Persuasion Produces the (Diamond) Paradox

Mark Whitmeyer*

July 19, 2021

Abstract

This paper extends the sequential search model of Wolinsky (1986) by allowing firms to choose how much match value information to disclose to visiting consumers. This restores the Diamond paradox (Diamond (1971)): there exist no symmetric equilibria in which consumers engage in active search, so consumers obtain zero surplus and firms obtain monopoly profits. Modifying the scenario to one in which prices are advertised, we discover that the no-active-search result persists, although the resulting symmetric equilibria are ones in which firms price at marginal cost.

Keywords: Diamond Paradox, Sequential Search, Information Design, Bayesian Persuasion

JEL Classifications: C72; D82; D83

*Arizona State University

Email: mark.whitmeyer@gmail.com.

I am grateful to an anonymous referee for a different paper for the comment that sparked this idea. I also thank Özlem Bedre-Defolie, Vasudha Jain, Stephan Lauermann, Benny Moldovanu, Joseph Whitmeyer, and Thomas Wiseman for helpful comments and feedback. I wrote much of this paper while I was at the University of Bonn and acknowledge their generous support.

I'm gonna stop, uh, shoppin' around.

Elvis Presley (G.I. Blues)

1 Introduction

The Diamond paradox (Diamond (1971)) is a stark result that highlights the importance of search frictions in models of price competition. Infamously, Diamond establishes that even in a market with a large number of firms, an arbitrarily small (yet positive) search cost ensures that the firms behave like monopolists and that consumers do not search. The intuition behind this result is well-known: for any price strictly below the monopoly price, demand is locally perfectly inelastic and so a firm can always improve its lot by raising its price slightly.

As a number of subsequent papers illustrate, small modifications to the model can overturn the result. Two such works stand out in particular: Stahl (1989) assumes that a fraction of the market's consumers have zero search cost and therefore freely buy from whomever sets the lowest price. As a result, demand for each firm slopes down again and they behave like monopolists no longer. Wolinsky (1986) takes a different approach. In his model, although each consumer faces a positive search cost, the firms' products are differentiated and consumers have imperfect information about the products. Each consumer's match value at any firm is an i.i.d. random variable, about which the firms and consumer are *ex-ante* uninformed. Upon visiting a firm, a consumer discovers both the firm's price as well as the realized match value. This uncertainty begets equilibria with active search.

In this paper, we revisit the framework of Wolinsky (1986) but alter it by assuming that each firm may choose how much information to provide to a consumer during her visit, in addition to setting a price. Each firm commits to a signal that maps a consumer's match value to a (conditional) distribution over signal realizations. In the spirit of the original Diamond paper, neither the signal nor the price chosen by a firm is observable until a consumer incurs the search cost. That is, after paying the search cost to visit a firm, a consumer then observes that firm's signal, price, and draws a signal realization that yields her posterior belief about the firm's match value.

This modification has drastic consequences. We find that even in a model with product differentiation and imperfect information, information provision restores the Diamond paradox. Namely, the main result of this paper, Theorem 2.4, states that there are no symmetric equilibria with active consumer search. Each consumer visits at most one firm and purchases from it.[1] In all symmetric equilibria, firms leave consumers with zero rents.

The intuition to this result is similar to that in the original paper by Diamond. There, in any purported equilibrium with active search, firms can always exploit visiting consumers by raising their prices slightly, which does not affect their purchase decisions due to the positive search cost. Here, firms deviate by providing slightly less information and pooling beliefs above consumers' stopping thresholds. Subsequently, given that any equilibrium must involve no search, the classic Diamond paradox incentive kicks in and firms must obtain monopoly profits.

This incentive to pool beliefs is extremely strong, and continues to drive the results even when prices are posted and so can be observed before consumers embark on their searches. The second finding of this paper, Theorem 3.6, is that even when prices are posted–and therefore shape consumers' search behavior directly–there exist no symmetric equilibria with active search. Again, any supposed equilibrium in which there is active search would allow firms to deviate profitably by providing less information. In the unique symmetric equilibrium, because prices are posted, the usual Bertrand forces apply such that firms price at marginal cost and obtain no profits. Consumer surplus is merely the expected match value of the firm and no useful information is provided.

This paper belongs to the growing collection of papers that explore information design and persuasion in consumer search settings. This literature includes Board and Lu (2018), who also provide conditions for an analog of the Diamond paradox to arise. Their model is completely different, however: the uncertain state of the world is common and

[1] There is a trivial equilibrium in which consumers conjecture extremely high prices and do not visit any firms. As is convention, throughout this paper, we ignore this equilibrium.

so the competing sellers each provide information about the common state to prospective consumers. If sellers can perfectly observe a consumer's current belief about the state upon her visit, or coordinate their persuasion strategies, then there is an equilibrium in which they provide the monopoly level of information and set the monopoly price. This equilibrium is not generally unique and requires (mild) additional assumptions that pertain to the competitive persuasion problem with a common state. Furthermore, the sellers in Board and Lu (2018) do not set prices but compete through information alone.

In contrast, this paper is a direct adaptation of Wolinsky (1986), in which the distributions over match values are endogenous objects. Importantly, a consumer's match value at any firm is i.i.d., so a firm's choice of distribution does not affect a consumer's belief at any other firm. This paper shows that the paradox occurs when firms set prices as well, and even holds when prices are posted (advertised) and so search is directed. Due to the independence of match values, firms know a visiting consumer's prior about their quality, but not her outside option, which is endogenously determined by her search behavior. Thus, although a firm may not know precisely where it is in a consumer's search order, it knows the upper bound of a consumer's outside option and can tailor its information accordingly.

