Can price dispersion be supported solely by information frictions?

José F. Tudón M.*

October 26, 2020

Abstract

Even with identical consumers and identical firms, if firms set prices in a first stage, and if consumers search sequentially in a second stage, price dispersion arises in the form of a mixed-strategy Nash equilibrium. One only needs to assume consumers know the realized price distribution and that they do not know which firm has what price. In contrast to Burdett and Judd (1983), price quotes are not required to be "noisy." Moreover, actual search is predicted to be nontrivial. (*JEL* L13, D83, D21)

Keywords: Price dispersion; information frictions; sequential search.

1 Introduction

In his economics-of-information article, Stigler (1961) wrote that "it would be metaphysical, and fruitless, to assert that all [price] dispersion is due to heterogeneity." Since then, several studies have confirmed the empirical significance of price dispersion.[1] However, after Stigler's seminal paper, Diamond (1971) presented a challenge: if firms and consumers are identical, and if consumers pay to sequentially search for prices, the only Nash equilibrium is the monopoly price. Intuitively, if all firms

*ITAM, Río Hondo 1, Ciudad de México 01080, Mexico (e-mail: jtudon@itam.mx). Thanks to Alex Frankel, Ali Hortaçsu, Philip Reny, Enrique Seira, Chad Syverson, Balázs Szentes, Gábor Virág, and workshop participants at the University of Chicago.

[1] See, for example, Kaplan and Menzio (2015), Hortaçsu and Syverson (2004), or Sorensen (2000).

set the monopoly price, the consumer will not search, but if the consumer will not search, all firms set the monopoly price. As Reinganum (1979) said, the Diamond paradox seemed to imply "imperfect information alone is insufficient to support price dispersion."

This paper adds to recent efforts by showing with minimal assumptions that imperfect information alone is indeed sufficient to support price dispersion. I assume firms fix their prices *before* consumers search, and that consumers know the distribution of actual prices being charged in the market, but do not know which firm is charging which price. This paper shows this information structure is sufficient to generate price dispersion, even with homogeneous firms and homogeneous consumers.

The informational assumption that drives the result is that the consumer knows the *realized* price distribution when she starts searching the market. As Menzio and Trachter (2015) explain,

> "This is the same assumption made in the vast majority of search models where the price distribution is exogenously given (e.g., Stigler, 1961; McCall, 1970). Indeed, only a few papers model the buyer's problem of learning about the price distribution while searching (see, e.g., Rothschild, 1974; Burdett and Vishwanath, 1988). The assumption is also common in models where prices are endogenous (see, e.g., Pratt, 1979; Rob, 1985; Stiglitz, 1987). Other models, such as Reinganum (1979), Burdett and Judd (1983), and Stahl II (1989), assume that buyers do not observe the actual price distribution, but have rational expectations about its equilibrium value."

For example, in a heterogeneous production-cost model, the consumer is assumed to know about the distribution of production costs in order to compute the equilibrium price distribution. However, this assumption seems as strong as the one used on this paper, in particular, in an extensive-form game. I thus motivate my assumption by thinking of consumers who learn the realized price distribution by reading the newspaper or searching online before engaging in physical search. For the modern consumer, who is connected to the internet, this assumption is entirely plausible. In fact, in the US auto market, a popular price-comparison website literally shows the actual price distribution to consumers.

Intuitively, in the Diamond model, firms cannot profit from a reduction in prices, because firms and consumers choose their strategies simultaneously. However, the passivity of the consumer is broken when we unfold the model into two stages, because her strategy now becomes a complete contingent plan. In this case, firms anticipate the consumer's reaction to a cut in prices, and thus have incentives to steal the market instead of sharing it. The main result is that even if we have no a priori reason to expect ex-post price heterogeneity, price dispersion arises in the form of a mixed-strategy Nash equilibrium. Finally, this paper sidesteps "a host of technical problems that haunt [this information structure], e.g. troublesome integer conditions (typically overlooked) in constructing dispersed equilibria" (Stahl II, 1996), which is a strong reason why the literature has favored other information structures.

This paper is related to the search literature interested in obtaining price dispersion under minimal assumptions.[2] For example, Burdett and Judd (1983) achieve price dispersion by allowing price quotes to be stochastic. In contrast to my paper, they predict no actual search in equilibrium.

Perhaps the closest paper is Menzio and Trachter (2015), where buyers search sequentially among a continuum of small firms and a single large firm. In their model, price dispersion arises as equilibrium mixed strategies, and price dispersion is obtained by exploiting the heterogeneity between small firms and a single large firm. In some sense, this paper shows the result in Menzio and Trachter (2015) can be extended to a market composed of identical large firms. However, in this paper, dispersion is obtained while maintaining ex ante homogeneity.

Beyond search models, this paper is also related to literature that establishes price dispersion as an equilibrium outcome of a Bertrand game. Kaplan and Wettstein (2000) find that mixed-strategy equilibria, which yield positive profits, are equilibrium strategies only for environments with unbounded revenues. Similarly, Baye and Morgan (1999) find any finite payoff vector can be achieved in a symmetric Nash equilibrium, when profits are unbounded.

2 The model

The model has two stages. In the first stage, n identical firms simultaneously produce an homogeneous good and fix their prices $p_i \in \mathcal{P}_i,\ i = 1, 2, \dots, n$, where $\mathcal{P}_i$ may be

[2] See Rauh (2007), who considers heterogeneity in search costs, demand functions, and production

either continuous or discrete. I assume no production costs.[3] After prices are chosen, the second stage begins, and a unique consumer enters the market.

The consumer has a unit demand, and a willingness to pay normalized to 1. That is, she will never pay more than 1 to buy the good. To be able to buy, the consumer must pay a fixed search cost s each time she arrives at a store. The first search is free.[4] In each search, the consumer randomly arrives at one of the n stores with equal probability. There, she must decide whether she buys the good and exits the market, or instead continues her search. Sampling is with replacement.[5]

To be consistent with random search, suppose the consumer has no recall (assuming perfect recall yields the same results).[6] Also note $s = 0$ is equivalent to perfect information, because the consumer could search forever.

Finally, the structure of the model and the rationality of the consumer and firms are common knowledge.

To be clear, the main assumption is that the consumer knows the *realized* price distribution when she starts her search. At this second stage of the game, such distribution is a realization of the equilibrium price distribution.

2.1 Optimal search

The goal of the consumer is to buy at the lowest price, so at each period, she chooses between either buying at the quoted price or paying the search cost and going to another store at random. The marginal cost of the extra search is s at any point. The marginal benefit given the current price at hand, p, is the expected discount $\frac{1}{n}\sum_{i=1}^{n}(p - p_i)\mathbb{1}\{p_i \leq p\}$, where $\mathbb{1}$ is the indicator function. A *reservation price*, $p^r(\mathbf{p}, s)$, is a function of the vector of prices in the economy, $\mathbf{p} = (p_1, \dots, p_n)$, and the search cost that sets the marginal benefit equal to the marginal cost of the extra search. Lemma 1 shows that such a price exists, and is implicitly defined by

functions, thus subsuming most of the literature as special cases. Current theoretical research focuses on the Stahl II (1989, 1996) and Wolinsky (1986) models, which lend themselves easily to structural estimation.

