. There is, however, substantial indeterminacy about details of the obfuscation pattern.

The fact that all equilibria must have $\bar{p} = p^m$ puts a lower bound on equilibrium search costs—consumers must not be willing to conduct a second search when they have found price p^m and know that prices are drawn from the distribution given in Proposition 2. If the exogenous component τ of the search costs is not too large, this implies that firms must obfuscate in equilibrium.

Corollary 1. If $g(2\tau) - g(\tau) < \int_p^{p^m} D(p)F(p) dp$, where $F(p)$ is given by equation (1), then in any nontrivial SE some firms set $t > 0$.

Proof. By Proposition 2, firms are willing to set price equal to p^m in any SE in which costly searchers enter. If $g(2\tau) - g(\tau) < \int_p^{p^m} D(p)F(p) dp$ and $t_i = 0$ for all i , then a costly searcher who first observes a price sufficiently close to p^m will search again, contradicting Lemma 4.

Q.E.D.

We next note the basic welfare consequences of obfuscation. Consumers suffer both directly from the effect of obfuscation on search costs and indirectly because obfuscation leads to higher prices. Firms benefit from the higher prices.

Corollary 2. Suppose $g(2\tau) - g(\tau) < \int_p^{p^m} D(p)F(p) dp$, where $F(p)$ is given by equation (1). Compared to the model in which obfuscation is impossible ($t = 0$ identically for all firms), the model in which obfuscation is possible leads to higher prices in sense of first-order stochastic dominance, higher profits for all firms, and lower utility for all consumers (both costly searchers and costless searchers) in every nontrivial SE.

Proof. In the model where $t = 0$, the possible nontrivial SE price distributions are given by Proposition 1. If obfuscation is possible, the nontrivial SE price distribution is given by Proposition 2. Since the formula for $F(p)$ is the same in both cases and is decreasing in \bar{p} for all p , prices are higher in the sense of first-order stochastic dominance in the latter SE. This and the fact that obfuscation reduces consumer welfare directly imply that in the latter SE firms earn higher profits and consumer welfare is lower.

Q.E.D.

Although we did not emphasize this when we first stated it, Proposition 2 in fact fully characterizes the range of possible obfuscation levels via entry and stopping conditions in equations (2) and (3). The entry condition requires that the expected cost of the first search be sufficiently low, which places an upper bound on how much obfuscation firms can be doing.

And the stopping condition requires that consumers do not want to carry out a second search after each price p that is observed in equilibrium, which provides a continuum of lower bounds. All of these bounds are bounds on average obfuscation in some sense. In addition to the distance between the lower and upper bounds, there is substantial indeterminacy in where the obfuscation occurs. For example, all firms could set the same obfuscation level, or firms could mix over high and low obfuscation levels. And there is a lot of indeterminacy in the cross section relationship between obfuscation and price. We return to this issue in Section 4, where the possibility that firms may prefer not to use too much obfuscation greatly reduces this indeterminacy.

As noted earlier another unappealing consequence of our assumption that obfuscation is costless is that setting $t = \infty$ is a weakly dominant deviation from the equilibrium. Indeed, if one were to modify our model so that the heterogeneity in consumer search costs had full support, then there would be no nontrivial equilibria because firms would always be strictly better off setting $t = \infty$ (thereby deterring some marginal searcher). In practice, we think this concern is less troubling because of another effect we have left out—consumers would move on from one firm to the next before learning the first firm’s price if the search was taking much longer than they had expected. A working paper available from the authors includes a variant of the model along these lines.

□ **Comparative statics with price-independent obfuscation.** In this section, we provide some comparative statics on obfuscation by restricting attention to equilibria in which all firms use the same pure obfuscation strategy (while mixing over prices). When obfuscation is independent of price, the entry and stopping conditions reduce to equations (A1) and (A2) in the Appendix. Multiple obfuscation levels are consistent with these bounds so our comparative statics are on sets of equilibria with respect to the strong set order. Recall that a (one-dimensional) set X is higher than Y in the strong set order if, given elements $x \in X$ and $y \in Y$, the maximum of x and y is in X while the minimum of x and y is in Y .¹⁹

Our first result identifies a sense in which obfuscation levels must rise when the exogenous component of search costs falls.

Proposition 3. The set of obfuscation values t^u (for uniform) played in any nontrivial SE in which firms do not mix over obfuscation levels is decreasing in τ , the exogenous component of consumer search costs, in the strong set order.

Proof. As in the proof of Proposition 2, the lower bound on t^u is given by $g(2(\tau + t^u)) - g(\tau + t^u) = \int_{\underline{p}}^{p^m} D(p)F(p) dp$ and the upper bound on t^u is given by $g(\tau + t^u) = \int_{\underline{p}}^{\infty} D(p)F(p) dp$, and an increase in τ causes both of these bounds to decrease. Q.E.D.

Proposition 3 shows that changes in equilibrium obfuscation offset changes in the exogenous component of search costs. This follows because high exogenous search costs rule out equilibria with high obfuscation by the entry condition, and eliminate the need for high obfuscation by the stopping condition. That is, costly searchers will not be willing to obtain a price quote if they face both high exogenous search costs and high obfuscation, and firms have no need to set high obfuscation when consumers are already deterred from comparison-shopping by high exogenous search costs. This effect, however, is weak enough that an increase in τ must nonetheless lead to a decrease in the set of nontrivial SE values of consumer welfare. The intuition is that prices are fixed by Proposition 2, and any nontrivial SE value of consumer welfare given $\tau' \geq \tau$ can be reproduced by uniformly increasing obfuscation by $\tau' - \tau$. Note that Proposition 4 does not restrict attention to equilibria with price-independent obfuscation.

Proposition 4. The set of nontrivial SE values of the costly searchers’ welfare is decreasing in the exogenous search cost τ in the strong set order.

The proof is deferred to the Appendix.

¹⁹ See Milgrom and Shannon (1994) for more on the strong set order.

In addition, one can use the comparative statics presented in Stahl (1989) to derive a number of other comparative statics results. For example, we can show that the lower and upper bounds on the SE obfuscation level are both increasing in the proportion of costless searchers, and, if there are enough firms in the market, decreasing in the number of firms. The argument again involves considering the entry and stopping conditions. When there are more costless searchers, Stahl shows that SE prices are lower in the sense of first-order stochastic dominance. Therefore, more obfuscation is needed to prevent costly searchers who first observe price p^m from searching again, so the lower bound on equilibrium obfuscation increases by the stopping condition. Similarly, lower prices imply that costly searchers would be willing to enter despite higher obfuscation, so the upper bound on equilibrium obfuscation increases by the entry condition. We state this result below as Proposition 5. Proposition 6, which gives comparative statics with respect to the number of firms, is closely related. One starts by recalling Stahl's finding that an increase in the number of firms increases SE prices once the number of firms is sufficiently large. Hence, obfuscation will be lower by the same argument as in Proposition 5. The details of both proofs are deferred to the appendix.

Proposition 5. The set of obfuscation levels t^u played in any nontrivial SE where firms do not mix over obfuscation levels is increasing in μ , the proportion of costless searchers, in the strong set order.

Proposition 6. There exists $\bar{N} > 0$ such that, if $N > \bar{N}$, then the set of obfuscation levels t^u played in any nontrivial SE where firms do not mix over obfuscation levels is decreasing in N , the number of firms, in the strong set order.

Note that the comparative static above only applies when the number of firms is above some threshold. The monopoly versus oligopoly comparison goes in the opposite direction. The lower bound on equilibrium obfuscation would always be zero in a monopoly model because there is no benefit to obfuscation—detering consumers from obtaining additional price quotes is not relevant when there are no other firms. And the upper bound on possible obfuscation levels would also be lower for a monopolist because the entry condition must be satisfied despite the monopolists' higher price.

One can easily show that the set of values of consumer welfare is increasing in μ and decreasing in N for $N > \bar{N}$. The connection between Propositions 5 and 6 and this fact is the same as that between Propositions 3 and 4: raising μ (for example) leads to higher obfuscation only because costly searchers benefit more from entering the market when μ is high and are thus willing to tolerate more obfuscation, so obfuscation cannot be so much higher that costly searchers benefit less on net from entering.

4. Costly obfuscation

■ We argued above that obfuscation is sometimes costless: it can, for example, consist simply of not taking actions that would help consumers. In other applications, however, obfuscation seems costly. For example, mattress stores make price comparisons more difficult by getting manufacturers to label the same mattress with a different model name at each store that sells the product. There must be some cost associated with this. And even in examples without direct costs like a bank's adoption of a complex fee structure, there may be indirect costs such as an increase in customer service costs to deal with questions and complaints. In this section we examine models in which the cost $c(t)$ that firms must incur in order to raise the time cost of search to t is not identically zero. Obfuscation costs make the model less tractable, but also eliminate much of the indeterminacy of the previous section making "equilibrium obfuscation" more sharply defined. We discuss the extent to which some results of the previous section carry over, provide a characterization of equilibrium obfuscation, and note some interesting patterns that may arise.