Three other papers similar to this one are Au and Whitmeyer (2020a), He and Li (2021), and Au and Whitmeyer (2020b). The first explores a game of pure information provision (without prices) between sellers in Weitzman (1979)'s classic sequential search setting. There, if consumers must incur a search cost to discover a firm's match value distribution, firms do not provide any (useful) information and consumers do not search actively, as in this paper. He and Li (2021) investigate a similar problem in which consumers observe the firms' match values for free but may not direct their search. They find that, despite the relative transparency of the market, firms (again) provide no (useful) information. Au and Whitmeyer (2020b), in turn, modifies Au and Whitmeyer (2020a) by allowing firms to compete both through prices and (advertised) information.

Another relevant paper is Dogan and Hu (2018), who look at consumer-optimal information structures in Wolinsky's model of sequential search. Related to that is recent work by Zhou (2020), who looks broadly at how improved information affects consumer welfare in Wolinsky's setting. Armstrong and Zhou (2019), in turn, focus on consumer-optimal information structures in a (frictionless) duopoly market. Like this paper, Hwang et al. (2018) also look at competitive information design and price setting in an oligopoly market, albeit one without search frictions. In their framework, absent such frictions, they show that firms provide (some) information to consumers and that both consumers and firms obtain some surplus in the market.

Finally, this paper relates naturally to the vein of research that focuses on obfuscation. One particularly relevant paper is Ellison and Wolitzky (2012), who modify the model of Stahl (1989) by allowing firms to choose the length of time necessary for consumers to learn their prices. Like firms' information policies in this paper, the required time to learn the price at a firm is hidden from consumers until their visits. A pair of important distinctions between their paper and this one; however, are that products in their model are undifferentiated and the level of obfuscation chosen by a firm affects consumers by altering their *future* search costs.

In contrast, here, firms choose how much information about their products consumers acquire during their visits, and therefore shape consumers' valuations directly. Thus, any alteration by a firm to its information policy may come at a cost, since it could potential limit the rents the firm could extract. Moreover, prices in their model are hidden (given the market's homogeneity, a posted price version of their model is just the canonical Bertrand setup), whereas we find that the same economic intuition drives this paper's results regardless of whether prices are posted. Firms could (costlessly) provide useful information, yet do not do so in any symmetric equilibrium.

2 Model

The foundation for this paper is the workhorse sequential search model of Wolinsky (1986). There are n symmetric firms and a unit mass of consumers with unit demand. The match value of a consumer at firm i is an i.i.d. random variable, X_i, distributed according to (Borel) cdf G on $[0,1]$. Let μ denote the expectation of each X_i. Given a realized match value of x_i and price p_i, a consumer's utility from purchasing from firm i is

$$u(x_i, p_i) = x_i - p_i \ .$$

In contrast to the original model of Wolinsky (1986), a consumer does not directly observe her match value at firm i, but instead observes a signal realization that is correlated with it. Each firm has a compact metric space of signal realizations, S, and commits to a signal, Borel measurable function $\pi_i : [0,1] \to \Delta(S)$. As is well known, each signal realization begets–via Bayes' law–a posterior distribution over values, and thus a signal begets a (Bayes-plausible) distribution over posteriors. Alternatively, a signal begets a distribution over posterior means, and the following remark is now standard in the literature:

Remark 2.1. Each firm's choice of signal, π_i, is equivalent to a choice of distribution $F_i \in \mathcal{M}(G)$, where $\mathcal{M}(G)$ is the set of all mean-preserving contractions of G.

Thus, each firm chooses a distribution over values, F_i, and sets a price p_i. Importantly, a consumer only observes these choices upon her visit to a firm. Following Wolinsky (1986), search is sequential and with recall. At a cost of $c > 0$, a consumer may visit a firm and observe its price and realized draw from distribution F_i. As does Diamond (1971), we assume that a consumer incurs no search cost for her first visit. For simplicity, we impose that the marginal cost for each firm is 0.[2]

We look for symmetric equilibria in which each firm sets the same price p and chooses the same distribution over values F.[3] Our equilibrium concept is (weak) Perfect Bayesian Equilibrium (PBE), in which consumers have the same beliefs about the (theretofore unobserved) prices and information policies (signals) chosen by unvisited firms on and off the equilibrium path. This is customary in the literature.[4] For simplicity, a consumer has

[2]This assumption is relatively unimportant in the main specification and in the alternative setting in which prices are posted. As we discover in Section 3.1 it may be significant when consumers are heterogeneous.

[3]Wolinsky (1986) focuses on symmetric pure strategy pricing equilibria, which emphasis is echoed in this paper. Nevertheless, the main results of this paper–Theorems 2.4 and 3.6–hold for all symmetric equilibria, in both pure and mixed strategies.

[4]This assumption; however, is not innocuous. See Janssen and Shelegia (2020) for an in depth exploration of the ramifications of this assumption. As they note there–and in an earlier paper, Janssen and Shelegia (2015)–this stipulation may be more objectionable in markets with vertical relations.

an outside option of 0.