[3] Janssen, Pichler and Weidenholzer (2011) consider stochastic production costs.

[4] Ellison and Wolitzky (2012) consider firms that can increase the search cost by obfuscating. Janssen, Moraga-González and Wildenbeest (2005) consider a costly first search.

[5] See Carlson and Mcafee (1983).

[6] Daughety and Reinganum (1992) consider endogenous recall.

$$p^r(\mathbf{p}, s) = s + \frac{1}{n} \sum_{i: p_i \leq p^r(\mathbf{p},s)} p_i + \frac{\#\{i : p_i > p^r(\mathbf{p}, s)\}}{n} p^r(\mathbf{p}, s), \tag{1}$$

where # is the cardinality of a set.

Lemma 1 (Properties of the reservation price $p^r(\mathbf{p}, s)$)**.**

1. $p^r(\mathbf{p}, s)$ exists and is increasing in s; in particular, $p^r(\mathbf{0}, s) = s$.
2. $p^r(\mathbf{p}, s)$ is bounded above by $s + \frac{1}{n} \sum_{j=1}^{n} p_j$.
3. $p^r(\mathbf{p}, s)$ is continuous in $\mathbf{p}$ and s.
4. $\mathbf{p} \geq \mathbf{p}^* \Rightarrow p^r(\mathbf{p}, s) \geq p^r(\mathbf{p}^*, s)$; where $\geq$ is element-wise.

Proof. Existence comes from the fact that the right-hand side of equation (1) is continuous, concave in p^r, and equal to $s > 0$ when $p^r = 0$. Note also that the right-hand side of equation (1) becomes $s + \frac{1}{n} \sum_{j=1}^{n} p_j$ when $p^r > \max_j\{p_j\}$.

Continuity of $p^r(\mathbf{p}, s)$ is straightforward. Finally, increasing any price will not decrease the right-hand side of equation (1). □

The optimal stopping rule and ending of the model is the standard condition found in search models: stop searching when $p \leq p^r(\mathbf{p}, s)$.[7] Let R^* denote this optimal stopping rule, which is a dominant strategy. Also, let $\mathcal{R}$ denote the space of strategies for the consumer's stopping rules. Therefore,

Assumption 1. The consumer always follows the dominant strategy R^*.

2.2 Optimal firm's strategy

Let the profits of firm i be denoted by $\pi_i(\mathbf{p}, R)$, which is a function of prices and the stopping rule of the consumer. Assume firms are risk neutral. Then, because of Assumption 1, profits can be stated as

$$\pi_i^*(p_i, p_{-i}, R^*) = \begin{cases} p_i & \text{if } p_i \leq p^r(\mathbf{p}, s) \text{ and } p_j > p^r(\mathbf{p}, s)\ \forall j \neq i \\ \dfrac{p_i}{k} & \text{if } p_i \leq p^r(\mathbf{p}, s) \text{ and } p_j \leq p^r(\mathbf{p}, s) \\ & \quad \forall j \in J_i \text{ with } k = \#J + 1 \\ 0 & \text{in any other case,} \end{cases} \tag{2}$$

[7] To see the optimality of this rule, recall that for prices above p^r, the marginal benefit of a search is higher that its marginal cost. But when the consumer finds a price below p^r, searching is no longer

where p_{-i} is the vector of prices p_j, with $j \neq i$, and $J_i = \{j | j \neq i,\ p_j \leq p^r(\mathbf{p}, s)\}$. Using notation from Reny (1999), the game that the firms play is defined next.

Definition 1 (Game 1)**.** Given the consumer's stopping rule R^*, the strategic-form game that the n firms play is given by $\{(\Delta\mathcal{P}_i, \pi_i(\cdot))_{i=1}^n\}$, where n is the number of firms, $\mathcal{P}_i$ is the strategy space for each firm, and π_i is firm i's profit function.

Because the firms will never choose a price above 1, henceforth, consider two cases: (1) a continuous $\mathcal{P}_i = [0, 1]$ for all i; or (2) a discrete $\mathcal{P}_i = \left\{0, \frac{1}{v}, \frac{2}{v}, \ldots, 1\right\}$, where $v + 1$ is the number of grid points.

Define price dispersion as a nontrivial mixed strategy that arises in equilibrium. We are interested in the cases where *pure*-strategy equilibria do not exist, but *mixed*-strategy equilibria do. The following theorem shows Game 1 always has a Nash equilibrium, and the equilibrium must be mixed if the following price-dispersion conditions are met.

Definition 2 (Price-dispersion conditions)**.**

- For continuous pricing: $s > 0$ and $ns < \frac{n-1}{n}$.
- For discrete pricing: $s \geq \frac{n-1}{n} \cdot \frac{1}{v}$ and $\left(ns, \frac{n-1}{n}\right) \cap \mathcal{P}_i \neq \emptyset$.

Theorem 3 (Price dispersion)**.** Game 1 always has a Nash equilibrium. Moreover, (a) Game 1 has no Nash equilibria in *pure* strategies if and only if the conditions in Definition 2 hold, and (b) Game 1 has no Nash equilibria in *mixed* strategies if and only if the conditions in Definition 2 do not hold. That is, a nontrivial price dispersion arises only when the conditions in Definition 2 hold.

Theorem 3 formally discards *pure*-strategy equilibria from the game under some "Goldilocks" conditions: an intermediate range where the search cost is not too low and not too high. Intuitively, a small search cost amounts to perfect information, which yields competition à la Bertrand. On the other hand, if the search cost is too high, firms enjoy market power, because information is too expensive for the consumer, which results in a captive costumer. Thus, a high search cost yields the Diamond equilibrium. Finally, and aligned with the intuition of Stahl II (1989), costs accumulate when searching among many firms. Therefore, the search cost must be weighted by the number of firms in the market.

worthwhile. Note the optimal search rule is myopic (Weitzman, 1979).

In the proof, I show Game 1 always has a Nash equilibrium regardless of the set of parameters. However, if the search cost meets with the theorem's conditions, the equilibrium must be in mixed strategies for the firm.

2.3 Proof of Theorem 3

First, I prove existence of a Nash equilibrium for Game 1.

If $\mathcal{P}_i$ is discrete for each i, the sets $\mathcal{P}_i$ have a finite number of elements. Furthermore, Game 1 is a finite strategic-form game. Therefore, it has a mixed-strategy Nash equilibrium. Corollary 7 characterizes the equilibrium.