Our first proposition notes that our previous characterization of the equilibrium distribution of prices sometimes still applies to models with costly obfuscation and sometimes does not. Whether the costless-obfuscation equilibrium carries over depends on the consumer search-disutility function $g(t)$ and other aspects of the model, but is largely independent of the obfuscation cost function $c(t)$.

Proposition 7. Let $V(p)$ be the consumer benefit from search assuming that prices are distributed as in the nontrivial SE of the costless-obfuscation model. Suppose that the obfuscation cost function $c(t)$ is twice continuously differentiable with $c'(t) > 0$ for $t > 0$.

- (a) If $g(\tau) < V(\infty)$ and $g(2\tau) - g(\tau) > V(p^m)$, then the model with costly obfuscation has a unique nontrivial SE, in which no obfuscation occurs and the equilibrium price distribution coincides with that of the costless-obfuscation model.
- (b) If $g(2\tau) - g(\tau) < V(p^m)$, then any nontrivial SE of the costly obfuscation model must have a price distribution with supremum \bar{p} strictly less than p^m .

Part (a) of the proposition notes that our previous characterization of the equilibrium price distribution remains valid with costly obfuscation if the unavoidable search costs τ and other factors are such that no obfuscation is necessary to deter consumers from searching a second time. In this case, firms will not engage in costly obfuscation.

Part (b) implies that the equilibrium must be different from that characterized in the previous section when consumers will conduct a second search if firms do not obfuscate—one difference is that the upper bound of the price distribution is now strictly less than the monopoly price instead of being equal to the monopoly price. We defer the proof to the Appendix, but the two-step intuition is fairly simple. The first step is showing that firms pricing near \bar{p} must sell to those consumers who visit them first, as otherwise these firms would make no sales. The second is that the upper bound cannot be p^m because otherwise firms would profit from slightly decreasing both prices and obfuscation: lowering prices has a second-order effect on profits conditional on making a sale and allows a first-order reduction in obfuscation costs. Note that this second step does not go through if consumers have unit demands.


In the costless-obfuscation model there was substantial indeterminacy in obfuscation levels (and hence in the price-obfuscation relationship). A second basic observation is that making obfuscation costly eliminates much of this indeterminacy because firms will never do more obfuscation than is necessary to deter consumers who visit them from searching for a second time. However, we show below that firms may do *strictly less* obfuscation than is needed to deter their customers from conducting a second search, because the loss in sales from not obfuscating may be outweighed by the costs that would be required to deter future searches. Therefore, unlike in the textbook model with constant search costs, the consumers in our model with convex search costs may conduct a second search but then return to the first firm to purchase before exhausting all possibilities. We feel that the fact that our model can generate patterns of search behavior in which consumers return to a previous store to purchase is an attractive feature consistent with empirical evidence on search such as that in De los Santos, Hortacsu, and Wildenbeest (forthcoming).

To show that choosing not to deter search is something that will happen for some specifications (as opposed to just being something that we can't show doesn't happen) we present the example below.

Proposition 8. Suppose that the demand function is $D(p) = 1 - p$, consumer disutility of search is $g(t) = 0.15t + \max\{t - 1, 0\}$, and that the obfuscation cost function is such that the firm can choose obfuscation levels 0, t^* and t^{**} at costs 0, c^* , and c^{**} , respectively. Then for some parameter values the model has a symmetric SE in which firms use a strategy of the form shown in Figure 1. In this equilibrium the support of the price distribution is the union of four intervals: $[\underline{p}, p_1] \cup [\hat{p}_1, p_2] \cup [\hat{p}_2, p_3] \cup [p_3, \bar{p}]$. Firms choosing prices in the third interval do no obfuscation and costly searchers who see a price in this interval on their first search conduct

FIGURE 1

AN EXAMPLE WITH SECOND SEARCHES AND NONMONOTONE OBFUSCATION



a second search. Firms setting prices in the other three intervals use obfuscation 0, t^* , and t^{**} , respectively, and thereby deter costly searchers who visit them from conducting a second search.

The proof of Proposition 8 is numerical: the figure was generated by solving the model numerically for particular parameter values. In particular, the equilibrium shown is for a model with four firms, 30 percent of consumers being costless searchers, and obfuscation cost function having parameters $\tau \approx 0.347$, $t^* \approx 0.023$, $t^{**} \approx 0.308$, $c^* \approx 0.0044$, and $c^{**} \approx 0.308$.²⁰

Another striking feature of the example above is that obfuscation levels are not monotonically related to prices. We think it is useful to point out the possibility both because it may be unexpected and because it calls attention to the fact that the cost and benefit of obfuscation are both increasing in a firm's price (it also implies that an outside observer—such as a regulator—cannot conclude that the highest-priced firms in a market are necessarily the ones who use the most obfuscation, or vice versa). But there are also senses in which our model does predict that obfuscation will be increasing in prices. One such sense is a comparison of extremes: if we suppose that $c(t)$ is strictly increasing then there will be an interval of positive length $[\underline{p}, \bar{p}]$ above the lowest price on which firms set zero obfuscation, and the highest obfuscation level chosen will be that of the firm that sets $p = \bar{p}$. No obfuscation is used when prices are very low because the exogenous search costs are by themselves sufficient to deter a second search for consumers who first observe a price sufficiently close to \underline{p} . The fact about the top of the distribution follows from something we noted in the proof of Proposition 7: a firm that sets price \bar{p} must choose a sufficiently high obfuscation level so that consumers who visit it first purchase with probability one.

Another result with the same flavor is that below which considers what happens when obfuscation costs are small. An intuition for the result is that when obfuscation costs are sufficiently small all firms will choose obfuscation levels that are just large enough to deter consumers from conducting a second search. This level is lower when a firm's price is lower.

²⁰ The disutility of search and obfuscation cost functions do not satisfy some of our regularity conditions— g is nondifferentiable and only weakly convex and the obfuscation cost is discontinuous and only weakly increasing. The departures make the model easier to solve numerically, but nearby smooth models should have similar equilibria.

An increasing price-obfuscation relationship contrasts with the predictions of the model of Wilson (2010). In Wilson's two-firm model, the firm that obfuscates chooses prices from a distribution that is lower than the distribution from which the non-obfuscating firm chooses its prices.

Proposition 9. Suppose $g(t)$ is an increasing, strictly convex function with $\lim_{t \rightarrow \infty} g'(t) = \infty$. Let $C(t)$ be strictly increasing. Then there exists a $\underline{\delta} > 0$ such that for any $\delta \in (0, \underline{\delta})$, any nontrivial SE of the model with obfuscation cost function $c(t) = \delta C(t)$ is such that firms do no obfuscation when they choose prices in some interval $[\underline{p}, \hat{p}]$ and obfuscation is strictly increasing in price at all higher prices.

One assumption that we have maintained throughout the article is that $c'(t) \geq 0$. In applications where obfuscation consists of not taking actions that would help consumers it is plausible that obfuscation costs might be decreasing in t at least for small t (indeed, it is natural that it is costly for firms to reduce consumer waiting time to literally zero). We do not analyze this case formally, but some observations are that firms would always obfuscate up to the point where costs start increasing and that this can lead to nonexistence of nontrivial equilibria. In particular, any model in which obfuscation costs are decreasing up to some point and increasing thereafter is equivalent (in equilibrium) to a model where exogenous search costs are higher and obfuscation costs are everywhere increasing, as one can reclassify the “cost-saving” obfuscation that all firms will do as part of the exogenous search cost.

This prediction that firms will in equilibrium take any cost-saving action that increases consumer search costs is of course counterfactual. In our view, the key element that is missing from the model that would overturn this prediction is the possibility that consumers can abandon search at a firm before learning its price. To take an extreme example, Bank of America would not benefit from reducing the size of its customer service staff to one person, even though this would surely save costs and increase waiting times, because consumers would not wait on the line for long enough to get through to the lone representative. An earlier version of this article (available from the authors) investigated the possibility of abandoning search formally. The key implication of that model was that, under some conditions—in particular, the condition that consumers complete all searches they begin *in equilibrium*—many of the qualitative conclusions of the baseline model continue to hold.

