

Search, Information, and Prices

Dirk Bergemann

Yale University

Benjamin Brooks

University of Chicago

Stephen Morris

Massachusetts Institute of Technology

Consider a market with identical firms offering a homogeneous good. For any given ex ante distribution of the price count (the number of firms from which a consumer obtains a quote), we derive a tight upper bound on the equilibrium distribution of sales prices. The bound holds across all models of firms' common-prior higher-order beliefs about the price count, including the extreme cases of full information and no information. One implication of our results is that a small ex ante probability that the price count is equal to one can lead to a large increase in the expected price. The bound also applies in a large class of models where the price count distribution is endogenously determined.

I. Introduction

When two or more identical firms engage in Bertrand competition, the standard prediction of economic theory is that the price of a homogenous

We acknowledge financial support through National Science Foundation grants SES-1459899 and SES-2001208. We thank Roi Orzach for valuable research assistance and Mark Armstrong, Glenn Ellison, Ben Golub, Jason Hartline, Ali Hortaçsu, and Phil Reny for helpful conversations, as well as seminar audiences at Université de Montréal, Stanford University, and the University of Chicago.

Electronically published May 26, 2021

[*Journal of Political Economy*, 2021, vol. 129, no. 8]

© 2021 by The University of Chicago. All rights reserved. 0022-3808/2021/12908-0003\$10.00

good will be competed down to cost. The model of Bertrand competition is therefore consistent with the "law of one price." But as Varian (1980, 651) noted 40 years ago, "the law of one price is no law at all," and price dispersion, even for homogenous goods, seems to be a ubiquitous feature of the market economy.¹

Since the time of Varian's writing, a large literature on imperfect price competition and consumer search has developed numerous models of equilibrium price dispersion. A central object in these models is the number of price quotes that a consumer receives, which we refer to as the *price count*. The price count and the firms' strategies induce a distribution over the *sales price*, that is, the price at which trade occurs. The sales price (or equivalently transaction price) is the lowest price offered by the quoted firms. This distribution of the sales price determines the consumer and producer welfare. In the classical Bertrand model, firms have full information about the price count, which leads to the law of one price. In models with price dispersion, the noise in prices is driven by firms' uncertainty about the price count and, in particular, the failure of common knowledge of whether there are at least two firms competing for a given consumer. Many explanations have been proposed for this failure of common knowledge, including unobserved consumer search, advertising, and informational frictions due to market intermediaries.

The equilibrium price distribution depends on two critical features of the model: the distribution of the price count and the firms' information about the price count.² Significantly, most existing models make a strong simplifying assumption about firms' information, namely, that firms have no information beyond the *ex ante* distribution of the number of firms competing for a given consumer. While the no-information assumption leads to a nondegenerate price distribution in equilibrium, it places strong restrictions on the range of welfare outcomes that can arise. In particular, in symmetric and simultaneous-search models, the no-information assumption implies that the expected sales price is the same as it would be if firms had full information about the number of competing firms. Thus, while these models are consistent with failure of the law of one price, they predict the same welfare outcome as if the law of one price held.

¹ A substantial empirical literature has studied cross-sectional price dispersion; see Baye, Morgan, and Scholten (2006) for an early survey. A conclusion of this literature is that price dispersion has persisted and sometimes increased in the internet age.

² These objects are sufficient to determine the equilibrium price distribution in models in which firms' realized prices do not directly affect price counts or other firms' prices, such as in models of simultaneous consumer search. If search is sequential or if firms play a Stackelberg game, then the equilibrium price distribution can depend on further strategic considerations. For example, in the sequential search model of Stahl (1989), prices are lower than in a simultaneous-search model with the same price count distribution. This is discussed in greater detail in sec. V.

This paper develops new predictions for equilibrium prices that do not depend on how the price count is determined or on firms' information about price counts. In particular, we take as given a particular distribution of the price count. (We subsequently give conditions for our results to hold when price counts are endogenized.) At the same time, we are agnostic about firms' information, and we consider the full range of equilibrium outcomes that might obtain for all common-prior beliefs that firms might have about the price count. Our main result, theorem 1, is a tight upper bound on the equilibrium sales price distribution, in the sense of first-order stochastic dominance.³ This bound holds across all informational models of firms' beliefs that are consistent with the given price count distribution. This theorem immediately implies a tight upper bound on producer surplus and a tight lower bound on consumer surplus.

The bound we construct is based on the following logic. If the sales price distribution were too high—for example, if all firms priced at the monopoly level—then firms would obviously have an incentive to undercut and thereby gain more sales. This suggests that there are nontrivial bounds on how high the sales price distribution can go and that the critical equilibrium constraints are those associated with cutting prices. We focus on a particular class of such deviations, wherein, for some fixed price, a firm deviates by setting the minimum of that fixed price and whatever price they would have set in equilibrium. We refer to this as a *uniform price cut*. We show that the requirement that firms do not want to adopt a uniform price cut can be expressed as a constraint on the sales price distribution. We further show that there is a highest sales price distribution that satisfies all of the uniform price cut incentive constraints and is associated with firms being indifferent to all price cuts (uniform or otherwise).