On the other hand, the case where $\mathcal{P}_i$ is continuous presents serious challenges, because the payoffs $\pi_i^*(p_i, p_{-i}, R^*)$ are not continuous in (p_i, p_{-i}) nor quasiconcave in p_i. It follows that the convexity and upper hemicontinuity of the best-response correspondence cannot be assured. Therefore, the existence of a mixed-strategy Nash equilibrium cannot be obtained with standard arguments. Fortunately, Reny (1999) establishes conditions under which equilibria exist in a discontinuous game.

Let q_i be the mixed strategy of firm i and $\mathbf{q} = (q_1, \ldots, q_n)$. Formally, $q_i \in \Delta\mathcal{P}_i$ is a probability measure on Borel subsets of $\mathcal{P}_i$, $\mathcal{B}(\mathcal{P}_i)$. Moreover, $\Delta\mathcal{P}_i$ is a compact metric space when endowed with the Prohorov metric (Billingsley, 1999). Also, define $\pi_i^*(\mathbf{q}) \equiv \int_{\mathcal{P}} \pi_i^*(\mathbf{p}, R^*) d\mathbf{q}$ for all $\mathbf{q} \in \Delta\mathcal{P} \equiv \times_i \Delta\mathcal{P}_i$, and endow all product sets with the product topology.

To show a Nash Equilibrium in mixed strategies exists, we need to show the mixed extension of Game 1 is *better-reply secure*, as in Reny (1999).[8] I conduct the proof in a series of steps. First, I show the sum of the payoffs is upper semicontinuous in Lemma 2. Second, I prove a property that implies the game is payoff secure in Lemma 3. Third, I combine Lemmas 2 and 3 to show the game is better-reply secure, and conclude it must have a Nash Equilibrium.

Lemma 2. $\sum_{i=1}^{N} \pi_i^*(\mathbf{p})$ is upper semicontinuous in $\mathbf{p}$ on $\mathcal{P}$.

Proof. By definition, $\sum_{i=1}^{n} \pi_i^*(\mathbf{p})$ is equal to the mean of the prices that are equal to or less than the reservation price. Fix any $\mathbf{p}$ and any $\varepsilon > 0$. Note that increasing any element of $\mathbf{p}$ by some small amount $\delta > 0$ will either increase $\sum_{i=1}^{n} \pi_i^*(\mathbf{p})$ by δ

[8] For the definitions of better-reply security and payoff security, please refer to Reny (1999). The reader might ask if Game 1 is better-reply secure. It is not. Consider $n = 2$, $s = 1/10$, $p_1 = 2/5$, and $p_2 = 3/5$ as a counterexample.

at the most, or decrease $\sum_{i=1}^{n} \pi_i^*(\mathbf{p})$, because some $p_i + \delta$ might become higher than $p^r(\mathbf{p}, s)$. Note also that decreasing any element of $\mathbf{p}$ by some small amount $\delta > 0$ will unambiguously decrease $\sum_{i=1}^{n} \pi_i^*(\mathbf{p})$ by two channels: (1) because the average of prices decreases; and (2) because low prices being reduced may cause some of the high prices to become higher than $p^r(\mathbf{p}, s)$. Therefore, a neighborhood $N_{\delta(\varepsilon)}(\mathbf{p})$ always exists such that

$$\sum_{i=1}^{n} \pi_i^*(\mathbf{p}') \leq \sum_{i=1}^{n} \pi_i^*(\mathbf{p}) + \varepsilon \qquad \forall \mathbf{p}' \in N_{\delta(\varepsilon)}(\mathbf{p}).$$

$\therefore \sum_{i=1}^{N} \pi_i^*(\mathbf{p})$ is upper semicontinuous in $\mathbf{p}$ on $\mathcal{P}$. □

Lemma 3. Game 1 satisfies the following property: For all $i = 1, \ldots, n$, $\varepsilon > 0$, $p_i \in \mathcal{P}_i$ and $q_{-i} \in \Delta\mathcal{P}_{-i}$, there exists $\widehat{p}_i \in \mathcal{P}_i$ such that

$$q_{-i}\Big(\{p_{-i} \in \mathcal{P}_{-i} : \pi_i^* \text{ is discontinuous at } (\widehat{p}_i, p_{-i})\} \Big) = 0$$

and $\pi_i^*(\widehat{p}_i, q_{-i}) \geq \pi_i^*(p_i, q_{-i}) - \varepsilon$.

Proof. Note π_i^* is discontinuous in a countable set. Moreover, i can always find a deviation $\widehat{p}_i$ that is slightly worse than p_i but satisfies $\pi_i^*(\widehat{p}_i, q_{-i}) \geq \pi_i^*(p_i, q_{-i}) - \varepsilon$. A deviation such as $\widehat{p}_i \equiv p_i - \delta(\varepsilon)$ works for a well-chosen $\delta(\varepsilon)$. See Figure 1 for reference. Finally, such deviation can always be chosen away from those discontinuity points. □

As the last step, combining Lemma 2 and Proposition 5.1 in Reny (1999), we get that Game 1 is reciprocally upper semicontinuous. Combining Lemma 3 with Theorem 3.33 in Carmona (2013), we get that Game 1 is payoff secure. Finally, by Proposition 3.2 in Reny (1999), Game 1 is better-reply secure.

To complete the existence part of Theorem 3, by Corollary 5.2 in Reny (1999) and better-reply security, Game 1 possesses a mixed-strategy Nash equilibrium.

In the second part of the proof, I show the equilibrium cannot be in *pure* strategies if and only if the conditions of Definition 2 hold.

First, sufficiency. Assume the conditions of Definition 2 hold. I will show that for any $\mathbf{p} = (p_1, \ldots, p_n)$, a firm with incentives to deviate exists. The following lemma has the purpose of narrowing the strategies that may arise in a pure-strategy equilibrium.

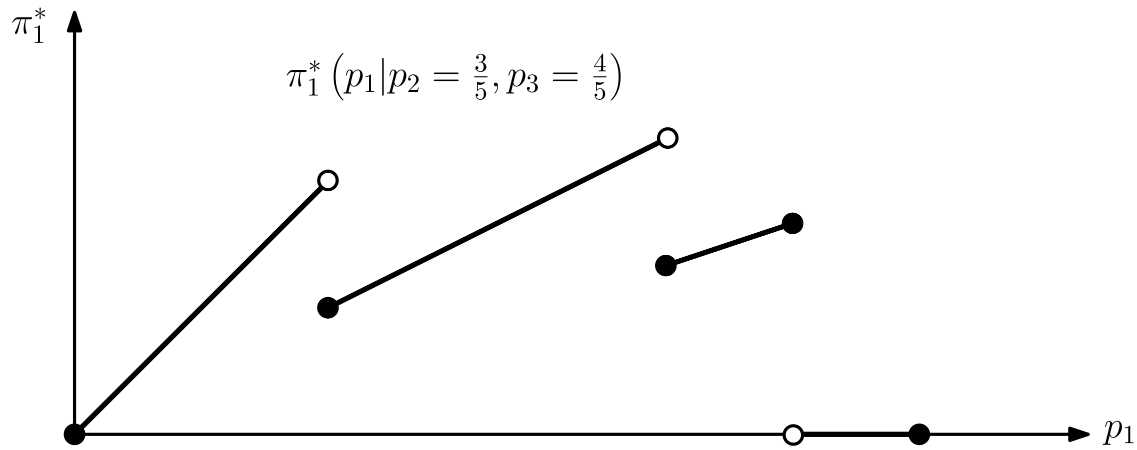

Note: Qualitatively, all payoff functions are the same. Discontinuities are from the left and "jump down" from the left with the exception of the last segment, which becomes irrelevant.