5. A signal-jamming model

■ In this section we explore an alternate mechanism through which obfuscation can affect consumer search. In particular, we show that allowing the exogenous component of consumer search costs τ to be uncertain makes obfuscation individually rational for firms even if search costs are *linear*. The basic idea behind this *signal-jamming* mechanism is straightforward: if τ is initially unknown, consumers learn about τ from their first shopping experience, so obfuscation raises consumer expectations about the search costs they will incur on future searches.^{21,22} This mechanism seems plausible for many applications. For example, if a home-improvement contractor spends a long time with a consumer discussing details about the job and takes a long time to prepare and submit his or her bid, then it seems plausible that consumers will expect that getting a second bid will entail similar time costs. While obfuscation is individually rational in this model, the differences in the mechanics of this model and that in Section 3 lead to some differences in the results. Among these are an *excess obfuscation* problem that leads to both lower

²¹ We call this a *signal-jamming* model because obfuscation distorts the signals consumers get about the exogenous search costs. Of course, consumers account for this in equilibrium when forming their beliefs. See Holmström (1982) or Fudenberg and Tirole (1986) for seminal models of signal-jamming.

²² In standard signal-jamming models, the marginal benefit of signal jamming is positive and the equilibrium quantity is determined by marginal costs that increase in the quantity of signal-jamming. In our model, the marginal cost of signal-jamming is zero, so equilibrium requires that the marginal benefit be zero.

prices and lower consumer welfare, and a different selection among the possible equilibria of Stahl's model.

Formally, assume that $g(t) = t$, and that there is an underlying parameter θ with expectation zero distributed with continuous density $h(\theta)$ on $[\underline{\theta}, \infty)$ with $\underline{\theta} > -\tau$, such that it costs a consumer $\tilde{t} = \tau + \theta + t_i$ to visit a firm that sets obfuscation level t_i . We assume that when a consumer visits firm i , she observes only \tilde{t} and p_i , so that she must draw inferences about θ in equilibrium. The timing of the game is almost as before. The one amendment is that we assume that θ is drawn once and for all at the beginning of the game and is unobserved by both firms and consumers.

We focus in this section on costless obfuscation. As before, this lets us bring out our main observations most simply. We also restrict our attention to strategies for firms which do not mix over obfuscation levels for a given price. That is, we consider equilibria in which there exists a function $t^*(p)$ such that the support of firms' mixed strategies is contained in the set of ordered pairs (p, t) with $t = t^*(p)$. This restriction can be motivated by thinking of obfuscation as being "slightly costly" for firms, so that firms do the smallest amount of obfuscation that maximizes profit, taking price as given. The importance of this assumption is that equilibrium implies that consumers believe with probability one that $\theta = \hat{\theta} \equiv \tilde{t} - (\tau + t^*(p))$ after observing total search cost \tilde{t} and price p .²³ The only exception to this, of course, is if the observed (p, t) is inconsistent with equilibrium, which can happen if p is not in the support of the equilibrium price distribution or if $\tilde{t} < \tau + t^*(p) + \underline{\theta}$. Consumers' beliefs about θ are unrestricted in this case, as well as after observing longer sequences of price-obfuscation pairs that are jointly inconsistent with equilibrium (e.g., if the implied values θ resulting from visiting two firms are not the same). Throughout this section, "SE" should be understood to mean, "SE in which firms do not mix over obfuscation levels for a given price."

In this section, consumer search behavior in the signal-jamming model is similar to search behavior in our previous model. First, we have a straightforward application of standard results.

Lemma 5. In any SE, a costly searcher searches for the first time if $\tau + \mathbb{E}[t(p)] < V(\infty)$ and continues to search if $\tau + \hat{\theta} + \mathbb{E}[t(p)] < V(p_0)$ and there are previously unsearched firms remaining. Conversely, a costly searcher does not search for the first time if $\tau + \mathbb{E}[t(p)] > V(\infty)$ and stops searching if $\tau + \hat{\theta} + \mathbb{E}[t(p)] > V(p_0)$.

Next, we observe that costless obfuscation implies that costly searchers search at most once in equilibrium, as in the convex search cost model of Section 3.²⁴

Lemma 6. In any SE of the signal-jamming model with costless obfuscation, all costly searchers search at most once.

Proof. Fix a SE obfuscation strategy $t(p)$ and suppose that a firm does not sell to all costly searchers that visit it first when it sets price p_0 and obfuscation level $t(p_0)$. A consumer who first visits a firm with price p_0 and incurs total search costs \tilde{t} will buy if $\tilde{t} - t(p_0) + \mathbb{E}[t(p)] > V(p_0)$, since in equilibrium the consumer infers that $\theta = \tilde{t} - (t(p_0) + \tau)$. Therefore, a firm can always induce those consumers who visit it first to buy with probability 1 by setting $t > t(p_0) + V(p_0) - \mathbb{E}[t(p)]$. By Lemma 5, this maximizes the market share of the firm, so there cannot be a SE in which a firm does not sell to all consumers who visit it first. Q.E.D.

The fact that consumers search at most once will again allow us to provide simple, closed-form expressions for the possible SE price distributions. It is the primary place where we use the assumption that the distribution of θ is unbounded. If θ is bounded, it may be that consumers sometimes search multiple times.

²³ Without this result, the consumer search problem could become much more complicated, as consumers might have an incentive to search multiple times in order to learn more about θ .

²⁴ Unlike our earlier results, this result does rely on assumptions that imply that firms can always dissuade consumers from future search by setting very high obfuscation.

The first main result of this section is that we get an incomplete selection from the equilibria of Stahl’s model, and that many price distributions may be possible; in addition, there may not exist an equilibrium in which $\bar{p} = p^m$.

Proposition 10. The signal-jamming model has a nontrivial SE price distribution with supremum p^m if and only if $\tau \leq V(\infty)$ and $-\underline{\theta} \leq \int_{p^m}^{\infty} D(p) dp$. For every $p^* \in [0, p^m)$, the signal-jamming model has a nontrivial SE price distribution $F_{p^*}(p)$ with supremum p^* only if

$$F_{p^*}(p) = 1 - \left[\left(\frac{1 - \mu}{N\mu} \right) \left(\frac{R(p^*)}{R(p)} - 1 \right) \right]^{\frac{1}{N-1}} \tag{4}$$

for all $p \in [\underline{p}^*, p^*]$, where \underline{p}^* is given by $R(\underline{p}^*) = \left[\frac{1-\mu}{1+(N-1)\mu} \right] R(p^*)$. Such a SE exists if and only if $\tau + \underline{\theta} \leq \int_{\underline{p}^*}^{p^*} D(p) F_{p^*}(p) dp$ and $-\underline{\theta} \leq \int_{p^*}^{\infty} D(p) dp$. Under these conditions, the set of nontrivial SE with the supremum of $F_{p^*}(p)$ equal to p^* is the set of joint distributions over p and t such that the marginal distribution over p equals $F_{p^*}(p)$ and the marginal distribution over t is such that

$$\tau + \mathbb{E}[t] \leq \int_{\underline{p}^*}^{\infty} D(p) F_{p^*}(p) dp \tag{5}$$

and such that for all $p_0 \in [\underline{p}^*, p^*]$,

$$\tau + \underline{\theta} + \mathbb{E}[t] \geq \int_{p^*}^{p_0} D(p) F_{p^*}(p) dp, \tag{6}$$

with equality for $p_0 = p^*$. Furthermore, some nontrivial SE exists if $\tau \leq \int_0^{\infty} D(p) F_{p^m}(p) dp$.

The proof is given in the Appendix. To understand the result, consider first the conditions for the existence of an equilibrium with $\bar{p} = p^m$. The first condition, $\tau \leq V(\infty)$, is simply the requirement that the exogenous search costs are not high enough to prevent consumers from searching at least once. It is analogous to the sole condition required for the existence of a nontrivial SE in the convex search cost model. The second condition, $-\underline{\theta} \leq \int_{p^m}^{\infty} D(p) dp$, is an additional restriction requiring that the consumer surplus from purchasing at price p^m also be sufficiently large relative to the uncertainty about θ . This reflects that there is what one can think of as “excess obfuscation” in the signal-jamming model with costless obfuscation. In equilibrium, firms must obfuscate to the point where consumers will not want to search again even when the exogenous component of search costs turns out to take on its lowest possible value. This implies that, with probability one, the average obfuscation level in any SE is higher than the minimal average obfuscation level needed to keep all costly searchers from conducting a second search, conditional on θ . This excess obfuscation makes it harder to sustain equilibria in which consumers search, as the lowest $\mathbb{E}[t]$ such that $\tau + \underline{\theta} + \mathbb{E}[t] \geq V(p^m)$ (the stopping condition) may be so high that $\tau + \mathbb{E}[t] > V(\infty)$ (the negation of the entry condition), precluding costly searcher entry. To put this another way, in the convex search costs model we could always simultaneously satisfy the entry and stopping conditions as long as $g(\tau)$ was less than $V(\infty)$. Now, the difference between $\underline{\theta}$ and $\mathbb{E}[\theta]$ drives a wedge between the entry and stopping conditions, so we may not be able to satisfy them both simultaneously.