To show that the bound is tight, we explicitly construct a model of beliefs and equilibrium pricing strategies for firms that attain the bound. The critical beliefs are induced by signals of the following form. Each firm observes a positive integer that is a lower bound on the realized price count, and at least one firm's signal is equal to the true price count. An interpretation of this information structure is that firms are quoted in a random order, and each firm observes a subset of the firms that were quoted before them, with firms that are quoted last seeing all other quoted firms. In equilibrium, firms randomize prices over intervals that are decreasing in the number of other firms they observe, so that the sale is always

³ Our analysis is focused on the distribution of the sales price, rather than the distribution of prices posted by firms. Posted prices are necessarily higher than the sales price and can be considerably so. For example, if there is common knowledge that there are at least three firms, then there is an equilibrium outcome in which two firms price at cost and the remaining firms set arbitrary higher prices. The study of informationally robust predictions for posted prices will require new assumptions or analytical techniques and is an important direction for future work.

made by a firm that observes the true price count. This information structure pushes prices up, because firms that think that the price count is low price high and firms that know that the price count is higher are willing to price higher because they anticipate that other firms believe that the price count is lower. By carefully tuning the distribution of firms' signals, it is possible to make firms indifferent to all price cuts, thereby attaining the maximal sales price distribution.

Theorem 1 can be interpreted as providing an empirical test for collusion: if the observed sales price distribution is not below the theoretical bound generated by the observed price count distribution, then prices cannot be explained by Bertrand competition under incomplete information. This test for collusion is informationally robust, as its validity does not depend on the nature of the information held by the market participants.⁴ As the bound is expressed in terms of first-order stochastic dominance, it simultaneously contains information about many moments of the price distribution. This is a feature shared with previously suggested screens for collusion, such as Abrantes-Metz et al. (2006), where competition is identified by lower means and higher variances in the price distribution.⁵

Theorem 1 also gives a global upper bound on the effect of monopoly power on prices, as we depart from the benchmark of perfect competition: if the probability of monopoly (i.e., a price count of one) is μ , then revenue can reach a proportion $[\mu(2 - \mu)]^{1/2}$ of monopoly revenue. Thus, if we allow for firms to have partial information about the price count, producer surplus is nonlinear in the probability of monopoly, and in fact, marginal revenue in the probability of monopoly is unbounded at $\mu = 0$. Thus, a small amount of monopoly power may translate into rents for firms that are much larger than what would obtain under the benchmarks of no information and full information, in which revenue is linear in μ . This finding complements the conclusion of Diamond (1971) and others that small search frictions can translate into a large degree of monopoly power.

Theorem 1 takes the price count distribution as given. A critical question is whether our bounds remain valid when price counts are endogenously determined in equilibrium. Among the various models of price count formation that have been suggested in the literature, it is useful to distinguish two classes. First, there are models in which price counts and firms' prices depend on expectations of firms' equilibrium pricing behavior but price counts and other firms' prices do not react when a firm deviates from their

⁴ Of course, the bound on the sales price distribution cannot account for features that are outside of our model, such as bounded rationality, product differentiation, or heterogeneous production costs.

⁵ By contrast, Bajari and Ye (2003) use notions of statistical independence to distinguish between competition and collusion in bidding games. In our setting, these notions could not distinguish between competition and collusion, as private information can lead to correlated prices even under competition.

equilibrium strategy. In this case, we say that the model has no feedback. This category includes any model in which price counts and prices are simultaneously determined, such as Butters (1977), Varian (1980), Burdett and Judd (1983), Baye and Morgan (2001), Ellison and Wolitzky (2012), and de Clippel, Eliaz, and Rozen (2014). Our bounds immediately apply to any model with no feedback. In contrast, models with feedback have the feature that firms' deviations can affect price counts or other firms' prices. Whether the bounds hold on the presence of feedback depends on details of the model. The bounds need not be satisfied in Stackelberg games, wherein firms price sequentially and observe past prices. In contrast, theorem 2 shows that the bounds must be satisfied in models of sequential search, such as Stahl (1989, 1996), in which a consumer solicits price quotes one at a time and the decision of whether to solicit more quotes depends on past prices. The reason is that consumer search makes price cuts more attractive, compared to the benchmark with fixed price counts, and this reinforces the logic that leads to the bound in theorem 1.

Our main results concern a model where consumers have unit demand and firms are all equally likely to be quoted, conditional on the price count. Both assumptions can be relaxed. First, we generalize the upper bound on sales prices to the case where consumers have downward-sloping demand. Second, we argue that the bounds still hold if firms are heterogeneous in their probabilities of being quoted, although the bound is not necessarily tight.

We interpret our results as applying to a market in which many firms sell to a single consumer with an unknown price count. A mathematically equivalent interpretation is that there is a continuum of consumers with heterogeneous price counts, and firms can imperfectly price discriminate on the basis of characteristics that are statistically linked to the price count. Armstrong and Vickers (2019) study such a model, with an emphasis on the case where quotation probabilities are asymmetric across firms. They compare welfare under price discrimination with what would obtain if firms had to set the same price for all consumers. Myatt and Ronayne (2019) offer a two-stage version of the shopper and captive-consumer model that attains price dispersion in pure strategy equilibria and thus "stable" price dispersion.

As mentioned above, theorems 1 and 2 could be used to test whether the empirical distributions of price counts and sales prices are consistent with Bertrand competition under incomplete information. This exercise would complement the approach taken by Hong and Shum (2006), who postulate that prices are generated by an equilibrium under the standard sequential consumer search model with constant cost per search and no information. Under these assumptions, they show that it is possible to identify the price count distribution and consumers' search costs from the empirical price distribution. The virtue of our approach is that it

makes weaker assumptions about firms' information, although we no longer have a one-to-one relationship between the price count and sales price distributions. There are two challenges in bringing our results to the data. First, we assume throughout that goods are homogenous. While this is an important benchmark, there are many markets of interest in which firms offer heterogeneous goods, and generalizing our results to allow for such heterogeneity is an important direction for future work. Second, in order to compute our bounds, the analyst needs to have data on the price count distribution. Recent work by De los Santos, Hortaçsu, and Wildenbeest (2012) has leveraged browser history data to provide such direct evidence on price counts.