FIGURE 1: Example of a payoff function with three firms and $s = 0.5$.

Lemma 4. Suppose $s > 0$ and $(p_1, \ldots, p_n)$ are such that $p_i > p^r(\mathbf{p}, s)$ for some i. Then, the pure-strategy profile $(p_1, \ldots, p_n)$ is not a Nash equilibrium for Game 1.

Proof. I will propose a deviation p_i^* such that $p_i^* \leq p^r(\mathbf{p}^*, s)$. If $\mathcal{P}_i$ is discrete, because $p_i > p^r(\mathbf{p}, s)$ for some firm i, let $p_i^* = 1/v$ be a deviation. If $\mathcal{P}_i$ is continuous, let $p_i^* = \epsilon$ be a deviation with $0 < \epsilon < s$.

By way of contradiction, suppose $p_i^* > p^r(\mathbf{p}^*, s)$. Then, $p_j \geq p^r(\mathbf{p}^*, s)\ \forall j$, and, by equation (1), we have that $p^r(\mathbf{p}^*, s) = s + p^r(\mathbf{p}^*, s)$, which is a contradiction. Therefore, $p_i^* \leq p^r(\mathbf{p}^*, s)$ and $\pi_i^*(p_i^*, p_{-i}) > \pi_i^*(p_i, p_{-i}) = 0$. □

Because of Lemma 4, it suffices to review the following cases.

Case 1: $p_i < p^r(\mathbf{p}, s)\ \forall i$

- Case 1.1: For some j, i with $j \neq i$, it is true that $p_j = \min_k\{p_k\} < \max_k\{p_k\} = p_i$. Consider the deviation $p_j^* = \max_k\{p_k\} = p_i$. Then, $p_j^* = p_i < p^r(\mathbf{p}, s) \leq p^r(\mathbf{p}^*, s)$, because $p^r(\mathbf{p}, s)$ is nondecreasing in $\mathbf{p}$.
- Case 1.2: $p_i = p_j < 1\ \forall\ i, j$.

– If $\mathcal{P}_j$ is discrete, consider the deviation $p_j^* = p_j + \frac{1}{v}$. By way of contradiction, suppose $p_j^* > p^r(\mathbf{p}^*, s)$. Then, equation (1) implies

$$\begin{aligned} p^r(\mathbf{p}^*, s) &= s + \frac{n-1}{n} p_j + \frac{p^r(\mathbf{p}^*, s)}{n} \\ \Leftrightarrow p^r(\mathbf{p}^*, s)\left(\frac{n-1}{n}\right) &= s + \frac{n-1}{n} p_j \\ \Leftrightarrow p^r(\mathbf{p}^*, s) &= \frac{s}{\frac{n-1}{n}} + p_j \\ \Rightarrow s &< \frac{n-1}{n} \cdot \frac{1}{v}. \qquad \text{(contradiction)} \end{aligned}$$

The contradiction implies $p_j^* \leq p^r(\mathbf{p}^*, s)$. Thus, $\pi_i^*(p_i^*, p_{-i}) > \pi_i^*(p_i, p_{-i})$.

– If $\mathcal{P}_j$ is continuous, consider the deviation $p_j^* = p_j + \epsilon$, such that $0 < \epsilon < s$ and $p_j^* \leq v$. Analogous to the discrete case, equation (1) implies $s < \frac{n-1}{n}\varepsilon$, which is a contradiction. Therefore, $p_j^* \leq p^r(\mathbf{p}^*, s)$ and $\pi_i^*(p_i^*, p_{-i}) > \pi_i^*(p_i, p_{-i})$.

- Case 1.3: $p_i = 1 \; \forall i$.

  I will construct a deviation $p_j^* = p_j - \delta$ with $\delta \in \left(ns, \frac{n-1}{n}\right) \cap \mathcal{P}_j$, such that the following conditions are satisfied:

  1. $p_j^* \leq p^r(\mathbf{p}^*, s)$.
  2. $p_i^* > p^r(\mathbf{p}^*, s) \; \forall \; i \neq j$.
  3. $\pi_j^*(p_j^*) > \pi_j^*(p_j)$.

  Condition 1 is satisfied because $p^r(\mathbf{p}, s)$ is nondecreasing (Lemma 1).

  Condition 2 requires firm j to win the whole market. Suppose, by way of contradiction, that $p_i \leq p^r(\mathbf{p}^*, s)$, for $i \neq j$. Recall that $p_i = 1$ for all i. Then,

$$p^r(\mathbf{p}^*, s) = s + \frac{1}{n}\sum_{i=1}^{n} p_i = s + 1 - \frac{\delta}{n} \geq p_i = 1 \Rightarrow \delta \leq ns. \qquad \text{(contradiction)}$$

  The contradiction implies $p_i^* > p^r(\mathbf{p}^*, s) \; \forall \; i \neq j$.

  Finally, because firm j wins the entire market under the first two conditions, Condition 3 requires the firm to be better off:

$$\pi_j^*(p_j^*) > \pi_j^*(p_j) \Leftrightarrow 1 - \delta > \frac{1}{n} \Leftrightarrow \delta < \frac{n-1}{n},$$

  but due to the hypothesis, such δ exists, and $p_j^* = p_j - \delta \in \mathcal{P}_j$. Therefore, p_j^* is a deviation.

Case 2: $p_i < p^r(\mathbf{p}, s)$ and $p_j = p^r(\mathbf{p}, s)$ for some $i \in I$ and some $j \in J$ such that

$I \cup J = \{1, 2, \ldots, n\}$ and $I \cap J = \emptyset$. For any $i \in I$, consider the deviation $p_i^* = p_j > p_i$. Then, $p_i^* = p^r(\mathbf{p}, s) \leq p^r(\mathbf{p}^*, s)$, which implies $\pi_i^*(p_i^*) > \pi_i^*(p_i)$.

Therefore, Game 1 has no Nash equilibria in pure strategies if the conditions of Definition 2 hold.