The set of nontrivial SE in the signal-jamming model also differs from that in the convex costs model in that the selection from equilibrium price distributions of Stahl’s model may now include price distributions with $\bar{p} < p^m$. In the convex costs model, the key deviation that prevented these distributions from being equilibria was that a firm could charge a price slightly above \bar{p} and obfuscate slightly more so that consumers would not search again. This constraint on the SE set no longer exists in the signal-jamming model: the above deviation is a deviation to an out-of-equilibrium price, so the firm cannot necessarily induce consumers to have the beliefs about θ that it would like them to have after such deviations. In particular, a consumer who observes a price above \bar{p} may believe that the firm that set this price also set very high obfuscation, which would lead her to believe that θ is very low. Indeed, the conditions for the

existence of an equilibrium with $\bar{p} = p^* < p^m$ are analogous to the conditions for the existence of an equilibrium with $\bar{p} = p^m$, with the exception that the condition that $\tau \leq \int_{\underline{p}}^{\infty} D(p)F_{p^m}(p) dp$ is strengthened to $\tau + \underline{\theta} \leq \int_{\underline{p}^*}^{p^*} D(p)F_{p^*}(p) dp$ (this condition is stronger given the second condition that $-\underline{\theta} \leq \int_{p^*}^{\infty} D(p) dp$). This stronger condition reflects the fact that the stopping condition (equation (6)) must hold with equality at $p_0 = p^*$ if $p^* < p^m$, as otherwise a firm with price p^* would deviate to a slightly higher price, while if $p^* = p^m$ then the stopping condition may hold with strict inequality for all p_0 in equilibrium.

Finally, observe that even though Proposition 10 places both lower and upper bounds on equilibria p^* , the last sentence of the proposition shows that existence of some nontrivial SE is guaranteed under conditions identical to those in the convex search cost model.

The working paper version of this article (Ellison and Wolitzky 2009) contains several observations about the emergence of equilibrium obfuscation and its effect on prices, profits, and consumer welfare, in analogy with the analysis of the convex costs model in Section 3. The results here are slightly more subtle due to the range of SE described in Proposition 10, but two important ideas—that obfuscation occurs in equilibrium, and that obfuscation offsets changes in the exogenous component of consumer search costs—continue to come through. Here, we only present a pair of results on the excess obfuscation problem discussed above, which is the most important effect that is present here but not in the convex costs model: prices actually fall as the excess obfuscation problem becomes more severe, i.e. as $\underline{\theta}$ decreases, holding $\mathbb{E}[\theta]$ fixed at 0; but nonetheless consumer welfare falls as $\underline{\theta}$ decreases. Thus, the excess obfuscation problem hurts both firm and consumer welfare, even though it makes markets more competitive in the sense of lowering prices.

For these last results, we impose the following assumption:²⁵

Assumption 1. $\int_{\underline{p}^*}^{p^*} D(p)F_{p^*}(p) dp$ is increasing in p^* for $p^* < p^m$, where F_{p^*} is given by equation (4).

Recalling that an increase in $\underline{\theta}$ corresponds to a decrease in the severity of the excess obfuscation problem, Proposition 11 shows that an increase in the severity of this problem leads to a decrease in prices, which is perhaps a surprising result. The intuition here is that the excess obfuscation problem rules out equilibria with high prices, because under Assumption 1 higher prices correspond to both lower consumer welfare and a smaller gap between consumer welfare and the benefit of a second search conditional on observing p^* and $\underline{\theta}$, which makes it more likely that the excess obfuscation needed to satisfy the stopping condition (6) is so great as to violate the entry condition (5). Proposition 12 (proof in Appendix) shows that, if all equilibria have $\bar{p} < p^m$, this effect cannot overturn the direct welfare costs to consumers of an increase in excess obfuscation, because excess obfuscation leads to lower prices only by making consumers sufficiently worse off that they refuse to enter when prices are high. So long as equilibria with $\bar{p} = p^m$ do not exist, then Propositions 11 and 12 show that, while equilibrium requires that firms extract some of the additional surplus that comes with an increase in $\underline{\theta}$ through reduced excess obfuscation, consumers are still better off after such a reduction in uncertainty.²⁶

Proposition 11. The set of nontrivial SE values of \bar{p} is increasing in $\underline{\theta}$ in the strong set order.

Proof. Recall that p^* is a nontrivial SE value of \bar{p} if and only if $\int_{\underline{p}^*}^{p^*} D(p)F_{p^*}(p) dp \geq \tau + \underline{\theta}$ and $\int_{p^*}^{\infty} D(p) dp \geq -\underline{\theta}$. And $\int_{\underline{p}^*}^{p^*} D(p)F_{p^*}(p) dp$ is increasing in p^* for $p^* < p^m$ by Assumption 1,

²⁵ This assumption is implied by Assumption C in Stahl (1989), which is the same as the “Revenue Condition” in Stahl (1996). As Stahl (1996) points out, this condition holds “for all concave (and linear) demand functions, as well as many convex demand functions.”

²⁶ The assumption that equilibria with $\bar{p} = p^m$ do not exist is needed because equilibria with $\bar{p} = p^m$ can have very high obfuscation, whereas equilibria with $\bar{p} < p^m$ must have low enough obfuscation so that firms are not tempted to deviate to prices slightly above \bar{p} . This makes comparing consumer welfare across equilibria with $\bar{p} < p^m$ and $\bar{p} = p^m$ difficult.

while $\int_{p^*}^{\infty} D(p) dp$ is decreasing in p^* , so an increase in $\underline{\theta}$ raises the lower bound on \bar{p} given by the first inequality and raises the upper bound of \bar{p} given by the second. *Q.E.D.*

Proposition 12. Suppose that no SE with $\bar{p} = p^m$ exist when the lower bound on θ equals $\underline{\theta}$ or $\underline{\theta}'$ for some $\underline{\theta}' \geq \underline{\theta}$. Then increasing the lower bound on θ from $\underline{\theta}$ to $\underline{\theta}'$ increases the set of SE values of consumer welfare in the strong set order.

6. Conclusion

■ In this article we have explored obfuscation using two related models in which obfuscation is treated as an action that increases the amount of time that consumers must spend to learn a firm's price. In both cases, the key impact of such actions is that they lead consumers to behave as if future search costs will be higher. In the convex costs model this is because obfuscation directly increases the incremental costs that consumers would incur to perform another search. In the signal-jamming model there is no real effect on the future, but an informational linkage implies that increased obfuscation leads consumers to expect higher future search costs.

In both models, we show that obfuscation must occur in an equilibrium unless the exogenous component of consumer search costs is high enough that consumers are willing to purchase at the highest equilibrium price in the absence of obfuscation. And we show that obfuscation has the same qualitative impact on welfare. It is bad for consumers both because it directly imposes costs on them and because it leads to higher prices. The higher prices make obfuscation beneficial for firms, except in the case when excess obfuscation makes the market completely collapse. Note that obfuscation benefits all firms, not only those who engage in it; even transparent firms benefit from serving an obfuscation-rich market, as their customers are prevented from comparison-shopping by other firms' obfuscation.

The mechanics of our models are similar to those of Stahl (1989). In both cases obfuscation can be seen as selecting among the dispersed price equilibria of Stahl's model. In the convex costs model, the selection is that obfuscation must be sufficiently high to result in an equilibrium price distribution that goes all the way up to the monopoly price. In the signal-jamming model, the constraints are that overall obfuscation levels must be sufficiently high so that consumers are willing to search once, but never more than once. This can leave a range of possible dispersed-price equilibria.

Our two models also have similar comparative statics implications. In both, equilibrium obfuscation adjusts to offset changes in the exogenous component of consumer search costs, though in equilibrium consumers still benefit from reductions in exogenous search costs and are hurt by increases in these costs (see Ellison and Wolitzky, 2009). The signal-jamming model is also distinguished by the fact that it displays excessive obfuscation with probability one; this effect leads to Pareto inefficiency by decreasing both prices and consumer welfare.

Our analysis suggests a number of interesting avenues for future research. Our characterizations of the costly obfuscation model are limited. In reality we feel that it takes a great deal of cleverness for firms to devise effective obfuscation schemes, which could make such schemes quite costly. Such costs would be natural candidates for explaining why real-world obfuscation is limited. For example, we noted that whereas our convex costs model with costless obfuscation predicts that obfuscation will completely offset any technological reduction in search costs, Ellison and Ellison (2009) report that search is still fairly effective for at least some consumers in the environment they study. Developing models of costly obfuscation that are more tractable than ours could be challenging, but could have rewards both from a theoretical and from an applied perspective.