We also contribute to a growing literature on informationally robust predictions in Bayesian games (Bergemann and Morris 2013, 2016). Bergemann, Brooks, and Morris (2017) have applied a similar methodology to first-price auctions, where bidders do not necessarily know their values of the object being sold. We discuss the connection to this literature in greater detail in section VI.

The rest of this paper proceeds as follows. Section II describes our model. Section III presents a two-firm example that illustrates our results. Section IV describes our main result. Section V extends the analysis to sequential search. Section VI discusses further extensions and the connection to the literature on first-price auctions. Section VII is a conclusion.

II. Model

A single consumer has a willingness to pay $v > 0$ for a single unit of a homogeneous good.⁶ There are n firms, indexed by $i \in N \triangleq \{1, \dots, n\}$, who can produce the good at zero cost. The consumer receives price quotes from a subset K of those firms. The price count is the number of price quotes $k = |K|$ that the consumer receives. We write $\mu \in \Delta(N)$ for the ex ante distribution of the price count. Given a price count k , the conditional probability that the firms $K \subseteq N$ are quoted is denoted $\nu(K|k)$, where $\nu(K|k) > 0$ only if $|K| = k$. We assume that all firms are equally likely to be quoted conditional on k ; that is, for all $i \in N$, $\sum_{\{K \subseteq N | i \in K\}} \nu(K|k) = k/n$.⁷

We focus on the single-consumer interpretation of our model, but as mentioned in the introduction, there is an alternative interpretation in which there is a continuum of consumers and $\mu(k)$ is the proportion of

⁶ The value $v > 0$ plays no specific role in the analysis and could be normalized to 1. We have made it explicit merely to clarify the units in our formulas. The assumption of single-unit demand can be easily relaxed, and in sec. VI we report how the analysis extends to a setting with general downward-sloping demand for a homogenous good.

⁷ An earlier version of this paper made the stronger assumption that conditional on k , all sets $K \subseteq N$ with $|K| = k$ were equally likely. We thank Mark Armstrong for suggesting this weaker condition.

consumers who obtain k price quotes. We reference this interpretation occasionally in discussing our results.

An *information structure* (T, π) consists of measurable sets of signals T_i for each firm and a mapping π that associates with each set of quoted firms a joint probability over their signals. More specifically, when the set of quoted firms is K , the quoted firm observe signals $t_K = (t_i)_{i \in K}$. Each firm's signal represents characteristics of the consumer that are informative about the number of price quotes obtained by the consumer. The distribution of these signals is given by the joint probability measure $\pi(dt_K|K)$ on the measurable set $T_K = \prod_{i \in K} T_i$. Note that firms that are not quoted do not receive signals (as the firm is not active). This is without loss; see footnote 8 below.

Given an information structure, firms choose prices conditional on their signals. Thus, each firm is quoting a price only when prompted by a signal. In particular, no firm offers a standing or posted price independent of any signal. The pricing strategy of firm i , conditional on observing signal t_i , is described by the likelihood that firm i sets a price p_i greater than x :

$$F_i(x|t_i) \triangleq \Pr(p_i \geq x|t_i).$$

We use such upper cumulative distribution functions (also known as decumulative functions) for all price distributions throughout the paper. We let $F_i(dp_i|t_i)$ denote the measure over firm i 's price and $F_K(dp_K|t_K)$ denote the independent joint measure over prices of firms in K , given their respective signals t_K .⁸ When clear from the context, we write $F_K(dp|t)$ for $F_K(dp_K|t_K)$, and similarly we write $\pi(t|K)$ for $\pi(t_K|K)$ and simply drop the subscript when referring to the entire set N , so that $F(dp|t)$ represents $F_N(dp_N|t_N)$.

The consumer will buy from one of the firms offering the lowest price, with ties broken uniformly.⁹ Given a realized tuple of prices $p \in \mathbb{R}^K$ among a set of firms K , let $K(p)$ be the set of firms that offer the lowest price:

$$K(p) \triangleq \arg \min_{i \in N} p_i.$$

The revenue of firm i at given price profile $p \in \mathbb{R}^K$ is $p_i \mathbb{I}_{i \in K(p)} / |K(p)|$, where the indicator function $\mathbb{I}_{i \in K(p)} = 1$ if firm i is among the selling firms and is

⁸ In the definition of the information structure, we specify only the distribution of signals for firms that are quoted. We could have made explicit the distribution of signals for firms that are not quoted. But since the firms' payoff are always zero when they are not quoted, it is without loss for firms to condition on the event that they are quoted when determining the price they should set. As a result, the distribution of signals for firms that are not quoted is strategically irrelevant and can be omitted in our notation.

⁹ The uniform tie-breaking assumption is for simplicity of exposition. An asymmetric tie-breaking rule will not alter the fundamental inequality (9) that drives our results.