Finally, I prove the necessity part of the theorem: assume Game 1 has no Nash equilibria in pure strategies. Then, a profitable deviation always exists. Because $\mathbf{0}$ is not an equilibrium, $s > 0$ for the continuous case, or $s > \frac{n-1}{n} \cdot \frac{1}{v}$ for the discrete case. Moreover, going again through Case 1.3 of the proof and using a contradiction argument, it follows that $\left(ns, \frac{n-1}{n}\right) \cap \mathcal{P}_i \neq \emptyset$. Therefore, the conditions of Definition 2 hold.

As a corollary, it is straightforward to show $s = 0$ implies a Bertrand equilibrium, and the Diamond equilibrium arises if $s > 0$ and the conditions of Definition 2 do not hold. □

3 Equilibrium properties and discussion

In general, the equilibrium has no analytical solution. However, we can use the discrete pricing equilibrium as an arbitrarily good approximation of the continuous-pricing equilibrium by increasing the number of grid points, v. For practical purposes, we can always think of the willingness to pay in cents of a dollar. Formally, we have a *strategic approximation* (Reny, 2011).[9]

Corollary 4. Suppose $\mathcal{P}_i$ is continuous for each i. Then, Game 1 has a strategic approximation given by the discrete-pricing equilibrium.

Proof. The existence of a strategic approximation is immediate given that the game is better-reply secure, and by Theorems 1 and 2 in Reny (2011). Moreover, the discrete-pricing equilibrium is an ε-equilibrium of the continuous-pricing game. To see this fact, note that for a fixed ε, a fine enough grid can approximate any price with arbitrary precision. Because the payoff functions are semicontinuous,[10] some

[9] A strategic approximation of the normal-form game $(\mathcal{P}_i, \pi_i^*)_{i=1}^n$ in which the firms take R^* as given is a countable set of pure strategies $\mathcal{P}^\infty = \mathcal{P}_1^\infty \times \cdots \times \mathcal{P}_n^\infty$ contained in $\mathcal{P} = \mathcal{P}_1 \times \cdots \times \mathcal{P}_n$, such that whenever for each player i, $\mathcal{P}_i^1 \subseteq \mathcal{P}_i^2 \subseteq \cdots$ is an increasing sequence of finite subsets of $\mathcal{P}_i$ whose union contains $\mathcal{P}_i^\infty$, any limit of equilibria of the sequence of finite games $(\mathcal{P}_i^1, \pi_i^*)_{i=1}^n, (\mathcal{P}_i^2, \pi_i^*)_{i=1}^n, \ldots$ is an equilibrium of $(\mathcal{P}_i, \pi_i^*)_{i=1}^n$.

[10] The payoffs are upper semicontinuous in own prices over some segments, and lower semicon-

$v(\varepsilon)$ large enough exists such that a price in the grid ε-approximates, either from above or below, the payoff of any price in [0,1]. Finally, limits of ε-equilibria are equilibria of the continuous-pricing game, because the game is better-reply secure (Reny, 1999). □

I henceforth focus on the discrete-pricing equilibrium, which is characterized by Corollary 7 in the appendix in a system of nonlinear equations.

Consider the following example. Let an economy consist of two firms, and $\mathcal{P}_i = \left\{0, \frac{1}{5}, \frac{2}{5}, \frac{3}{5}, \frac{4}{5}, 1\right\}$, that is, discrete pricing with $n = 2$ and $v = 5$. Assume $s = 1/10$, which meets the conditions of Definition 2. Therefore, no pure-strategy equilibria exist, but it can be shown that multiple symmetric mixed Nash equilibria exist.[11]

Figure 2 shows the distribution of prices played with positive probability when the grid increases from $v = 5$ to 10, 25, and 50 points. Remarkably, the equilibrium price distribution is bimodal. In this example, firms charge high prices almost all the time, but, occasionally, we have a sale with a 50% discount.

In general, some basic properties of the equilibrium can be obtained, even without analytical solutions. First, the price distribution shifts to the right when the search cost increases, yielding higher profits for firms, and lower surplus for the consumer. Second, holding fixed the absolute size of the search cost, increasing the number of firms results in monopolistic pricing. The following corollaries formalize the results. The discussion that follows illustrates them.

Corollary 5. In the symmetric equilibrium, the expected price is nondecreasing in the search cost. Moreover, the equilibrium payoff of the firms is also nondecreasing in the search cost, whereas the consumer surplus is nonincreasing in the search cost.

Proof. From Lemma 1, the reservation price is nondecreasing in the search cost. Then, from equation (2), $\pi_i(\mathbf{p}, R^*)$ is nonincreasing in s when $p_i \leq p^r(\mathbf{p}, s)$ and $p_{-i} > p^r(\mathbf{p}, s)$. Analogously, $\pi_i(\mathbf{p}, R^*)$ is nondecreasing in s in any other case. A Nash equilibrium requires i to be indifferent between choosing p_i and p_i' whenever

tinuous over other segments (no isolated points are on the graph; see Figure 1). Moreover, because $p^r(\mathbf{p}, s)$ is continuous, the same is true with respect to the other firms' prices.

[11]This example has three Nash equilibria: playing $\left\{\frac{2}{5}, \frac{3}{5}, \frac{4}{5}\right\}$ with probabilities $\left(\frac{1}{4}, \frac{1}{4}, \frac{1}{2}\right)$; $\left\{\frac{3}{5}, \frac{4}{5}, 1\right\}$ with probabilities $\left(\frac{1}{5}, \frac{7}{15}, \frac{1}{3}\right)$; and, $\left\{\frac{2}{5}, \frac{3}{5}, \frac{4}{5}, 1\right\}$ with probabilities $\left(\frac{1}{6}, \frac{1}{6}, \frac{5}{9}, \frac{1}{9}\right)$. These equilibria can be verified in the usual manner. This example was first presented in García P. and Tudón M. (2010).