It is a standard result that a consumer follows a reservation price policy. Specifically, a consumer's search problem is a special case of the one explored in Weitzman (1979). Define the reservation value, z, induced by the conjectured price $\tilde{p}$ and distribution $\tilde{F}$ by

$$c = \int_{z+\tilde{p}}^{1} (x - \tilde{p} - z)\, d\tilde{F}(x)\ . \tag{$\star$}$$

In a symmetric equilibrium, a consumer's optimal search protocol is to visit firms in random order and stop and purchase from firm i if and only if the realized value at that firm, x_i, satisfies $x_i - p \geq z$, where p is the actual price set by firm i. If $x_i < z + p$ for all i then a consumer selects the firm whose realized value x_i is highest. At equilibrium, the conjectured price, $\tilde{p}$, must equal the actual price set by each firm, p; and the conjectured distribution, $\tilde{F}$, must equal the actual distribution chosen by each firm, F.

Note that we are using Weitzman (1979)'s formulation of a consumer's stopping problem, which is slightly nonstandard but equivalent to that of Wolinsky (1986). This allows for an easy transition into the next section, wherein we allow for advertised (posted) prices and heterogeneous consumers.

Now, let us establish the main result. The first step is to derive the following lemma.

Lemma 2.2. *There are no symmetric equilibria in which consumers visit more than one firm.*

Proof. Suppose for the sake of contradiction that there is such an equilibrium. For expositional ease, let us begin by assuming that the firms' choices of F and p are deterministic. Because $c > 0$ and by the definition of the reservation value, the conjectured distribution, $\tilde{F}$, must be such that both values strictly below and strictly above $z + \tilde{p}$ occur with strictly positive probability. Accordingly, let $[\underline{a}, \bar{a}]$ and $\left[\underline{b}, \bar{b}\right]$ be intervals such that $\int_{\underline{b}}^{\bar{b}} d\tilde{F}(x) > 0$ and $\int_{\underline{a}}^{\bar{a}} d\tilde{F}(x) > 0$, where $\underline{a} \leq \bar{a} < z + \tilde{p}$ and $z + \tilde{p} < \underline{b} \leq \bar{b}$.

Given price $\tilde{p} > 0$,[5] a firm's payoff from any value, x, that is weakly greater than $z + \tilde{p}$ is $\tilde{p}$. Moreover, its payoff from any value, x, that is strictly lower than $z + \tilde{p}$ is strictly less than $\tilde{p}$. Let $\alpha < \tilde{p}$ denote its average payoff (under $\tilde{F}$) for values in the interval $[\underline{a}, \bar{a}]$.

It is easy to see then that a firm can deviate profitably by choosing distribution $\hat{F}$, where $\hat{F}$ is constructed from $\tilde{F}$ by taking the measure on $\left[\underline{b}, \bar{b}\right]$ and some fraction $\epsilon > 0$

[5] Obviously, since $c > 0$, there are no symmetric equilibria in which the market price is 0.

of the measure on $[\underline{a}, \bar{a}]$ and collapsing them to their barycenter, and is set equal to $\tilde{F}$ everywhere else.[6]

$\hat{F}$ will have a point mass on some $\hat{x}$, which will occur with probability $\int_{\underline{b}}^{\bar{b}} d\tilde{F}(x) + \epsilon \int_{\underline{a}}^{\bar{a}} d\tilde{F}(x)$. For ϵ sufficiently small, $\hat{x} > z + \tilde{p}$, and by construction $\hat{F} \in \mathcal{M}\left(\tilde{F}\right)$. Thus, since $\tilde{F} \in \mathcal{M}(G)$, $\hat{F} \in \mathcal{M}(G)$. Finally, the net change in the firm's payoff is $\epsilon \int_{\underline{a}}^{\bar{a}} d\tilde{F}(x)(\tilde{p} - \alpha) > 0$, which concludes the pure strategy portion of the proof.

The proof that there are no symmetric mixed strategy equilibria in which consumers visit more than one firm is almost identical to its pure strategy analog, and so the proof may be found in Appendix A.1. ∎

In the standard Diamond paradox, firms have an incentive to raise prices slightly in order to take advantage of the quasi-monopoly power inferred upon them by the search cost. A similar effect occurs with regard to their information provision policies: they have an overwhelming incentive to provide slightly less information to consumers in order to increase the probability that they do not continue their searches.

Next, we find that, given that no consumer visits more than one firm, there are no equilibria in which consumers obtain any rents.

Lemma 2.3. *There are no symmetric equilibria in which consumers obtain any surplus.*

Proof. From the previous lemma, we know that there are no symmetric equilibria in which consumers visit more than one firm. First, suppose that there is an equilibrium in which the reservation value chosen by each firm, z, is greater than 0. Evidently, since a consumer visits just one firm, the reservation value must equal $\mu - c - \tilde{p}$, which is greater than 0 by assumption.

A firm's profit, conditional on a consumer's visit, is $\tilde{p} \leq \mu - c$. Clearly, then, a firm can deviate by charging price $p = \tilde{p} + c - \epsilon$. For $\epsilon > 0$ sufficiently small, this is a profitable deviation. Thus, there are no equilibria in which firms induce non-negative reservation values.

[6] In the parlance of Elton and Hill (1992), $\hat{F}$ is a fusion of distribution $\tilde{F}$. The notions of fusion and mean-preserving-contraction are (in this paper) equivalent.

Second, suppose that there is an equilibrium in which the reservation value chosen by each firm, z, is strictly less than 0. Now, each firm is facing a monopolist's problem where a consumer has an outside option of 0. Accordingly, a firm can extract all of a consumer's surplus by providing no information (choosing the degenerate distribution $\delta_X(\mu)$) and charging $p = \mu$. ■

Combining the two lemmata, we have the first main result.