0 otherwise. Given the strategy profile $F = (F_1, \dots, F_n)$, the expected revenue of firm i is

$$R_i(F) = \sum_{k=1}^n \mu(k) \sum_{K \subseteq N} \nu(K|k) \left(\int_{t \in T_K} \int_{p \in [0, v]^K} p_i \frac{\mathbb{I}_{i \in K(p)}}{|K(p)|} F_K(dp|t) \pi(dt|K) \right). \quad (1)$$

The strategy profile F is a (*Bayes-Nash*) *equilibrium* if and only if $R_i(F) \geq R_i(F'_i, F_{-i})$ for each i and strategy F'_i .

III. A Two-Firm Example

We first illustrate our approach and results for the case of two firms, $n = 2$. We normalize $v = 1$ and let the price count be 1 with probability μ and 2 with probability $1 - \mu$. The consumer collects a single (monopoly) quote with probability $\mu \in (0, 1)$ and two (competitive) quotes with probability $1 - \mu$. Thus, the consumer gets a quote from only firm 1 with probability $\mu/2$, from only firm 2 with probability $\mu/2$, and from both firms with probability $1 - \mu$. In the continuum interpretation of the model, a proportion $\mu/2$ are “captive” consumers of firm 1, a proportion $\mu/2$ are captives of firm 2, and a proportion μ of consumers are “contested” and can purchase from either firm.¹⁰

A. Full Information

First, suppose that there is full information about the price count. If there is one quote, the quoted firm is a monopolist and charges the monopoly price of 1. If there are two quotes, both firms charge the competitive price of 0. Thus, the sales price is 1 with probability μ , and it is 0 with probability $1 - \mu$. The ex ante (upper cumulative) distribution of the sales price is denoted by $S(\cdot)$; that is, $S(x)$ is the probability that the lowest price is at least x . This function is depicted as the thin solid curve in figure 3 for $\mu = 1/2$: the price is at least 0 with probability 1, and the price is at least x for any $0 < x \leq 1$ with probability $1/2$. In the continuum interpretation, this corresponds to the case where the firms can see whether a customer is captive or contested and price discriminate accordingly.

B. No Information

Now suppose that the firms have no information about the price count. In the continuum interpretation, this corresponds to the assumption that firms cannot price discriminate and must offer a uniform price. A firm asked to quote a price will therefore assign conditional probability

¹⁰ This terminology follows Armstrong and Vickers (2019).

$$\frac{\mu/2}{\mu/2 + (1 - \mu)} = \frac{\mu}{2 - \mu}$$

to being the monopolist.

This model has a unique mixed-strategy equilibrium where firms randomize over prices in the interval $[\mu/(2 - \mu), 1]$. Both firms use the same mixing probabilities, wherein the price is at least $p_i \in [\mu/(2 - \mu), 1]$ with probability

$$F_i(p_i) = \frac{\mu(1 - p_i)}{2(1 - \mu)p_i}.$$

To verify that this is an equilibrium, observe that the expected profit from quoting price p_i in the support of F_j is

$$\left[\frac{\mu}{2 - \mu} + \frac{2(1 - \mu)}{2 - \mu} F_j(p_i) \right] p_i = \frac{\mu}{2 - \mu}.$$

Prices outside the support of F_j yield a weakly lower payoff, and we conclude that these strategies are an equilibrium. The resulting ex ante sales price distribution is

$$\begin{aligned} S(x) &= \frac{\mu}{2} (F_1(x) + F_2(x)) + (1 - \mu) F_1(x) F_2(x) \\ &= \mu \left[\frac{\mu(1 - x)}{2(1 - \mu)x} \right] + (1 - \mu) \left[\frac{\mu(1 - x)}{2(1 - \mu)x} \right]^2 \end{aligned}$$

for $x \in [\mu/(2 - \mu), 1]$, and $S(x) = 1$ for $x < \mu/(2 - \mu)$. The no-information sales price distribution is the dotted curve in figure 3, again for $\mu = 1/2$. Note that the monopoly price is a best response for each firm, and by setting the monopoly price, a firm would sell only when they are the monopolist. Thus, each firm's ex ante payoff must be $\mu/2$, and producer surplus is μ . As a result, the no-information and full-information sales price distributions have the same expectation.

In section IV.D, we return to the comparison between no information and full information. It is a general result that when the price count distribution is held fixed, expected sales prices are the same under no information and full information. A key message of our paper is that these two extreme cases are not representative. Indeed, these two information structures minimize the expected sales price across all information structures and equilibria.

C. Maximum Prices with Public Information

Now suppose that firms observe a public signal about the price count (by which we mean that quoted firms observe the same signal with probability

one). In particular, firms observe the same signal $t_i = t_j \in \{i, j\}$. If firm i is a monopolist, the signal is i ; if the market is competitive, the signal is equally likely to be i or j . This information structure is depicted in figure 1, where $t_i = \emptyset$ represents the case where a firm is not quoted (and hence does not receive a signal).

We now describe firms' equilibrium strategies. If the public signal is i , both firms know that firm i is a monopolist with probability μ and that the market is competitive with probability $1 - \mu$. In this case, the firms mix over the interval $[\mu, 1]$. Firm i 's strategy is

$$F_i(p_i) = \frac{\mu}{p_i},$$

which has a mass point of size μ on the monopoly price $p = 1$, while firm j follows strategy without a mass point

$$F_j(p_j) = \frac{\mu(1 - p_j)}{(1 - \mu)p_j}.$$

As we did with our analysis of the no-information equilibrium, it is straightforward to verify that, given the pair of strategies (F_i, F_j) , each firm is indifferent across all prices in the common support $[\mu, 1]$ and strictly prefers them to any price outside the support, so that these strategies are an equilibrium.