Finally, we note that there are more basic related areas of search theory that have not been fully explored. We showed in Section 4 that the combination of convex search costs and costly obfuscation creates an environment in which search strategies are more interesting and realistic, with different consumers searching different numbers of times. Wolinsky's (1986) model

of search with product differentiation provides a natural way to account for such behavior in some applications, and Stahl (1996) explores one way to get such behavior without product differentiation by adding heterogeneous search costs, but further analyses of the convex cost model—with or without obfuscation—could be a valuable complement and provide additional insights.²⁷

Appendix

Proof of Lemma 1. We proceed by induction on the number of remaining unsearched stores. The fact that any strategy not of this form yields a lower payoff is immediate when one store remains unsearched. Now, assume that we have shown the result for all numbers of remaining unsearched stores up to m and consider possible continuation strategies at a history $x_n = ((p_1, t_1), (p_2, t_2), \dots, (p_n, t_n))$ at which $m + 1$ stores remain unsearched. Let $p_0 = \min\{p_1, \dots, p_n\}$ and $t_0 = n\tau + \sum_{i=1}^n t_i$.

First, suppose that $V(p_0) > \mathbb{E}_t[g(t_0 + \tau + t) - g(t_0)]$. Here, we show that any strategy that involves stopping at x_n cannot be optimal. To see this, note that stopping at x_n yields utility $\int_{p_0}^{\infty} D(p) dp - g(t_0)$. Continuing at x_n and stopping at x_{n+1} regardless of (p_{n+1}, t_{n+1}) yields expected utility $\int_{p_0}^{\infty} D(p) dp + V(p_0) - \mathbb{E}_t[g(t_0 + \tau + t)]$ which is larger. Hence, stopping at x_n cannot be optimal. And by the inductive hypothesis after continuing to x_{n+1} the optimal strategy is the strategy given in the proposition.

Now suppose that $V(p_0) < \mathbb{E}_t[g(t_0 + \tau + t) - g(t_0)]$. Consider the alternate model where search costs are fixed at $c \equiv \mathbb{E}_t[g(t_0 + \tau + t) - g(t_0)]$. In the alternate model, it is well known that in any optimal strategy the consumer stops at x_n .²⁸ But, relative to expected continuation payoffs conditional on reaching x_n in the alternate model, expected continuation payoffs conditional on reaching x_n in the original model are the same for the strategy that stops at x_n and are lower for any strategy that continues at x_n . So in any best response in our model the consumer stops at x_n as well.

Therefore, any strategy of the desired form yields a strictly higher expected continuation payoff than any strategy not of this form when there are $m + 1$ remaining unsearched firms, so the result for $m = N$ follows. *Q.E.D.*

Proof of Lemma 2. If prices were ever negative in a nontrivial SE, then a firm would have a profitable deviation by replacing the negative prices in the price distribution with zero price. So in any nontrivial SE all prices are weakly positive. Therefore, if the first price a costly searcher observes is p_0 , then in SE her benefit from searching again if every firm sets $t_i = 0$ is $\int_{p_0}^{\infty} D(x)F(x) dx - (g(2\tau) - g(\tau)) \leq \int_0^{p_0} D(x) dx - (g(2\tau) - g(\tau))$, which is negative for p_0 sufficiently close to 0. By convexity of g , her benefit from searching again is no greater than this if any firm sets positive obfuscation.²⁹ Therefore, in SE any firm can guarantee itself positive profits by choosing such a sufficiently small but strictly positive p_0 , so every firm must make positive profits in any SE. *Q.E.D.*

Proof of Lemma 3. By Lemma 2, no firm sets $p = 0$ in any SE in which costly searchers enter. So if $F(p)$ has an atom, it must have an atom at some $p > 0$. But then pricing slightly below this atom yields strictly higher profits than pricing at the atom, as it yields a discrete gain in profits from the costless searchers and an arbitrarily small loss in profits from the costly searchers. *Q.E.D.*

Proof of Lemma 4. Let $F(p)$ be a nontrivial SE price distribution for a model with costless obfuscation. Let \bar{p} be the maximum of the support of $F(p)$. Consider a firm that sets price equal to \bar{p} . If this firm does not sell to any of the costly searchers that visit it first, then with probability 1 it will not sell to any consumers as, by Lemma 3, every other firm has a lower price with probability 1 and consumers buy from the lowest-priced firm they visit. This would contradict Lemma 2, so a firm that sets price equal to \bar{p} must sell to some costly searchers that visit it first. Furthermore, if consumers mix between buying and not buying from a firm with price equal to \bar{p} , then by lowering prices by an arbitrarily small amount the firm could sell to these consumers with probability 1, by Lemma 1 strictly increasing profits. So if $F(p)$ is a SE price distribution then every costly searcher who visits a firm with price equal to \bar{p} first buys immediately.

Since those costly searchers who first visit a highest-priced firm buy from it, any lower-priced firm could sell to those costly searchers who visit it first by setting the same obfuscation level as the highest-priced firm. And raising one's obfuscation level only increases the number of consumers one sells to, by Lemma 1, so if a lower-priced firm did not sell to those costly searchers who visited it first it could strictly increase profits by raising its obfuscation level to that of the highest-priced firm. *Q.E.D.*

²⁷ Anderson and Renault (2006) build on Wolinsky's framework to examine the related topic of the information content of advertising.

²⁸ See Kohn and Shavell (1974) or Weitzman (1979).

²⁹ For this result, unlike many of our main results, it suffices that g is weakly convex (i.e., that obfuscation does not decrease future search costs).

Proof of Proposition 2. We first show that any such joint distribution over p and t is a SE price-obfuscation distribution. By Lemma 1, all consumers search at least once if (2) holds, and all consumers search at most once if (3) holds. A firm that chooses price p^m obtains profit $\frac{1-\mu}{N}R(p^m)$, while a firm that chooses price $p < p^m$ obtains profit $[\mu(1 - F(p))]^{N-1} + \frac{1-\mu}{N}R(p)$. It is easy to check that these are equal for all $p \in [\underline{p}, p^m]$ when $F(p)$ is given by equation (1), while profits associated with a deviation to any p_i outside this interval are strictly smaller, regardless of the chosen obfuscation level t_i . So any such distribution over p and t is a SE price-obfuscation distribution.

Next, we show that the set of these joint distributions is nonempty if and only if $g(\tau) \leq V(\infty)$. If firms choose p according to $F(p)$ and all choose the same obfuscation level, t , equations (2) and (3) become

$$g(\tau + t) \leq \int_{\underline{p}}^{\infty} D(p)F(p) dp, \tag{A1}$$

$$g(2(\tau + t)) - g(\tau + t) \geq \int_{\underline{p}}^{p^m} D(p)F(p) dp, \tag{A2}$$

Since $g(\cdot)$ is strictly increasing and convex with $g(0) = 0$, we have $g(2(\tau + t)) - g(\tau + t) \geq g(\tau + t)$ and $\lim_{t \rightarrow \infty} g(2(\tau + t)) - g(\tau + t) = \infty$. If $g(\tau) \leq V(\infty)$, continuity of g implies that there exist $t \in \mathbb{R}^+$ that satisfy both (A1) and (A2). If $g(\tau) > V(\infty)$, then (A1) does not hold for any marginal distribution over t .

We next note that the marginal distribution over t must satisfy (2) and (3) in any nontrivial SE in which the marginal distribution over p equals $F(p)$. This is immediate: if (2) does not hold, costly searchers will not search once, and if (3) does not hold, some consumers will search twice, by Lemma 1, which cannot occur in a nontrivial SE by Lemma 4.

Finally, we come to the main part of the proof: showing that equation (1) defines the only possible nontrivial SE price distribution. Suppose that costly searchers enter and that $F(p)$ is not given by equation (1). We consider two cases:

First, suppose that $\bar{p} \neq p^m$ is such that $F(p) = 1 - [(\frac{1-\mu}{N\mu})(\frac{R(\bar{p})}{R(p)} - 1)]^{\frac{1}{N-1}}$ for all p in the support of F . If $\bar{p} > p^m$, a firm could deviate to p^m and make strictly higher profits from both costless and costly searchers. Suppose $\bar{p} < p^m$. By Lemma 4, if a firm plays (t_i, \bar{p}) in SE then $\mathbb{E}_i[g(2\tau + t_i + t) - g(\tau + t_i)] \geq \int_{\underline{p}}^{\bar{p}} D(p)F(p) dp$. So, by strict convexity of g , there exist $\varepsilon, \varepsilon' > 0$ such that $\mathbb{E}_i[g(2\tau + t_i + \varepsilon + t) - g(\tau + t_i + \varepsilon)] > \int_{\underline{p}}^{\bar{p} + \varepsilon'} D(p)F(p) dp$, so if such a firm deviated to playing $(t_i + \varepsilon, \bar{p} + \varepsilon')$ then a consumer will prefer to buy at price $\bar{p} + \varepsilon'$ rather than searching again when the firm's obfuscation level is $t_i + \varepsilon$. This deviation gives the firm strictly higher profits from the costly searchers and makes no difference to its profits from the costless searchers, since at price \bar{p} it had zero probability of selling to these consumers (by Lemma 3) and still has zero probability of selling to them at price $\bar{p} + \varepsilon$.