Theorem 2.4. *In any symmetric equilibrium in which firms are visited, each firm makes a sale with certainty, conditional on being visited, and obtains the monopoly profit of μ. There is no active search and consumer surplus equals* 0.

3 Extensions/Alternative Specifications

In this section, we explore two additional scenarios. In the first, we maintain the assumption that prices are not publicly posted but allow for consumers to have private types that affect their match values. In the second, we return to the original specification in which consumers do not have private types but we impose that prices now are publicly posted.

3.1 Private Information

We begin this subsection with an example. The set-up is a special case of the general model set out in Section 2 with one modification. There are just two firms in the market and the random variable X_i (previously termed the match value)–which we now call the "quality"–is binary, taking value 1 with probability $\mu \in (0,1)$ and 0 with its complement. Now, consumers have private types: fraction $\rho \in (0,1)$ of the populace is type l and the rest are type h. These types affect consumers' match values, which are

$$u_l(0,p) = u_h(0,p) = -p, \quad u_l(1,p) = 1-p, \quad \text{and} \quad u_h(1,p) = 2-p\ .$$

It costs $c \in (0,\mu)$ for a consumer to visit a second firm.

Define

$$\bar{\rho} := \frac{(2-\mu)(1-(\mu-c))}{3(1-\mu)+\mu(\mu-c)+2c}\ .$$

Then,

Remark 3.1. If ρ is sufficiently large ($\rho \geq \bar{\rho}$), there is an equilibrium in the example in which a portion of consumers search actively and obtain positive surplus.

Proof. In this equilibrium, both firms set a price of 1 and provide full information. Low types (l) do not search but purchase immediately from the first firm they visit. High types (h) search: if they do not obtain a high match value at the first firm they visit the second. Low types obtain 0 surplus and high types obtain a surplus of $\mu + (1-\mu)(\mu - c)$.

Given the conjectured strategies of the firms, type 1's continuation value is $-c$ and type 2's continuation value is $\mu - c$. Each firm's problem is straightforward: for any price a firm sets, its payoff as a function of the induced posterior is a step function. Consequently, it maximizes its profits by providing full information and either setting a price so high that only consumers of type 2 purchase, or setting a lower price that is appealing to both types. It is simple to further reduce the number of potential optima to two: either a firm sets a price of 1 and serves both types, which is the conjectured behavior; or it sets a price of $2 - (\mu - c)$, which is just low enough that the high type will buy without continuing her search. The cutoff $\bar{\rho}$ is precisely when a firm's payoffs from these two approaches coincide. ■

In this example, the spirit of the paradox lingers. Although consumers of type h benefit from competition, consumers of type l neither search nor obtain any surplus. Nevertheless, consumers of type l do participate in the market–they purchase if their match value realization is high (they merely refrain from searching). This suggests that the fact that a consumer's first visit is free is important. Pursuing this hunch, if we modify the example so that a consumer's visit costs some $\epsilon > 0$, it is easy to see that an equilibrium in which firms provide full information and sell to both types is no longer an equilibrium–a firm can always pool beliefs above the stopping threshold of type l or raise prices slightly and obtain a strictly higher profit.

As we shortly discover, this observation is a special case of a more general result. Suppose that each consumer has a private type, random variable W, which takes values in the interval in some compact subset of $\mathbb{R}$, without loss of generality $[0,1]$. Let W be distributed according to some (Borel) cdf H and let W and each firm's quality random variable X_i be mutually independent (and recall that each X_i is i.i.d.). Furthermore, assume that

$$u(x_i, w, p_i) = v(x_i, w) - p_i \ ,$$

where v is linear and increasing in x_i.[7] Given this setup, the Diamond paradox reemerges, provided visiting the first firm is not free. For simplicity we establish the result for the case in which the cost of visiting the first firm is the same as the cost of subsequent visits.

Proposition 3.2. *If a consumer incurs cost $c > 0$ to visit the first firm, there are no symmetric pure strategy equilibria in which a nonzero measure of consumers visit any firm.*

The proof of this proposition is quite similar to that of Lemma 2.2, and so we relegate it to Appendix A.2.

A couple things come to mind when reflecting on this result and the assumptions and conditions that give rise to it. First, when consumers are homogeneous, the assumption that consumers incur no visit cost at the first firm is relatively innocuous–regardless of whether the first visit is free, there is no active search; and the existence of a "first visit cost" merely determines whether trade occurs between an effective monopolist and its captive consumers. In contrast, with private information, as our example reveals, there may exist equilibria with active search when consumers' first visits are free. Second, we assume in the proposition (and in the example) that the consumer heterogeneity is of a certain sort–different consumers have different valuations for a firm's product, but they all agree on the ranking of products. This seems reasonable in some cases, but there are other instances in which this assumption seems less appropriate.

3.2 Advertised Prices

Next, we restore the initial assumption that consumers are homogeneous. Perhaps surprisingly, the no active search result continues to hold, *even if prices are advertised* and so search is directed. Specifically, we assume that a consumer observes the price set by each

[7] Note that this formulation justifies the moniker "quality" for a firm's random variable X_i: all types of consumers have a higher match value with a higher quality product.

firm before starting her search, but still must visit a firm to observe its (expected) match value. As noted by Armstrong and Zhou (2011), "price comparison websites are now a major part of the retailing website." In many markets, consumers can obtain price quotes for free, before embarking on their costly searches. Information, on the other hand, is perhaps more difficult to advertise, so it is plausible that firms can advertise prices but not how much information they provide.