The resulting ex ante sales price distribution is

$$S(x) = \mu(F_i(x)) + (1 - \mu)F_i(x)F_j(x) = \left(\frac{\mu}{x}\right)^2,$$

for $x \in [\mu, 1]$, and $S(x) = 1$ for $x < \mu$. It is depicted as the dashed curve in figure 3, again for $\mu = 1/2$. The expected sales price is $\mu(2 - \mu)$.

In the continuum interpretation, Armstrong and Vickers (2019) labeled the ex post markets (where only one firm has captive consumers) as "nested." They show that the expected price is higher, compared to the situation with no information. Indeed, even the "weak" firm that knows that it has no captive customers charges a higher average price than it does under no information. In a short note (Bergemann, Brooks, and Morris 2020), we show that the expected sales price under the above public information

$\pi(t K = \{i\})$				$\pi(t K = \{j\})$				$\pi(t K = \{i, j\})$			
t_i/t_j	$t_j = \emptyset$	$t_j = i$	$t_j = j$	t_i/t_j	$t_j = \emptyset$	$t_j = i$	$t_j = j$	t_i/t_j	$t_j = \emptyset$	$t_j = i$	$t_j = j$
$t_i = \emptyset$	0	0	0	$t_i = \emptyset$	0	0	1	$t_i = \emptyset$	0	0	0
$t_i = i$	1	0	0	$t_i = i$	0	0	0	$t_i = i$	0	1/2	0
$t_i = j$	0	0	0	$t_i = j$	0	0	0	$t_i = j$	0	0	1/2

FIG. 1.—Public signal distribution with two firms.

structure is indeed higher than that under any other public signal structure, using a standard concavification argument. The intuition is that we maximize the probability that one firm assigns to being a monopolist by considering nested markets. This firm will price less aggressively, which allows the rival to raise prices too.

D. Maximum Prices with Private Information

The expected price can be driven even higher with private signals. To see why, observe that under the nested public information structure, both the firm that thinks that it might be a monopolist and the firm that knows that it is not a monopolist follow mixed strategies with the same support. The common support is a necessary feature of models with public information. With private signals, however, we can similarly have a firm that thinks that it may be a monopolist always competing with a firm that knows that it is not, but with distinct supports. In particular, we describe in this paper a class of information structures and equilibria where a firm that thinks that it may be a monopolist always charges price 1 and a firm that knows that it not a monopolist always charges a price less than 1. It turns out that this class contains an information structure whose equilibrium sales price distribution first-order stochastically dominates not only the previous three examples but also any equilibrium sales price distribution under any information structure.

The information structure is as follows. Each quoted firm i receives a signal $t_i \in \{1, 2\}$. This signal can be interpreted as conveying to firm i a lower bound on the total number of quoted firms—including firm i . If only a single firm i is quoted, that is $K = \{i\}$, then only firm i receives a signal and $t_i = 1$ with probability one. If both firms are quoted, that is, $K = \{1, 2\}$, then with probability $1 - 2\alpha$ both firms receive the signal $t_i = 2$, with probability α the signals $t_1 = 1$ and $t_2 = 2$, and with probability α the signals $t_1 = 2$ and $t_2 = 1$.¹¹ These probabilities are summarized in figure 2, where again the null signal is used to represent an event where a firm is not quoted.

In effect, signal $t_i = 1$ means that there is at least one quoted firm, and signal $t_i = 2$ means that there are two quoted firms. Note that when both firms are quoted, there is a positive probability of 2α that one of the two firms receives signal $t_i = 1$ and believes that it is possible that only one firm was quoted. The parameter $\alpha \in [0, 1/2)$ therefore controls the dispersion in the beliefs of the market participants. If α is close to 0, then the information

¹¹ We note that the signals carry a different meaning in the public and private information environments. In the case of public information, the signal identifies the firm that is possibly a monopolist, thus i or j . In the case of private information, the signal identifies a lower bound on the number of competitors in the market.

$\pi(t K = \{1\})$				$\pi(t K = \{2\})$				$\pi(t K = \{1, 2\})$			
t_1/t_2	$t_2 = \emptyset$	$t_2 = 1$	$t_2 = 2$	t_1/t_2	$t_2 = \emptyset$	$t_2 = 1$	$t_2 = 2$	t_1/t_2	$t_2 = \emptyset$	$t_2 = 1$	$t_2 = 2$
$t_1 = \emptyset$	0	0	0	$t_1 = \emptyset$	0	1	0	$t_1 = \emptyset$	0	0	0
$t_1 = 1$	1	0	0	$t_1 = 1$	0	0	0	$t_1 = 1$	0	0	α
$t_1 = 2$	0	0	0	$t_1 = 2$	0	0	0	$t_1 = 2$	0	α	$1 - 2\alpha$

FIG. 2.—Private signal distribution with two firms.

structure is close to full information, and with high probability both of the firms learn that they are in a competitive environment. If α is close to $1/2$, then the information structure is close to that we constructed with public signals, and with high probability exactly one firm learns that the environment is competitive.

We now describe an equilibrium where the firm that has received signal $t_i = 1$ charges the monopoly price, $p_i = 1$, and the firm that receives the signal $t_i = 2$ mixes according to the upper cumulative distribution

$$F_i(p_i|t_i = 2) \triangleq F_i(p_i) = \frac{\alpha}{1 - 2\alpha} \frac{1 - p_i}{p_i}, \quad (2)$$

with support $p_i \in [\alpha/(1 - \alpha), 1]$. We refer to the firm that receives the signal $t_i = 2$ as “informed,” as such a firm knows the price count. Conversely, a firm that receives the signal $t_i = 1$ is “uninformed,” as the firm is uncertain whether it is in a monopoly or a competitive environment.