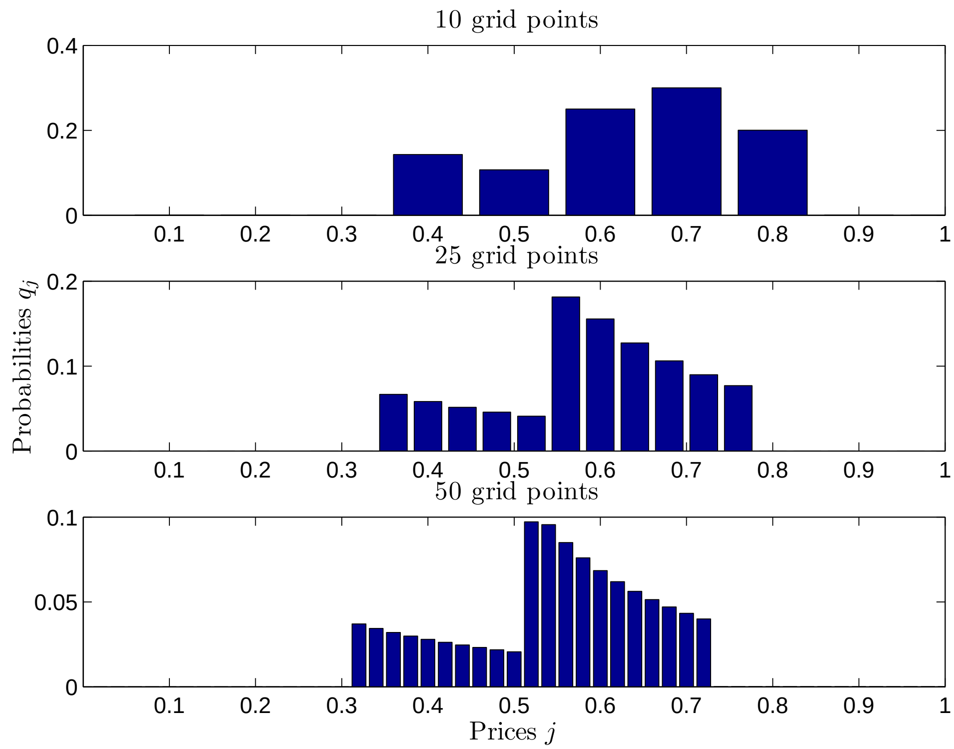

FIGURE 2: Equilibrium distributions with two firms and $s = 1/10$ for grids of size 10, 25, and 50.

both p_i and p'_i have positive weights in the equilibrium mixed strategy. It follows that strategies yielding $p_i \leq p^r(\mathbf{p}, s)$ and $p_{-i} > p^r(\mathbf{p}, s)$ must receive zero weight when the search cost is high enough. Then, the equilibrium weights shift to the right as s increases. Then, the expected price is nondecreasing. Because the consumer faces weakly higher prices, her surplus weakly decreases. □

Corollary 6. For a fixed $s > 0$, as $n \to \infty$, the unique equilibrium is the Diamond equilibrium. That is, $p_i = 1$, for all i.

Proof. Follows immediately from Theorem 3. □

The next figures illustrate the effects of increasing the search cost and entry through simulation. I only consider sets of parameters that create price dispersion, because the discussion becomes trivial otherwise. All the results are qualitatively robust to different specifications of the parameters, and are selected for clarity of exposition.

Figure 3 considers different search costs when the grid size $v = 10$, and when the number of firms $n = 2$. Increasing the search cost has a clear effect: a higher search cost moves the price distribution to the right, where firms benefit from more monopolistic power.

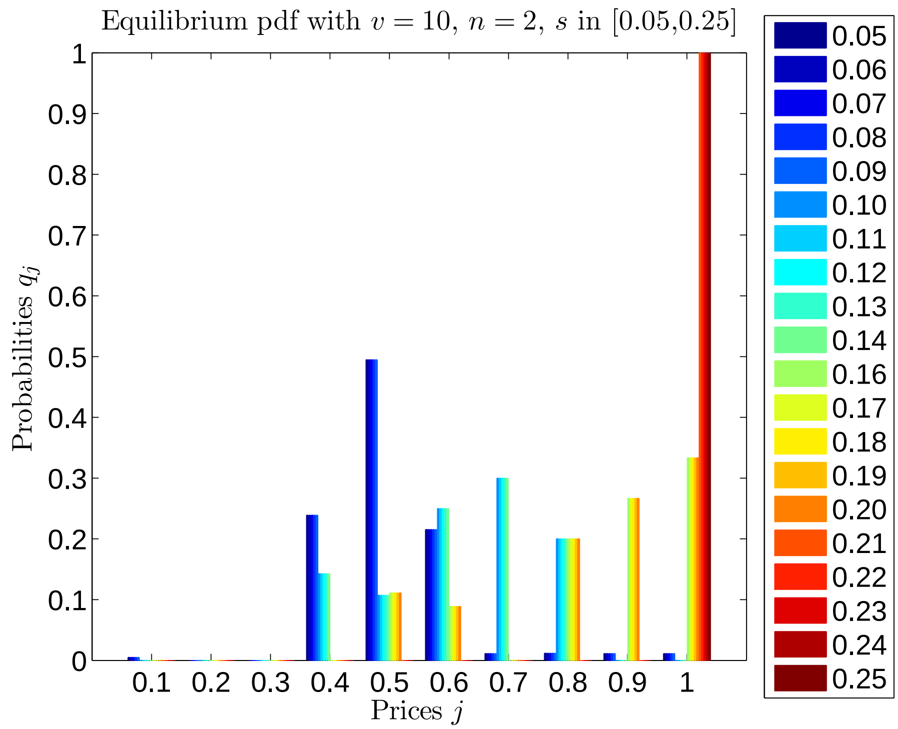

FIGURE 3: Comparative statics on the search cost.

In a simple model such as this one, the total surplus from the Bertrand and Diamond equilibria are the same, because no search exists in either equilibrium. Either the consumer or the producer extract all surplus, and no dead-weight loss exists. However, in the dispersed equilibrium, the search cost must be subtracted from the consumer surplus, which in turn decreases total surplus.

The producer surplus, π, is obtained with Corollary 7. Once the equilibrium distribution is obtained, one can construct the consumer surplus by Monte Carlo simulation.[12] Finally, total surplus is equal to the consumer surplus plus n times the producer surplus. Figure 4 summarizes these observations where the price discreteness causes the five-points pattern in which firms use the same strategy. As

[12]Simulate draws of equilibrium prices, construct the reservation price, and calculate the expected number of searches and the expected price paid. The Matlab codes used for these calculations, and for

expected, the consumer is worse off when the search cost increases, but the producer is better off. Moreover, the total surplus shows the dead-weight loss due to search for intermediate search-cost levels.