We continue to use the weak PBE solution concept. Since price deviations are now observable, we allow consumers the freedom to have any belief about a firm's signal following an off-path price choice by that firm. In the same spirit of the assumption in the hidden prices case, we stipulate that a price deviation by a firm does not affect consumers' conjectures about the distributions chosen by firms that choose on-path prices. Similarly, during a consumer's search, her beliefs about the signals chosen by theretofore unvisited firms are unaffected by other firms' deviations.

Lemma 3.3. *In any symmetric equilibrium in which firms make strictly positive profits, there is no active search. That is, consumers visit no more than one firm.*

The detailed proof of this result is left to Appendix A.3, and the intuition behind it mirrors that of Lemma 2.2: firms prefer to pool values right above their reservation values since there is no benefit to inducing any higher values.

Next, because there is no active search, we find that firms cannot make positive profits in any symmetric equilibrium.

Lemma 3.4. *There are no symmetric equilibria in which firms make nonzero profits. Equivalently, there are no symmetric equilibria in which firms charge any price other than* 0*.*

The full proof of this result is rather tedious and is, therefore, banished to Appendix A.4. The basic logic is straightforward; however: if firms are making positive profits, a firm can always deviate by providing no information and setting a price just lower than the maximal price chosen at equilibrium. Because there is no active search (Lemma 3.3), this guarantees that firm a discrete jump in its profits.

The next lemma extinguishes the last possibility for there to be a symmetric equilibrium with active search.

Lemma 3.5. *There exist no symmetric equilibria with active search in which firms make zero profits.*

The sketch of the proof is as follows: in any purported equilibrium with active search, firms' distributions must place positive measure on values that are strictly below the no information (and $p = 0$) reservation value $\mu - c$. A firm can deviate by choosing a small (but strictly positive) price such that no matter a consumer's conjecture about the firm's information provision–which, recall, is hidden until a consumer's visit to that firm–she will visit it. When she does, which occurs with positive probability, she will purchase from the deviator, yielding it a positive profit. The details may be found in Appendix A.5.

At last, we arrive at the second main result.

Theorem 3.6. *Any symmetric equilibrium must be one in which each firm posts price $p = 0$ and chooses a distribution that induces the minimum reservation value, $\mu - c$.*

Proof. Trivially, a firm cannot deviate profitably by choosing a different distribution over values (and keeping price $p = 0$). Next, we stipulate that a consumer assumes that any firm that deviates to a price $p \neq 0$ chooses a completely uninformative distribution. Accordingly, that firm will never be visited and therefore such a deviation is not profitable.

Second, uniqueness follows from Lemmata 3.3, 3.4, and 3.5. ∎

In any symmetric equilibrium, each firm's profit is 0 and each consumer's welfare is just μ (they do not incur a search cost, c, since that applies only to searches beyond the first firm). Although consumers obtain some surplus, there is no benefit from increased competition, i.e., a market with two firms is just as good for a consumer as one with infinitely many. Finally, note that the given timing specification of this subsection–in which firms post prices and choose information policies simultaneously (or at least are neither aware of the market vector of prices nor consumers' search histories)–is inessential to the main result. Regardless of what a firm may observe before choosing its information policy, it does not benefit from any values that are strictly above the conjectured reservation value the world (consumers and other firms) assigns to it, which drives the result.

4 Discussion

Wolinsky (1986) provides a compelling resolution for the Diamond paradox. Product differentiation and imperfect information allow for more realistic and empirically-relevant equilibria in which consumers search and in which competitive forces have real effects on the pricing decisions of firms. Indeed, a key aspect of Wolinsky's model is that consumers obtain information about their match values upon visiting firms.

Crucially, that information *must come from somewhere* and in particular, firms have a say in how much information a consumer's visit will glean about their respective products. This paper suggests a weakness in using imperfect information and horizontal differentiation alone to generate competitive pricing and active search in search models. Namely, if firms may choose how much information to provide about their products and that information is unadvertised then the Diamond paradox reemerges. Modifications to the model–like, e.g., assuming that a fraction of consumers can search without cost á la Stahl (1989) or assuming consumers are heterogeneous in other ways (as in Section 3.1 of this paper)–are needed to restore active search.

It is important to keep in mind that, just as the Diamond result requires that consumers learn prices only after paying a search cost, the hidden nature of the firms' information choices is essential to the results that we encounter here. If information itself can be advertised–i.e., consumers can observe each firm's chosen distribution over values without paying a visit cost–there are equilibria in which consumers search and obtain positive surplus. This situation with posted information is the subject of Au and Whitmeyer (2020b).

Furthermore, it seems quite plausible that although firms can advertise information policies *partially*, they have no way of fully specifying or committing to more information than a certain (minimal amount). Consequently, the strong incentives for firms to deviate and under-provide information we find here suggest that firms will not provide more information than they can (explicitly) commit to. Casual observation suggests that this is somewhat accurate–there are many anecdotes of car dealerships reneging on advertised test-drives[8] or real-estate agents whisking progressive tenants through their apartment tours.

References

Mark Armstrong and Jidong Zhou. Paying for prominence. *The Economic Journal*, 121 (556):F368–F395, 2011.

Mark Armstrong and Jidong Zhou. Consumer information and the limits to competition. *Mimeo*, November 2019.

Pak Hung Au and Mark Whitmeyer. Attraction versus persuasion. *Mimeo*, November 2020a.