We claim that these strategies are an equilibrium if α is sufficiently small. To see this, observe that the informed firm’s profit from charging price p_i is generated by two events: with probability α the other firm observed signal $t_j = 1$, and with probability $1 - 2\alpha$ the other firm observed $t_j = 2$. Interim expected profit from setting a price p_i in the support of F_j is therefore

$$\left(\frac{\alpha}{1 - \alpha} + \frac{1 - 2\alpha}{1 - \alpha} F_j(p_i) \right) p_i = \frac{\alpha}{1 - \alpha},$$

so that firm i with signal $t_i = 2$ is indeed willing to randomize. Similarly, the uninformed firm either is a monopolist with probability $\mu/2$ or is in a competitive environment with probability $(1 - \mu)\alpha$. We need to ensure that the uninformed firm receives a higher revenue from posting the monopoly price 1 rather than choosing a price $p_i \in [\alpha/(1 - \alpha), 1]$, which reduces to the following inequality:

$$\frac{\mu/2}{\mu/2 + (1 - \mu)\alpha} \geq \frac{\mu/2 + (1 - \mu)\alpha F_j(p_i)}{\mu/2 + (1 - \mu)\alpha} p_i.$$

After inserting the mixed strategy $F_j(\cdot)$ given by equation (2), the above inequality reduces to

$$\frac{\mu/2}{\mu/2 + (1 - \mu)\alpha} (1 - p_i) \geq \frac{(1 - \mu)\alpha}{\mu/2 + (1 - \mu)\alpha} \frac{\alpha}{1 - 2\alpha} (1 - p_i).$$

We can cancel terms and rearrange to obtain

$$\alpha \leq \alpha^* \triangleq \frac{1}{2} \frac{\sqrt{\mu(2 - \mu)} - \mu}{1 - \mu}. \quad (3)$$

We note for future reference that since $[\mu(2 - \mu)]^{1/2} < 1$, it must be that $\alpha^* < 1/2$, so that equation (3) is not redundant with the requirement that $\alpha \leq 1/2$. Moreover,

$$\sqrt{\mu(2 - \mu)} = \sqrt{\mu + \mu(1 - \mu)} > \sqrt{\mu} > \mu,$$

so that $\alpha^* > 0$. Thus, the proposed strategies are an equilibrium as long as equation (3) holds, and this equation is satisfied for a nontrivial interval of α 's.

It is straightforward to calculate the expected sales price and sales price distribution in this equilibrium. When a single firm is quoted, that firm receives signal 1, and the resulting sales price is 1. Thus, there is an atom of size μ on a sales price of 1. If two firms are quoted, then either one or two firms receive the signal 2, and they randomize according to F_n given above. The sales price distribution is therefore

$$S(x) = \mu + (1 - \mu) \frac{\alpha^2}{1 - 2\alpha} \left[2 \frac{1 - x}{x} + \left(\frac{1 - x}{x} \right)^2 \right],$$

for $x \in [\alpha/(1 - \alpha), 1]$, and $S(x) = 1$ for $x < \alpha/(1 - \alpha)$. As for the expected price, that is even easier to calculate: firms are always indifferent to pricing arbitrarily close to 1. A firm setting such a price would always sell the good when a monopolist and would also sell when they are the only informed firm. Equilibrium producer surplus is therefore $\mu + (1 - \mu)2\alpha$.

Note that the sales price distribution is increasing in α for every x , that is, in the first-order stochastic dominance order. Thus, sales prices are increasing in the noise in firms' signals about whether the consumer is contested, and the sales price distribution is maximized at $\alpha = \alpha^*$. The highest expected sales price is $[\mu(2 - \mu)]^{1/2}$, which is the upper bound on revenue that we referenced in the introduction. We observe that this is the square root of the expression of the sales price under public information. Indeed, theorem 1 below shows that this information structure and equilibrium maximize both the expected sales price and the sales price distribution. When $\mu = 1/2$, then the corresponding sales price distribution is the thick solid curve in figure 3.

Theorem 1 below establishes that a generalization of this construction delivers the highest possible sales price distribution for any number of firms and for any price count distribution. The analyses of this section




FIG. 3.—Sales price distributions with two firms, $v = 1$, $\mu = 1/2$.

and the next take the price count distribution as exogenous, as in Varian (1980). We continue the two-firm example in section V, where we extend the analysis to endogenous price count distributions.

IV. Bounds on Equilibrium Sales Prices

We now present our general results. We begin by stating our main result, theorem 1, which describes a tight upper bound on the equilibrium sales price distribution. The proof immediately follows. We then show a simple comparative static: the maximal sales price distribution is increasing in the price count distribution, where both distributions are ordered by first-order stochastic dominance. We use this observation to give a tight bound on the expected price as a function of the probability that the price count is one. Finally, we show that the minimum expected sales price is attained under both full information and no information.