Next, suppose that there exists p in the support of F such that $F(p) \neq 1 - [(\frac{1-\mu}{N\mu})(\frac{R(\bar{p})}{R(p)} - 1)]^{\frac{1}{N-1}}$. Then it is easy to check that profits at p do not equal profits at \bar{p} , contradicting that p and \bar{p} are both in the support of F . *Q.E.D.*

Proof of Proposition 4. Note that consumer welfare u is given by $u \equiv V(\infty) - \mathbb{E}[g(t + \tau)]$. Suppose that $\tau' \geq \tau$ and $u' \geq u$, where u is a nontrivial SE value of consumer welfare with fixed search cost τ and u' is a nontrivial SE value of consumer welfare with fixed search cost τ' . We must show that u' is a nontrivial SE value with search cost τ and that u is a nontrivial SE value with search cost τ' .³⁰

First, suppose that u' is the value of consumer welfare for a nontrivial SE with price distribution $F(p)$ and obfuscation strategies given as a function of price $t'(p)$. Note that $t'(p)$ can be a probability distribution over obfuscation levels, if firms mix over obfuscation levels given their prices. Consider the profile where firms price according to $F(p)$ and use obfuscation strategies $t(p) = t'(p) + \tau' - \tau \geq t'(p) \geq 0$, where if $t'(p)$ is a probability distribution over obfuscation levels this is interpreted as shifting this distribution up by $\tau' - \tau$. It is clear that this profile is a nontrivial SE when fixed search costs are given by τ , because at every history a consumer's expected future total search cost when the fixed component is given by τ and the variable component is given by $t(p)$ is the same as when the fixed component is given by τ' and the variable component is given by $t'(p)$. And consumer welfare in this SE is u' .

Next, suppose again that u, u' and $t'(p)$ are as above. To show that u is also a nontrivial SE value when the exogenous search cost is equal to τ' , we suppose that firms draw prices from $F(p)$ and obfuscate according to $t(p) = t'(p) + \delta$, where δ is such that $\mathbb{E}[g(t'(p) + \tau + \delta) - g(t'(p) + \tau)] = u' - u$. Note first that $t(p) \geq t'(p)$ implies that every costly searcher searches at most once, because compared to the original SE, search costs have increased while search benefits remain constant. Second, the fact that the equilibrium utility of each type of consumer is identical to the utility that the same consumer gets in the nontrivial SE with utility u in the game with exogenous search costs τ implies that costly searchers are willing to enter. Hence, this profile is a nontrivial SE with payoff u in the game with exogenous search costs τ' . *Q.E.D.*

³⁰ To see that this is what must be shown let X be the set of equilibrium values of consumer welfare given search cost τ and let Y be the set of equilibrium values of consumer welfare given τ' . X is higher than Y if the larger of u and u' is in X and the smaller is in Y whenever $u \in X$ and $u' \in Y$. For $u' < u$ this is trivially true, so what remains is to show it is also true when $u' \geq u$.

Proof of Proposition 5. By Proposition 2, $F(p) = 1 - [(\frac{1-\mu}{N\mu})(\frac{R(p^m)}{R(p)} - 1)]^{\frac{1}{N-1}}$ and $R(\underline{p}) = (\frac{1-\mu}{1+(N-1)\mu})R(p^m)$ in any pure-strategy nontrivial SE, so $F(p)$ is increasing in μ for all p and \underline{p} is decreasing in μ . The lower bound on t^u is given by $g(2(\tau + t^u)) - g(\tau + t^u) = \int_{\underline{p}}^{p^m} D(p)F(p) dp$, so it is increasing in $F(p)$ and therefore increasing in μ . Similarly, the upper bound on t^u is given by $g(\tau + t^u) = \int_{\underline{p}}^{\infty} D(p)F(p) dp$, so it is increasing in $F(p)$ and therefore in μ as well. Q.E.D.

Proof of Proposition 6. We claim that there exists $\bar{N} > 0$ such that $\int_{\underline{p}}^{p^m} D(p)F(p) dp$ is decreasing in N for $N > \bar{N}$. By Proposition 2, $F(p) = 1 - [(\frac{1-\mu}{N\mu})(\frac{R(p^m)}{R(p)} - 1)]^{\frac{1}{N-1}}$ and $R(\underline{p}) = (\frac{1-\mu}{1+(N-1)\mu})R(p^m)$. Treating N as a continuous variable, we have

$$\begin{aligned} \frac{\partial}{\partial N} \int_{\underline{p}}^{p^m} D(p)F(p) dp &= \int_{\underline{p}}^{p^m} D(p) \frac{\partial F(p)}{\partial N} dp - \frac{\partial \underline{p}}{\partial N} D(\underline{p})F(\underline{p}) \\ &= \int_{\underline{p}}^{p^m} D(p) \frac{\partial F(p)}{\partial N} dp. \end{aligned}$$

The derivative of $F(p)$ with respect to N is

$$\frac{1}{N(N-1)} \left[\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right]^{\frac{1}{N-1}} \left(1 + \frac{N}{N-1} \log \left(\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right) \right).$$

Therefore,

$$\begin{aligned} \text{sign} \left(\frac{\partial}{\partial N} \int_{\underline{p}}^{p^m} D(p)F(p) dp \right) &= \text{sign} \left(\int_{\underline{p}}^{p^m} D(p) \left[\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right]^{\frac{1}{N-1}} \left(1 + \frac{N}{N-1} \log \left(\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right) \right) dp \right). \tag{A3} \end{aligned}$$

For any p and N ,

$$\begin{aligned} &\left[\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right]^{\frac{1}{N-1}} \left(1 + \frac{N}{N-1} \log \left(\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right) \right) \\ &\leq \left[\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(\underline{p})} - 1 \right) \right]^{\frac{1}{N-1}} \left(1 + \frac{N}{N-1} \log \left(\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(\underline{p})} - 1 \right) \right) \right) \\ &= \left[\left(\frac{1-\mu}{N\mu} \right) \left(\frac{1+(N-1)\mu}{1-\mu} - 1 \right) \right]^{\frac{1}{N-1}} \left(1 + \frac{N}{N-1} \log \left(\left(\frac{1-\mu}{N\mu} \right) \left(\frac{1+(N-1)\mu}{1-\mu} - 1 \right) \right) \right) \\ &= 1. \end{aligned}$$

For any fixed $p \in (0, p^m)$,

$$\begin{aligned} &\lim_{N \rightarrow \infty} \left[\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right]^{\frac{1}{N-1}} \left(1 + \frac{N}{N-1} \log \left(\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right) \right) \\ &= \lim_{N \rightarrow \infty} 1 + \frac{N}{N-1} \log \left(\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right) \\ &= -\infty. \end{aligned}$$

Therefore, since $\underline{p} > 0$ for all N ,

$$\begin{aligned} &\lim_{N \rightarrow \infty} \int_{\underline{p}}^{p^m} D(p) \left[\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right]^{\frac{1}{N-1}} \left(1 + \frac{N}{N-1} \log \left(\left(\frac{1-\mu}{N\mu} \right) \left(\frac{R(p^m)}{R(p)} - 1 \right) \right) \right) dp \\ &= -\infty. \end{aligned}$$

By (A3), this implies that there exists $\bar{N} > 0$ such that $\int_{\underline{p}}^{p^m} D(p)F(p) dp$ is decreasing in N if $N > \bar{N}$.

Recall that the lower bound on t^u is given by $g(2(\tau + t^u)) - g(\tau + t^u) = \int_{\underline{p}}^{p^m} D(p)F(p) dp$, and the upper bound on t^u is given by $g(\tau + t^u) = \int_{\underline{p}}^{\infty} D(p)F(p) dp$. Both of these bounds are increasing in $\int_{\underline{p}}^{p^m} D(p)F(p) dp$, and are therefore decreasing in N if $N > \bar{N}$. Q.E.D.

Proof of Proposition 7. For part (a) note that the conditions on g imply that consumers will search once and will not search a second time if firms mix over prices as they do in the equilibrium of the costless obfuscation model. Hence, the same calculations as for the costless obfuscation model apply and show that there is no profitable deviation that involves zero obfuscation. Deviations that involve positive obfuscation are also not profitable, because they cannot lead to making greater sales (for any price weakly less than p^m) and any obfuscation costs incurred will only reduce profits.