Pak Hung Au and Mark Whitmeyer. Advertised prices and information. *Mimeo*, November 2020b.

Simon Board and Jay Lu. Competitive information disclosure in search markets. *Journal of Political Economy*, 126(5):1965–2010, 2018.

Peter A Diamond. A model of price adjustment. *Journal of Economic Theory*, 3(2):156 – 168, 1971.

Mustafa Dogan and Ju Hu. Consumer search and optimal information. *Mimeo*, August 2018.

Glenn Ellison and Alexander Wolitzky. A search cost model of obfuscation. *The RAND Journal of Economics*, 43(3):417–441, 2012.

J. Elton and T. P. Hill. Fusions of a probability distribution. *The Annals of Probability*, 20 (1):421–454, 1992.

[8] See, e.g., some of the stories detailed in the following thread: https://www.reddit.com/r/cars/comments/66r9ze/ever_had_a_dealer_not_let_you_test_drive_a_car/. One user describes a test-drive consisting merely of a drive around the dealership's parking lot.

Wei He and Jiangtao Li. Competitive information disclosure in random search markets. *Mimeo*, June 2021.

Ilwoo Hwang, Kyungmin Kim, and Raphael Boleslavsky. Competitive advertising and pricing. *Mimeo*, 2018.

Maarten Janssen and Sandro Shelegia. Consumer search and double marginalization. *American Economic Review*, 105(6):1683–1710, 2015.

Maarten Janssen and Sandro Shelegia. Beliefs and consumer search in a vertical industry. *Journal of the European Economic Association*, 18(5):2359–2393, 2020.

Dale O. Stahl. Oligopolistic pricing with sequential consumer search. *The American Economic Review*, 79(4):700–712, 1989.

Martin L. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–654, 1979.

Asher Wolinsky. True Monopolistic Competition as a Result of Imperfect Information*. *The Quarterly Journal of Economics*, 101(3):493–511, 08 1986.

Jidong Zhou. Improved information in search markets. *Mimeo*, October 2020.

A Omitted Proofs

A.1 Lemma 2.2 Mixed Strategy Portion Proof

Proof. It is simple to modify the (within-text) proof in which we assume pure strategies to accommodate mixing by firms. In such a (alleged) symmetric equilibrium, each firm chooses the joint distribution $F(x,p)$, over values and prices, $[\underline{x},\bar{x}]\times\left[\underline{p},\bar{p}\right]$, where each conditional distribution over values, $F_{X|P}(x|P=p)\in\mathcal{M}(G)$. No randomness is resolved, however, until consumers visit firms and so we may define a new random variable $Y := X-P$ with distribution H on some interval.[9]

[9] Clearly, because prices are non-negative, the upper bound of the support must be (weakly) less than 1.

Accordingly, the reservation value induced by a consumer's conjectured $\tilde{H}$ is

$$c = \int_z^1 (y - z)\, d\tilde{H}(y) \ .$$

Note that the induced reservation values, z, and distributions, H, chosen by firms are deterministic, and identical (because this a symmetric equilibrium).

The remainder is (more-or-less) identical to the pure strategy case with some minor subtleties. Because consumers are searching actively, $\tilde{H}$ must be such that a consumer strictly prefers to stop and strictly prefers to continue with strictly positive probability. There are two cases: first, it is possible that for some p' in the support of a firm's mixed strategy, the associated $\tilde{F}_{X|P}(x|p')$ has the same property as $\tilde{H}$–namely, both (strict) stopping and (strict) continuation values occur with strictly positive probability. As in the pure strategy case, a firm can deviate profitably by pooling these values carefully.

On the other hand, it may be possible that there is no such p'. However, then, there must exist some p'' and $\tilde{F}_{X|P}(x|p'')$ in the support of a firm's mixed strategy such that, given p'', stopping values occur with probability one and strict stopping values realize with strictly positive probability. Consequently, a firm can deviate by providing no information (choosing the distribution $\delta_X(\mu)$)[10]–which, at price p'', must leave the consumer with strict incentive to stop–and instead charging some price $p'' + \epsilon$, for a sufficiently small (yet strictly positive) ϵ. ■

A.2 Proposition 3.2 Proof

Proof. Suppose for the sake of contradiction that there is an equilibrium in pure strategies in which a nonzero measure of consumers visit at least one firm. The reservation value of type w, $z(w)$, is defined by the equation

$$c = \int_{T(z(w),w,\tilde{p})}^1 (v(x,w) - \tilde{p} - z(w))\, d\tilde{F}(x) \ ,$$

where $\tilde{F}$ and $\tilde{p}$ are (recall) the conjectured distribution and price, respectively, chosen by each firm at equilibrium and $T(\gamma, w, p)$ is implicitly defined by

$$v(T(\gamma, w, p), w) = \gamma + p \ \forall \ w \in W \ .$$

[10] $\delta_X(\mu)$ denotes the degenerate distribution $\mathbb{P}(X = \mu) = 1$.

Let A denote the subset of types who are actively searching:

$$A := \{w \in W : z(w) > \mathbb{E}_{\tilde{F}}[v(X,w)] - \tilde{p} - c,\ z(w) \geq 0\} \ ,$$

and B denote the subset of types who are purchasing but not actively searching:

$$B := \{w \in W : z(w) = \mathbb{E}_{\tilde{F}}[v(X,w)] - \tilde{p} - c \geq 0\} \ .$$

There are two cases to consider; that in which there is a strictly positive measure of consumers who are actively searching (A has nonzero measure), and that in which there is not a strictly positive measure of consumers who are actively searching (A has zero measure and B has nonzero measure).