A. Maximal Sales Price Distribution

Let us define a decreasing sequence of cutoff prices:

$$v = x_0 = x_1 > x_2 > \cdots > x_n,$$

where for $k > 1$,

$$x_k \triangleq v \left(\prod_{m=1}^k \left(\frac{Q_{m-1}}{Q_m} \right)^{(m-1)/m} \right),$$

where

$$Q_m \triangleq \sum_{l=1}^m l \mu(l), \quad (4)$$

for $m > 0$ and $Q_0 = 1$. We also define, for each $k \geq 1$, an upper cumulative distribution $\bar{S}(\cdot|k)$ whose support is $[x_k, x_{k-1}]$. In particular, $\bar{S}(\cdot|1)$ puts probability one on v , and for $k > 1$,

$$\bar{S}(x|k) \triangleq \frac{(x_k/x)^{k/(k-1)} - (x_k/x_{k-1})^{k/(k-1)}}{1 - (x_k/x_{k-1})^{k/(k-1)}} \quad (5)$$

for $x \in [x_k, x_{k-1}]$, $\bar{S}(x|k) = 1$ for $x < x_k$, and $\bar{S}(x) = 0$ for $x > x_{k-1}$. We then define $\bar{S}(x)$ according to

$$\bar{S}(x) \triangleq \sum_{l=1}^n \mu(l) \bar{S}(x|l), \quad (6)$$

or equivalently,

$$\bar{S}(x) = \mu(k) \bar{S}(x|k) + \sum_{m=1}^{k-1} \mu(m),$$

when $x \in [x_k, x_{k-1}]$. Finally, given sales price distributions $S(\cdot)$ and $S'(\cdot)$, we say that S *first-order stochastically dominates* S' if $S(x) \geq S'(x)$ for all x .

THEOREM 1 (First-order stochastic dominance). Fix a price count distribution μ . In any information structure $\{T, \pi\}$ and equilibrium F consistent with μ , the distribution of sales prices must be first-order stochastically dominated by $\bar{S}$, given by equations (5) and (6). Moreover, there exists an information structure and equilibrium consistent with μ for which $\bar{S}$ is the equilibrium sales price distribution.

An immediate corollary of theorem 1 is the following characterization of welfare.

COROLLARY 1 (Maximum producer surplus and minimum consumer surplus). Maximum producer surplus across all information structures and equilibria consistent with the price count distribution μ is $\bar{R} = \int_{x=0}^v x \bar{S}(dx)$. Minimum consumer surplus across all information structures and equilibria consistent with the price count distribution μ is $v - \bar{R}$.

Proof of corollary 1.—Clearly, producer surplus is the expected sales price, and since $\bar{S}$ is an equilibrium sales price distribution and first-order stochastically dominates every equilibrium sales price distribution, maximum expected producer surplus is the expected sales price under $\bar{S}$. Since the good is always sold, total surplus is always v , and hence minimum consumer surplus is $v - \bar{R}$. QED

To visualize the maximal sales price distribution, consider the case where $v = 1$ and there is a uniform distribution on the price count, so that $\mu(k) = 1/n$ for all k , that is, where the price count is uniformly distributed. We studied this example in the case where $n = 2$ in the previous section. Figure 4 plots $\bar{S}(x|k)$ in the case where $n = 5$. In figure 5, we plot the ex ante sales price distribution for n between 2 and 10.

B. Proof of Theorem 1

The proof of theorem 1 is divided into propositions 1–3, the formal proofs of which are in the appendix. Proposition 1 presents an integral inequality that any equilibrium sales price distribution must satisfy. Proposition 2 shows that any equilibrium sales price distribution that satisfies this inequality must be first-order stochastically dominated by $\bar{S}$. Finally, proposition 3 constructs an information structure and equilibrium for which the equilibrium sales price distribution is precisely $\bar{S}$.

Let us establish notation for the sales price distribution induced by an information structure (T, π) and a strategy profile F . Let $S_i(x|k)$ denote the conditional probability that the good is sold by firm i at a price greater than or equal to x , conditional on there being k firms quoting prices. This is the expression

$$S_i(x|k) \triangleq \sum_{K \subseteq N} \nu(K|k) \int_{t \in T_k} \int_{p \in [x, v]^K} \frac{\mathbb{I}_{i \in K(p)}}{|K(p)|} F_K(dp|t) \pi(dt|K)$$

that appeared earlier in the expected revenue formula (eq. [1]). Also, let

$$S(x|k) \triangleq \sum_{i=1}^n S_i(x|k)$$

denote the conditional sales price distribution, given a price count of k . Finally, let

$$S(x) \triangleq \sum_{k=1}^n \mu(k) S(x|k)$$

denote the ex ante sales price distribution.

The sales price distributions $S_i(x|k)$ are sufficient to determine revenue. In particular, firm i 's revenue is

$$\sum_{k=1}^n \mu(k) \int_{x=0}^v x S_i(dx|k), \quad (7)$$

and total revenue is

$$\int_{x=0}^v x S(dx). \quad (8)$$




FIG. 4.—Conditional and unconditional sales price distributions for $n = 5$ and uniformly distributed price counts.

Our first result is an integral inequality that must be satisfied by any equilibrium conditional sales price distribution.