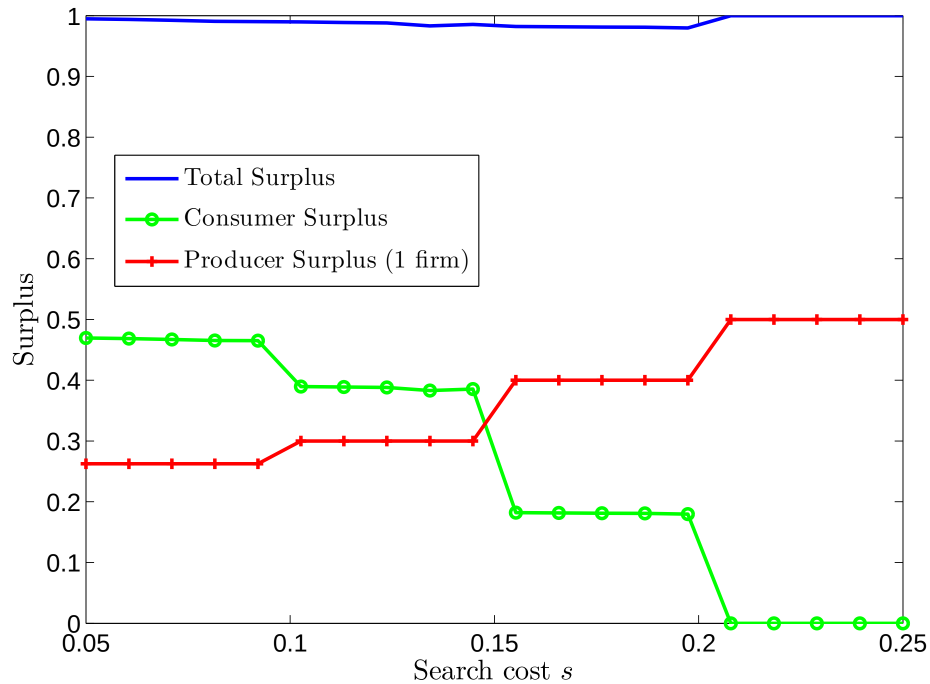

FIGURE 4: Welfare analysis. $v = 10$, $n = 2$ and $s \in [0.5, 2.5)$.

Finally, Figure 5 shows an example of the effect of increasing the number of firms. As more firms enter the market, the average price and the variance increase. Additionally, the equilibrium distribution shows nonmonotonic behavior on kurtosis, which measures the tails of the distribution. In this case, the distribution has fatter tails with three firms but is less dispersed than with four firms.

Admittedly, the noncompetitive effect of increasing the number of firms is a limitation but is not new to search models; see Robert and Stahl II (1993), Stahl II (1989) or Rosenthal (1980), for example. The probability of being the lowest-priced store decreases exponentially with entry, disrupting the ability of firms to steal the market. However, whereas the condition $ns < (n-1)/n$ becomes rapidly stringent as the number of firms in the market grows, real-life examples show search costs can be within a Goldilocks zone. De los Santos, Hortaçsu and Wildenbeest (2013) estimate

those of section 3, are available upon request or at jtudon.com.

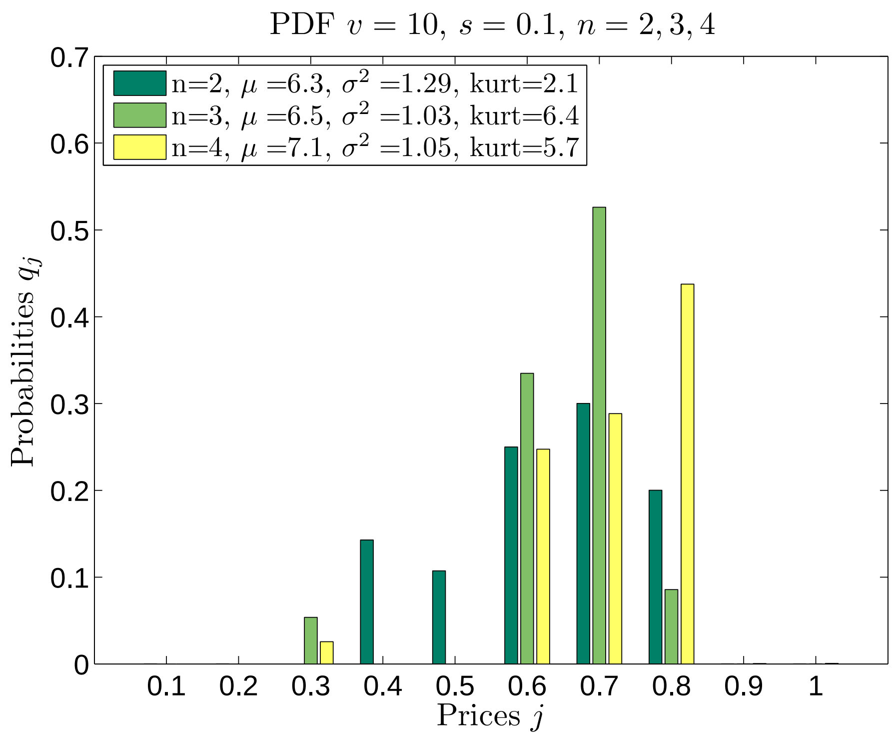

FIGURE 5: Increasing the number of firms

the search costs of consumers buying MP3 players. The authors consider up to 10 different firms offering several devices whose prices range considerably. Though their baseline model is different, their results show the constraint is plausible.[13] Moreover, note the result is not driven by a "small" n: even for two firms, if the firms and the consumer act simultaneously, the unique equilibrium is the Diamond paradox.

On the other hand, if the search cost decreases with the number of firms in the market, one could derive a pro competitive effect of entry. For instance, suppose $s_n \propto 1/n$. In a spatial model, one could think of business clusters where firms concentrate information in a certain area, thus decreasing the search cost per firm. In such an extension, price dispersion can be sustained even as $n \to \infty$.

Finally, note consumer search is not trivial. The expected number of searches is not limited to one or two, because realized prices can be above the reservation price.

[13]Specifically, from Table 4 in De los Santos, Hortaçsu and Wildenbeest (2013) one can obtain a lower bound on the willingness to pay, and an upper bound on the marginal cost by looking at the maximum and minimum price. Then, from their estimated search costs, we can conclude that some Goldilocks conditions hold for most of their products.

Possibly, consumers search for several periods, in contrast to previous approaches such as Burdett and Judd (1983).

4 Concluding remarks

The strongest assumption of this paper regards how the consumer knows the price distribution, and why firms do not react to it. Alternatively, suppose consumers search from an unknown distribution of which they have an uninformative Dirichlet prior. Rothschild (1974) showed the optimal search rule will have a reservation price that satisfies properties analogous to those of Lemma 1. Therefore, no great loss results from making such a simplifying assumption toward accomplishing the goal of this paper. Although searching from a known distribution is a strong assumption, one can only expect that relaxing the assumption would reinforce the result rather than reverse it.[14]

On the other hand, firms do not react to prices, because they move simultaneously. Yet, consumers do "react" because they move later in the game. Thus, I assume an informational asymmetry. If firms are allowed to react, the issue is commitment as in Daughety (1992). However, if the consumer is prohibited to react, we return to the Diamond setup. A more symmetric informational assumption is that neither firms nor the consumer know the distribution of prices, which implies firms cannot react to other prices. But as I discussed above, we can approximate the consumer's behavior in the second stage by assuming she knows the price distribution.