The second case is easy. If a consumer arrives at a firm she stops immediately and purchases if and only if she obtains a value weakly greater than $\mathbb{E}_{\tilde{F}}[v(X,w)] - \tilde{p} - c$. Accordingly, a firm can provide no information (the degenerate distribution $\delta_X(\mu)$) and set price $p = \tilde{p} + c/2$, obtaining a strictly higher profit.

In the first case, because each type $w_a \in A$ is actively searching, there must be an interval of (posterior expected) qualities $[\underline{t}, \bar{t}]$ with $\int_{\underline{t}}^{\bar{t}} d\tilde{F} > 0$ such that $v(\underline{t}, w) - \tilde{p} > z(w_a)$ for all $w_a \in A$. Analogously, there must be an interval of qualities $[\underline{s}, \bar{s}]$ with $\int_{\underline{s}}^{\bar{s}} d\tilde{F} > 0$ such that $v(\bar{s}, w) - \tilde{p} < z(w_a)$ for all $w_a \in A$. Because types in B are buying but not actively searching, $v(\underline{s}, w) - \tilde{p} \geq z(w_b)$ for all $w_b \in B$.

The remainder of the proof is the same as that of Lemma 2.2 *mutatis mutandis*, since $v(x', w) - \tilde{p} \geq z(w_b)$ for all $x' \in [\underline{s}, \bar{t}]$, for all $w_b \in B$. ■

A.3 Lemma 3.3 Proof

Proof. In any symmetric equilibrium, each firm chooses a joint distribution, $F(x,p)$, over values and prices, $[\underline{x}, \bar{x}] \times \left[\underline{p}, \bar{p}\right]$, where each conditional distribution over values, $F_{X|P}(x|P = p) \in \mathcal{M}(G)$. Given a firm's choice of (on-path) price $p^* \in \left[\underline{p}, \bar{p}\right]$, its (on-path) choice of distribution over values $F_{X|P}(x|P = p^*)$ must yield a maximal payoff for the firm, given the strategies of the other firms.

For each price, the corresponding conditional distribution over values corresponds to a reservation value, z, which is pinned down by Equation ⋆. Define $\mathcal{Z}$ to be the set of all reservation values that are induced on-path. Joint distribution $F(x,p)$ induces a distribution over reservation values $\Phi(z)$. Note that, in contrast to the previous section's case, in which both prices and distributions are hidden, the observability of firms' price choices means that the reservation values chosen by firms may be random.

Since firms are making strictly positive profits, any on-path price, p^*, must itself be strictly positive. Denote by $z^* \in \mathcal{Z}$ the corresponding reservation value. Evidently, a firm's payoff, conditional on being visited, from any value $x \geq \max\{0, z^*\} + p^*$, is p^*.

First, suppose that there exists an on-path price p^* and conjectured distribution $\tilde{F}_{X|P}(x|p^*)$ that induces a z^* that is weakly greater than 0. As in the proof for Lemma 2.2, suppose for the sake of contradiction that conditional on her arrival at the firm a consumer strictly prefers to stop and strictly prefers to continue her search (or select her outside option) with strictly positive probability. However, all values in the stopping set yield a payoff of p^* and all beliefs in the continuation set yield a payoff that is strictly below p^*. Consequently, as in the proof for Lemma 2.2, after posting price p^*, a firm can (secretly) deviate by providing slightly less information and fusing a subset of the beliefs in the continuation and stopping sets.

Second, suppose that there exists an on-path z^* that is strictly less than 0. In such an equilibrium, if a firm is visited, it is the first and only firm that is visited by the consumer, since the first visit is the only one that does not impose on the consumer a search cost. ■

A.4 Lemma 3.4 Proof

Proof. First, we establish an auxiliary claim:

Claim A.1. *In any equilibrium in which firms make strictly positive profits we must have the following:*

1. *If $1 \geq p^* > \mu$ is chosen on-path, distribution $F_{X|P}(x|p^*)$ has a mass point of size $q^* > 0$ on p^* and has no support strictly above p^*;*

2. *If $\mu \geq p^* > 0$ is chosen on-path, distribution $F_{X|P}(x|p^*)$ must have support entirely above p^*. A firm is selected with certainty, conditional on being visited.*

Proof. If $1 \geq p^* > \mu$ is chosen on-path, then if the associated reservation value, $z^* \geq 0$, the conjectured distribution must be such that $\int_{\underline{x}}^{\bar{x}} d\tilde{F}_{X|P}(x|p^*) > 0$ for some $\bar{x} \geq \underline{x} > z^* + p^*$. Moreover, because $p^* > \mu$, the conjectured distribution must be such that $\int_{\underline{w}}^{\bar{w}} d\tilde{F}_{X|P}(x|p^*) > 0$ for some $\underline{w} \leq \bar{w} < p^*$. A firm's payoff from any belief $x \geq z^* + p^*$ is p^* and from any belief $x < p^*$ is 0, so it can deviate profitably by fusing the measure on $[\underline{x}, \bar{x}]$ with a fraction, $\epsilon > 0$, of the measure on $[\underline{w}, \bar{w}]$ (collapsing them to their barycenter). We conclude that $z^* < 0$.