For part (b), note first that Lemmas 2 and 3 imply that prices are distributed according to some atomless distribution F^* on some interval $[\underline{p}, \bar{p}]$ with $\underline{p} > c$. We must have $\bar{p} \leq p^m$ because any price greater than $p > p^m$ is dominated by setting $p = p^m$ and using the same obfuscation level. We claim that firms setting prices p_0 in a neighborhood of \bar{p} must sell to all consumers who visit them first. To see this, first note that if they did not sell to any consumers who visit them first then they earn arbitrarily small profits (which is impossible, since all firms must earn the same, strictly positive level of profits in any SE), as consumers purchase only from the lowest-price firm they visit, and firms with prices close to \bar{p} have a vanishingly small probability of being the lowest-price firm visited by any consumer who has searched more than once. And if such a firm sells to only some of those consumers that visit it first, it can cut prices by an arbitrarily small amount and sell to all of these consumers, yielding a discrete gain in profits. Such firms must also be choosing obfuscation levels satisfying

$$\begin{aligned} \mathbb{E}_t[g(2\tau + t + t(p_0))] - g(\tau + t(p_0)) &= \int_{\underline{p}}^{p_0} D(p)F^*(p) dp \quad \text{if } t(p_0) > 0, \\ \mathbb{E}_t[g(2\tau + t)] - g(\tau) &\geq \int_{\underline{p}}^{p_0} D(p)F^*(p) dp \quad \text{if } t(p_0) = 0, \end{aligned}$$

where $F^*(p)$ is the equilibrium price distribution. Therefore, by the assumption that $g(2\tau) - g(\tau) < V(p^m)$, if $\bar{p} = p^m$ then $t(p_0) > 0$ for all p_0 just below \bar{p} . Hence, if $\bar{p} = p^m$, then for all p_0 just below \bar{p} a first-order approximation to the profit function is

$$R(p_0) \frac{1 - \mu}{N} - c(t(p_0)).$$

This expression is strictly decreasing for p_0 near p^m , because $R'(p^m) = 0$ and the cost term has a nonzero derivative. Hence, \bar{p} cannot be equal to p^m . Q.E.D.

Proof of Proposition 9. We note first that it suffices to show that there is a δ such that firms always obfuscate to the extent necessary so that consumers do not conduct a second search when $\delta < \delta$. This suffices because the benefit of a second search is strictly increasing in p so the convexity of g implies that a larger t is needed to deter obfuscation when p is higher (and $c(t)$ strictly increasing implies firms choose the smallest t that deters a second search). Showing that no consumers will search twice takes a few steps but is not difficult. Let \underline{p} be the lower bound of the support of the equilibrium price distribution in a nontrivial SE. Strict convexity of g and $\tau > 0$ imply that there is a positive Δ such that consumers will never do a second search if $p \leq \underline{p} + \Delta$. Let \bar{t} be such that a consumer who sees price p^m and obfuscation \bar{t} would never conduct a second search even if he expects to find $p = 0$ on the next search; such a \bar{t} exists by the assumption that $\lim_{t \rightarrow \infty} g'(t) = \infty$ and the fact that $\tau > 0$. Because a firm can deviate to setting price p^m and doing obfuscation \bar{t} we can give uniform lower bound on equilibrium profits when δ is small. This also implies a uniform lower bound on \underline{p} which is strictly above cost. The fact that profits are at least $(\mu(1 - F(p))^{N-1} + (1 - \mu)/N)R(p)$ when prices are in $[\underline{p}, \underline{p} + \Delta]$ then gives a uniform lower bound on what $F(\underline{p} + \Delta)$ can be. If some consumers did search for a second time in a nontrivial SE, then this must occur when they see prices at some \hat{p} above $\underline{p} + \Delta$ (and all such consumers must conduct a second search, as otherwise a firm would rather price slightly below \hat{p} than at \hat{p}). This is impossible, however, for δ small, because deviating to set price \hat{p} and do obfuscation \bar{t} would be more profitable: the cost of this obfuscation $\delta C(\bar{t})$ goes to zero as $\delta \rightarrow 0$, whereas the benefit is bounded below by $\frac{1-\mu}{N}F(\underline{p} + \Delta)R(\underline{p})$ and we have noted that all of these terms are uniformly bounded away from 0 when δ is small. Q.E.D.

Proof of Proposition 10. The proof for the $\bar{p} = p^m$ case is similar to the $\bar{p} < p^*$ case and is omitted; it may also be found in the working paper version.

The first part of the proposition is the usual condition for firms to be indifferent between charging any two prices in $[\underline{p}^*, p^*]$, as in Stahl (1989), for example.

For the second part, first note that the conditions $\tau + \underline{\theta} \leq \int_{\underline{p}^*}^{p^*} D(p)F_{p^*}(p) dp$ and $-\underline{\theta} \leq \int_{\underline{p}^*}^{\infty} D(p) dp$ hold if and only if there exists a $\bar{t} \geq 0$ such that $\tau + \bar{t} \leq \int_{\underline{p}^*}^{\infty} D(p)F_{p^*}(p) dp$ and $\tau + \underline{\theta} + \bar{t} = \int_{\underline{p}^*}^{p^*} D(p)F_{p^*}(p) dp$. To see this, note that if the first pair of conditions hold, then taking $\bar{t} = \int_{\underline{p}^*}^{p^*} D(p)F_{p^*}(p) dp - (\tau + \underline{\theta})$ gives a value for which the desired inequality and equality both hold using that $\int_{\underline{p}^*}^{p^*} D(p)F_{p^*}(p) dp + \int_{p^*}^{\infty} D(p) dp = \int_{\underline{p}^*}^{\infty} D(p)F_{p^*}(p) dp$. And conversely and if there exists a $\bar{t} \geq 0$ with the two desired properties then $\tau + \underline{\theta} \leq \int_{\underline{p}^*}^{p^*} D(p)F_{p^*}(p) dp$ is immediate and $-\underline{\theta} \leq \int_{p^*}^{\infty} D(p) dp$ follows by subtracting the equality that holds for \bar{t} from the inequality also assumed to hold for \bar{t} .

Now if there is no $\bar{t} \geq 0$ such that $\tau + \bar{t} \leq \int_{\underline{p}^*}^{\infty} D(p)F_{p^*}(p) dp$ and $\tau + \underline{\theta} + \bar{t} = \int_{\underline{p}^*}^{p^*} D(p)F_{p^*}(p) dp$, there can be no SE of the desired form. To see this, first note that costly searchers will not enter if the first inequality is violated and

costly searchers will search for a second time if they observe price p^* and face average obfuscation \bar{t} if the right-hand side of the equality is strictly greater than the left-hand side, which is impossible in SE by Lemma 5. And if the right-hand side of the equality is strictly less than the left-hand side, then a firm would be able to profitably deviate to setting price slightly above $p > p^*$, as costly searchers would still buy at such a price with probability one and such a firm makes zero profit from costless searchers.

If there exists a $\bar{t} \geq 0$ such that $\tau + \bar{t} \leq \int_{p^*}^{\infty} D(p)F_{p^*}(p) dp$ and $\tau + \underline{\theta} + \bar{t} = \int_{p^*}^{p^*} D(p)F_{p^*}(p) dp$, consider the strategy profile where all firms set obfuscation level equal to $\int_{p^*}^{p^*} D(p)F_{p^*}(p) dp - (\tau + \underline{\theta}) \geq 0$ and randomize their prices according to $F_{p^*}(p)$, and suppose that consumers search optimally and have the off-equilibrium path belief that $\theta = \underline{\theta}$ if they ever observe $p \notin [p^*, p^*]$. Under these strategies, firms sell to all consumers who visit them first and are indifferent among all prices in $[p^*, p^*]$, so the only deviation that could possibly be profitable would be to a price greater than p^* . But a consumer that observes $\hat{p} > p^*$ expects to face search cost $\int_{p^*}^{p^*} D(p)F_{p^*}(p) dp - (\tau + \underline{\theta}) + \tau + \underline{\theta} = \int_{p^*}^{p^*} D(p)F_{p^*}(p) dp$ from searching again and to receive expected benefit $\int_{p^*}^{\hat{p}} D(p)F_{p^*}(p) dp > \int_{p^*}^{p^*} D(p)F_{p^*}(p) dp$ from doing so, so a firm that deviated to such a price would not sell to consumers. So this is a SE.

In addition, it is immediate that the set of nontrivial SE with supremum of $F_{p^*}(p)$ given by p^* must be as stated in the proposition, as in the proof of Proposition 2.