On the plus side, the model offers a novel existence theorem that does not depend on exogenous heterogeneity or stochastic shocks. With minimal assumptions, this paper produces nontrivial consumer search and price dispersion. Finally, by proving information frictions are sufficient, and not only necessary for price dispersion, we can now think of ex-ante heterogeneity as a modeling tool and not as a necessity.

References

Baye, Michael R, and John Morgan. 1999. "A folk theorem for one-shot Bertrand games." *Economics Letters*, 65(1): 59–65.

[14]Indeed, using some uninformative prior will assume away the Diamond paradox since consumers will search more than once. Parakhonyak and Sobolev (2015) consider search without priors.

Billingsley, Patrick. 1999. *Convergence of Probability Measures.* Wiley-Interscience.

Burdett, Kenneth, and Kenneth L. Judd. 1983. "Equilibrium price dispersion." *Econometrica*, 51(4): 955–969.

Burdett, Kenneth, and Tara Vishwanath. 1988. "Declining Reservation Wages and Learning." *The Review of Economic Studies*, 55(4): 655.

Carlson, John A., and R. Preston Mcafee. 1983. "Discrete Equilibrium Price Dispersion." *Journal of Political Economy*, 91(3): 480–493.

Carmona, Guilherme. 2013. *Existence and Stability of Nash Equilibrium.* World Scientific Publishing Company.

Daughety, Andrew F. 1992. "A model of search and shopping by homogeneous customers without price precommitment by firms." *Journal of Economics and Management Strategy*, 1(3): 455–73.

Daughety, Andrew F., and Jennifer F. Reinganum. 1992. "Search equilibrium with endogenous recall." *RAND Journal of Economics*, 23(2): 184–202.

De los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest. 2013. "Search with Learning."

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

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

García P., José A., and José F. Tudón M. 2010. "Escape a la paradoja de Diamond." *Gaceta de Economía*, 28(1): 185–200.

Hortaçsu, Ali, and Chad Syverson. 2004. "Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds." *The Quarterly Journal of Economics*, 119(2): 403–456.

Janssen, Maarten C.W., José Luis Moraga-González, and Matthijs R. Wildenbeest. 2005. "Truly costly sequential search and oligopolistic pricing." *International Journal of Industrial Organization*, 23(5-6): 451–466.

Janssen, Maarten, Paul Pichler, and Simon Weidenholzer. 2011. "Oligopolistic markets with sequential search and production cost uncertainty." *RAND Journal of Economics*, 42(3): 444–470.

Jehle, Geoffrey, and Philip Reny. 2011. *Advanced Microeconomic Theory.* Prentice Hall.

Kaplan, Greg, and Guido Menzio. 2015. "The morphology of price dispersion." *International Economic Review*, 56(4): 1165–1206.

Kaplan, Todd R., and David Wettstein. 2000. "The possibility of mixed-strategy equilibria with constant-returns-to-scale technology under Bertrand competition." *Spanish Economic Review*, 2: 65–71.

McCall, John J. 1970. "Economics of Information and Job Search." *The Quarterly Journal of Economics*, 84(1): 113–126.

Menzio, Guido, and Nicholas Trachter. 2015. "Equilibrium price dispersion with sequential search." *Journal of Economic Theory*, 160: 188–215.

Nash, John F. 1950. "Equilibrium points in n-person games." *Proceedings of the National Academy of Sciences*, 36(1): 48–49.

Parakhonyak, Alexei, and Anton Sobolev. 2015. "Non-Reservation Price Equilibrium and Search without Priors." *The Economic Journal*, 125(584): 887–909.

Pratt, John W. 1979. "Price differences in almost competitive markets." *The Quarterly Journal of Economics*, 93(2): 189–211.

Rauh, Michael T. 2007. "Nonstandard foundations of equilibrium search models." *Journal of Economic Theory*, 132(1): 518–529.

Reinganum, Jennifer. 1979. "A Simple Model of Equilibrium Price Dispersion." *Journal of Political Economy*, 87(4): 851–858.

Reny, Philip. 1999. "On the existence of pure and mixed strategy Nash equilibria in discontinuous games." *Econometrica*, 67(5): 1029–1056.

Reny, Philip. 2011. "Strategic approximations of discontinuous games." *Economic Theory*, 17–29.

Robert, Jacques, and Dale O Stahl II. 1993. "Informative price advertising in a sequential search model." *Econometrica*, 61(3): 657–686.

Rob, Rafael. 1985. "Equilibrium price distributions." *The Review of Economic Studies*, 52(3): 487–504.

Rosenthal, Robert W. 1980. "A model in which an increase in the number of sellers leads to a higher price." *Econometrica*, 48(6): 1575–1579.

Rothschild, Michael. 1974. "Searching for the lowest price when the distribution of prices is unknown." *The Journal of Political Economy*, 82(4): 689–711.

Sorensen, Alan. 2000. "Equilibrium price dispersion in retail markets for prescription drugs." *Journal of Political Economy*, 108(4): 833–850.

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

Stahl II, Dale O. 1996. "Oligopolistic pricing with heterogeneous consumer search." *International Journal of Industrial Organization*, 14(2): 243–268.

Stigler, George. 1961. "The Economics of Information." *Journal of Political Economy*, 69(3): 213–225.

Stiglitz, Joseph E. 1987. "Competition and the number of firms in a market: Are duopolies more competitive than atomistic markets?" *Journal of Political Economy*, 1041–1061.

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

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

A Characterization of the discrete-pricing equilibrium

Corollary 7. (Symmetric equilibrium) Let $\mathcal{P}_i$ be discrete for all i. Define q_j^i as the symmetric equilibrium probabilities of firm i choosing the jth price in the set $\mathcal{P}_i$, and π as the symmetric equilibrium payoff for all firms. Then, the mixed strategy $(q_1, \ldots, q_v)$ for all firms is a Nash equilibrium of Game 1 if and only if $(q_1, \ldots, q_v, \pi)$ is the solution of the following system of nonlinear equations:

$$\left[\sum_{k_2=1}^{v} \cdots \sum_{k_n=1}^{v} q_{k_2} \cdots q_{k_n} \pi^* \left(\frac{1}{j}, \frac{1}{k_2}, \ldots, \frac{1}{k_n}, R^*\right) - \pi\right] q_j = 0 \quad j = 1, \ldots, v$$

$$\sum_{k_2=1}^{v} \cdots \sum_{k_n=1}^{v} q_{k_2} \cdots q_{k_n} \pi^* \left(\frac{1}{j}, \frac{1}{k_2}, \ldots, \frac{1}{k_n}, R^*\right) \leq \pi \quad j = 1, \ldots, v$$

$$q_j \geq 0 \quad j = 1, \ldots, v$$

$$\sum_{j=1}^{v} q_j = 1, \tag{3}$$

with $\pi^*\left(\frac{1}{j}, \ldots, R^*\right)$ as defined in equation (2).

Proof. For example, see Theorem 7.1 in Jehle and Reny (2011).