Define

$$F^-_{X|P}(a|p^*) := \sup_{w<a} F_{X|P}(w|p^*) \ ,$$

and observe that $F^-_{X|P}(p^*|p^*) < 1$, or else a firm would make 0 profit from choosing p^*. Furthermore, any values $x < p^*$ yield a firm a profit of 0 and since $p^* > \mu$, $F^-_{X|P}(p^*|p^*) > 0$. Consequently, a firm must have $F_{X|P}(p^*|p^*) = 1$ since otherwise it could (as in the previous paragraph) fuse a positive measure of values strictly below p^* with a positive measure of values strictly above p^* and obtain a higher payoff.

Define

$$q^* = q(p^*) := F_{X|P}(p^*|p^*) - F^-_{X|P}(p^*|p^*) \ ,$$

i.e., q^* is the size of the mass point that $F_{X|P}(x|p^*)$ places on p^*. Accordingly, $z^* = -c/q^* < 0$.

If $\mu \geq p^* > \mu - c$ is chosen on-path, then if the associated reservation value, $z^* \geq 0$, a firm must be selected with certainty, conditional on being visited. This is because $z^* \geq 0 > \mu - c - p^*$, which is the minimal reservation value that can be induced. Thus, if a firm were not selected with certainty, conditional on being visited, it could always fuse values strictly below and strictly above $z^* + p^*$. Accordingly, all values must be weakly greater than p^*, i.e., $F^-_{X|P}(p^*|p^*) = 0$. If the associated reservation value, z^* is strictly negative, the result is trivial since a firm could always just provide no information (choose the degenerate distribution $\delta_X(\mu)$) and be selected with certainty, conditional on being visited.

Finally, if $\mu - c \geq p^* > 0$, then the minimum reservation value that can be induced is $\mu - c - p^* \geq 0$. Either $F^-_{X|P}(z^* + p^*|p^*) = 0$, in which case $z^* = \mu - c - p^*$ and so a firm is clearly selected for sure, conditional on being visited; or $F^-_{X|P}(z^* + p^*|p^*) > 0$ and so $z^* > \mu - c - p^*$ and $F_{X|P}(z^* + p^*|p^*) < 1$. In that case either a firm is selected for sure or it can deviate profitably by fusing portions of the measure strictly above and below $z^* + p^*$. ■

Now let us finish proving the lemma. Let $\hat{p}$ be the maximal price that is chosen on-path, with associated reservation value $\hat{z}$.

First, suppose that $\hat{p} \leq \mu$. From Lemma 3.3 and Claim A.1, following any on-path p, a firm must be chosen for sure, conditional on being visited. Clearly, $\hat{p}$ must be chosen with strictly positive probability on-path, since there is no active search. Otherwise, choosing that price would guarantee that that firm is never visited, yielding a profit of 0.

Consequently, all values, x, must be such that $x - \hat{p} \geq \max\{0, \hat{z}\}$. Thus, $\hat{z} = \mu - c - \hat{p}$. But then a firm can deviate profitably to some price $\hat{p} - \eta$, where $\eta > 0$, and the degenerate distribution $\delta_X(\mu)$. Because $\mu - c - \hat{p} + \eta$ is the minimum reservation value such a deviation could induce, the deviating firm obtains a discrete jump up in its payoff.

Second, suppose that $\hat{p} > \mu$. The probability that firms are choosing prices that are weakly greater than μ must be strictly positive (or else choosing $\hat{p}$ would result in 0 profit for a firm). From a consumer's viewpoint, all on-path prices $p^* \geq \mu$ are equivalent. Indeed, for all such $p^* \geq \mu$, $F_{X|P}(p^*|p^*) = 1$.

A firm's payoff from choosing any $p^* \geq \mu$, conditional on being visited, is $q(p^*)p^*$, where, recall, $q(p^*)$ is the size of the mass point on p^*. Moreover, by definition, $q(p^*)p^* \leq \mu$. But then, a firm can deviate profitably by choosing some price $p = \mu - \epsilon$, for a small but strictly positive ϵ, and the degenerate distribution $\delta_X(\mu)$. No matter a consumer's belief about the firm's distribution after observing this price, a deviating firm must still be visited (and thus purchased from eventually) before any firm that is choosing $p^* \geq \mu$. Accordingly, for a sufficient small $\epsilon > 0$, a firm will receive a discrete jump in its payoff and so there exists a profitable deviation.

■

A.5 Lemma 3.5 Proof

Proof. Obviously, in any symmetric equilibrium in which firms make zero profits, then if a firm chooses any price $p > 0$ it must be visited with probability 0. Thus, we impose that firms post prices $p = 0$. If $\mu - c > 0$, it suffices to to show that there are no equilibria in which firms choose distributions over values that induce reservation values strictly greater than $\mu - c$.

Suppose for the sake of contradiction that such an equilibrium exists. Then, a firm's distribution over values must be such that the probability of a realization that is weakly below $\mu - c - \gamma$ (for some $\gamma \in (0, \mu - c]$) is strictly larger than 0. However, then a firm can deviate profitably to some price p that satisfies $0 < p < \min\{\gamma, c, \mu - c\}$ and the degenerate distribution $\delta_X(\mu)$. Even should a consumer assign it the most pessimistic belief about its distribution, yielding a reservation value of $\mu - c - p$, the firm would still be visited (and selected) with strictly positive probability, yielding a strictly positive profit.

If $\mu - c \leq 0$, there is no active search on-path since any value a consumer obtains at the first firm is at least weakly greater than the reservation values of the other firms.

■