Finally, suppose that $\tau \leq \int_{p^m}^{\infty} D(p)F_{p^m}(p) dp$. If $-\underline{\theta} \leq \int_{p^m}^{\infty} D(p) dp$, then a nontrivial SE with $p^* = p^m$ exists. So suppose that $-\underline{\theta} > \int_{p^m}^{\infty} D(p) dp$. Then $0 < \tau + \underline{\theta} < \int_{p^m}^{p^m} D(p)F_{p^m}(p) dp$. Note that $\int_{p^*}^{p^*} D(p)F_{p^*}(p) dp$ equals 0 if p^* equals 0, equals $\int_{p^m}^{p^m} D(p)F_{p^m}(p) dp$ if p^* equals p^m , and is continuous in p^* . Therefore, the Intermediate Value Theorem implies that there exists a p^* such that $\tau + \underline{\theta} = \int_{p^*}^{p^*} D(p)F_{p^*}(p) dp$. Therefore,

$$\begin{aligned} -\underline{\theta} &= \tau - \int_{p^*}^{p^*} D(p)F_{p^*}(p) dp \\ &\leq \int_{p^m}^{\infty} D(p)F_{p^m}(p) dp - \int_{p^*}^{p^*} D(p)F_{p^*}(p) dp \\ &\leq \int_{p^m}^{\infty} D(p) dp \\ &\leq \int_{p^*}^{\infty} D(p) dp, \end{aligned}$$

where the second inequality again uses the assumption that $\int_{p^*}^{p^*} D(p)F_{p^*}(p) dp$ is increasing in p^* for $p^* < p^m$. The characterization in the first part of the proposition then implies that a SE exists with $\bar{p} = p^*$. Q.E.D.

Proof of Proposition 12. Suppose that $\underline{\theta}' \geq \underline{\theta}$ and $u' \leq u$, where u is a SE value of consumer welfare with lower bound on θ given by $\underline{\theta}$ and u' is a SE value of consumer welfare with this lower bound given by $\underline{\theta}'$. Denote the upper bound of the price distribution yielding consumer welfare u by \bar{p} and denote the corresponding upper bound for u' by \bar{p}' . We must show that u' is a SE value when the lower bound is given by $\underline{\theta}$ and that u is a SE value when this bound is given by $\underline{\theta}'$.

We have that \bar{p} and \bar{p}' are both less than p^m , so the proof of Proposition 10 gives that $\int_{\bar{p}}^{\bar{p}} = \tau + \underline{\theta} + \bar{t}$ and $\int_{\bar{p}'}^{\bar{p}'} = \tau + \underline{\theta}' + \bar{t}'$, where \bar{t} and \bar{t}' are average obfuscation levels corresponding to SE with price upper bound \bar{p} and welfare u , and price upper bound \bar{p}' and welfare u' , respectively. Recall that $u = \int_{\bar{p}}^{\infty} F_{\bar{p}}(p)D(p) dp - \tau - \bar{t}$, so we have $u = \int_{\bar{p}}^{\infty} D(p) dp + \underline{\theta}$ and $u' = \int_{\bar{p}'}^{\infty} D(p) dp + \underline{\theta}'$. Since $u' \leq u$, this implies that $\bar{p}' \geq \bar{p}$. Therefore, we have

$$u' = \int_{\bar{p}'}^{\infty} D(p) dp + \underline{\theta}' \leq u \leq \int_{\bar{p}}^{\infty} D(p) dp + \underline{\theta}'.$$

The Intermediate Value Theorem then implies that there exists $p^* \in [\bar{p}, \bar{p}']$ such that $u = \int_{p^*}^{\infty} D(p) dp + \underline{\theta}'$, which then implies that u is a SE value of consumer welfare when the lower bound on θ is given by $\underline{\theta}'$ and the upper bound on p is given by p^* . The argument for u' is similar. Q.E.D.

References

ANDERSON, S. AND RENAULT, R. "Advertising Content." *American Economic Review*, Vol. 96 (2006), pp. 93–113.
 BAYE, M. AND MORGAN, J. "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets." *American Economic Review*, Vol. 91 (2001), pp. 454–474.
 ———, MORGAN, J., AND SCHOLTEN, P. "Information, Search, and Price Dispersion." In T. Hendershott, ed., *Handbook of Economics and Information Systems*. Amsterdam: Elsevier, 2006.
 BURDETT, K. AND JUDD, K.L. "Equilibrium Price Dispersion." *Econometrica*, Vol. 51 (1983), pp. 955–969.

- BUTTERS, G. "Equilibrium Distributions of Sales and Advertising Prices." *Review of Economic Studies*, Vol. 44 (1977), pp. 465–491.
- CARLIN, B.I. "Strategic Price Complexity in Retail Financial Markets." *Journal of Financial Economics*, Vol. 91 (2009), pp. 278–287.
- DE LOS SANTOS, B., HORTACSU, A., AND WILDENBEEST, M.R. "Testing Models of Consumer Search Behavior Using Data on Web Browsing and Consumer Purchases." *American Economic Review*, forthcoming.
- DIAMOND, P. "A Model of Price Adjustment." *Journal of Economic Theory*, Vol. 3 (1971), pp. 156–168.
- ELIAZ, K. AND SPIEGLER, R. "Consideration Sets and Competitive Marketing." *Review of Economic Studies*, Vol. 78 (2011), pp. 235–262.
- ELLISON, G. "A Model of Add-on Pricing." *Quarterly Journal of Economics*, Vol. 120 (2005), pp. 585–637.
- . "Bounded Rationality in Industrial Organization." In R. Blundell, W. Newey, and T. Persson, eds., *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*. Cambridge, UK: Cambridge University Press, 2006.
- ELLISON, G. AND ELLISON, S.F. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica*, Vol. 77 (2009), pp. 427–452.
- AND WOLITZKY, A. "A Search Cost Model of Obfuscation." Working Paper no. 15237, NBER, 2009.
- FUDENBERG, D. AND TIROLE, J. "A 'Signal-Jamming' Theory of Predation." *Rand Journal of Economics*, Vol. 17 (1986), pp. 366–376.
- GABAIX, X. AND LAIBSON, D. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics*, Vol. 121 (2006), pp. 505–540.
- GROSSMAN, S.J. "The Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics*, Vol. 24 (1981), pp. 461–483.
- HENDRICKS, K., SORENSEN, A., AND WISEMAN, T. "Observational Learning and Demand for Search Goods." *American Economic Journal: Microeconomics*, forthcoming.
- HOLMSTRÖM, B. "Managerial Incentive Problems: A Dynamic Perspective." In *Essays in Economics and Management in Honor of Lars Wahlbeck*. Helsinki: Swedish School of Economics, 1982. Reprinted in *Review of Economic Studies*, Vol. 66 (1999), pp. 169–182.
- JANSSEN, M.C.W., MORAGA-GONZALEZ, J.L., AND WILDENBEEST, M.R. "Truly Costly Sequential Search and Oligopolistic Pricing." *International Journal of Industrial Organization*, Vol. 23 (2005), pp. 451–466.
- KOHN, M. AND SHAVELL, S. "The Theory of Search." *Journal of Economic Theory*, Vol. 9 (1974), pp. 93–123.
- MILGROM, P. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics*, Vol. 12 (1981), pp. 380–391.
- AND SHANNON, C. "Monotone Comparative Statics." *Econometrica*, Vol. 62 (1994), pp. 157–180.
- ROBERT, J. AND STAHL, D.O. "Informative Price Advertising in a Sequential Search Model." *Econometrica*, Vol. 61 (1993), pp. 657–686.
- ROSENTHAL, R.W. "A Model in Which an Increase in the Number of Sellers Leads to a Higher Price." *Econometrica*, Vol. 48 (1980), pp. 1575–1579.
- SALOP, S. AND STIGLITZ, J. "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion." *Review of Economics Studies*, Vol. 49 (1977), pp. 3–13.
- SPIEGLER, R. "Competition over Agents with Boundedly Rational Expectations." *Theoretical Economics*, Vol. 1 (2006), pp. 207–231.
- STAHL, D.O. "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review*, Vol. 79 (1989), pp. 700–712.
- . "Oligopolistic Pricing with Heterogeneous Consumers Search." *International Journal of Industrial Organization*, Vol. 14 (1996), pp. 243–268.
- STIGLITZ, J.E. "Competition and the Number of Firms in a Market: Are Duopolies More Competitive than Atomistic Markets?" *Journal of Political Economy*, Vol. 95 (1987), pp. 1041–1061.
- VARIAN, H. "A Model of Sales." *American Economic Review*, Vol. 70 (1980), pp. 651–659.
- WEITZMAN, M. "Optimal Search for the Best Alternative." *Econometrica*, Vol. 47 (1979), pp. 641–654.
- WHITE, A. "Search Engines: Left Side Quality versus Right Side Profits." Mimeo, Toulouse School of Economics, 2009.
- WILSON, C.M. "Ordered Search and Equilibrium Obfuscation." *International Journal of Industrial Organization*, Vol. 28 (2010), pp. 496–506.
- WOLINSKY, A. "True Monopolistic Competition as a Result of Imperfect Information." *Quarterly Journal of Economics*, Vol. 101 (1986), pp. 493–512.