PROPOSITION 1 (Upper bound on sales price distribution). In any equilibrium, the sales price distributions must satisfy, for all $x \in [0, v]$,

$$x \sum_{k=1}^n \mu(k) k S(x|k) \leq \int_{y=x}^v y S(dy). \quad (9)$$

To develop some intuition for the above inequality, consider the case where there is zero ex ante probability that x is the sales price; that is, $S(\cdot)$ does not have an atom at x . Suppose that we first select a firm at random, and then the selected firm deviates by setting a price of x whenever they would have set a price greater than x in equilibrium. We refer to this as a uniform price cut to x . We claim that the expected profit resulting from such a deviation (where we also take expectation with respect to which firm is the deviator) is precisely

$$\sum_{k=1}^n \mu(k) \left(\frac{k}{n} x S(x|k) + \frac{1}{n} \int_{y=0}^x y S(dy|k) \right). \quad (10)$$

Expression (10) can be understood as follows: conditional on the price count being k , there is a k/n chance that the firm we picked to deviate is quoted. Conditional on being quoted, there is a probability $S(x|k)$ that the equilibrium sales price would have been above x (with zero mass on x itself), so that all firms set a price strictly greater than x . As a result, the deviating firm will set the lowest price, which is equal to x , and make a sale. But if the equilibrium sales price would have been less than x , either (1) the deviating firm would have set a price less than x , in which case they do not change their price, or (2) another firm has the lowest price, which is less than x . As a result, the deviation does not affect which firms have the lowest price, and the deviating firm's surplus is simply what they would have received in equilibrium. There is a $1/k$ likelihood of having the lowest price conditional on being quoted and hence a $1/n$ ex ante likelihood of having the lowest price, which is distributed according to $S(dy|k)$. If no firm wants to deviate in this manner, then it must be that the average surplus from this deviation across firms, given in expression (10), is less than the firms' average equilibrium revenue, which is $1/n$ of expression (8). Multiplying this inequality by n yields inequality (9).

Inequality (9) is central to our subsequent analysis. For future reference, we can integrate the right-hand side of (9) by parts and rearrange it into the following form:

$$x \sum_{k=1}^n \mu(k) (k-1) S(x|k) \leq \int_{y=x}^v S(y) dy. \quad (11)$$

Here we have used the facts that S is an upper cumulative distribution function, so $S(dx) = -dS(x)/dx$, and that $S(x) = 0$ for $x > v$.

To simplify terminology, we say that an ex ante sales price distribution $S(\cdot)$ *deters uniform price cuts* if there exist conditional sales price distributions $\{S(\cdot|k)\}$ whose expectation is $S(\cdot)$ and that satisfy inequality (9).

PROPOSITION 2 (First-order stochastic dominance). If the ex ante sales price distribution S deters uniform price cuts, then $\bar{S}$ first-order stochastically dominates S .

The three main steps in proving this proposition are as follows. First, we argue that when maximizing the ex ante sales price distribution, it is without loss to consider distributions that have *ordered supports*, meaning that the supports of $S(\cdot|k)$ are intervals of the form $[y_k, y_{k-1}]$, where $\{y_k\}_{k=1}^n$ is an increasing sequence. In other words, the sales price is perfectly “negatively correlated” with the price count. The reason is that holding fixed $S(\cdot)$, it is always possible to define new conditional distributions so that k and x are countermonotonic, which leaves the right-hand side of inequality (9) unchanged but decreases the left-hand side, thereby relaxing the constraint. Second, we argue that it is without loss to consider distributions for which (9) holds as an equality. If not, it is possible to push up the sales price distribution everywhere, while still satisfying (9). Third, we show that the ordered-supports property, together with (9) as an equality, reduce to a first-order differential equation whose unique solution is the distribution $\bar{S}$.

Finally, we construct an information structure and equilibrium that attain the upper-bound distribution of sales prices. Each firm receives signals in $\bar{T}_i \triangleq \{1, \dots, n\}$. For $k > 1$ and $k \geq l > 1$, we define

$$\alpha(l|k) \triangleq \frac{Q_{l-1} x_{l-1} [(x_{l-1}/x_l)^{1/(l-1)} - 1]}{Q_{k-1} x_{k-1} (x_{k-1}/x_k)^{1/(k-1)}}, \quad \text{and} \quad (12)$$

$$\alpha(1|k) \triangleq \frac{Q_1 v}{Q_{k-1} x_{k-1} (x_{k-1}/x_k)^{1/(k-1)}}, \quad \text{and}$$

$$\beta_k \triangleq 1 - (1 - \alpha(k|k))^k.$$

The signals are then generated according to the following distribution:

$$\bar{\pi}(t|K) \triangleq \begin{cases} \frac{1}{\beta_{|K|}} \prod_{i \in K} \alpha(t_i|K) & \text{if } |\{i \in K | t_i = |K|\}| > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Thus, when the realized price count is k , the signals are generated by taking independent draws from $\alpha(\cdot|k)$ and throwing out realizations where all firms draw numbers less than k . An interpretation is that each firm's signal t_i is a lower bound of the realized price count. Significantly, at least one firm observes the true price count, while the others observe numbers that are weakly lower. A key feature is that each firm receives a signal that is a hard lower bound on the price count. This is natural if we assume that the consumer collects prices sequentially and each firm sees a subset of the firms from whom the consumer has previously solicited prices.

Finally, firms use the pricing strategy in which conditional on receiving a signal k , the firm mixes on $[x_k, x_{k-1}]$ according to

$$\bar{F}_i(p_i|k) \triangleq G(p_i|k) \triangleq \frac{(x_k/p_i)^{1/(k-1)} - (x_k/x_{k-1})^{1/(k-1)}}{1 - (x_k/x_{k-1})^{1/(k-1)}}. \quad (14)$$

The pricing strategy of the individual firm mirrors the ex ante sales price distribution $\bar{S}(\cdot|k)$ given in equation (5) and simply lowers every exponent from $k/(k-1)$ to $1/(k-1)$. We are now ready to state the following result.

PROPOSITION 3 (Maximal information and equilibrium). The strategies $\bar{F}$ defined by equation (14) are an equilibrium for the information structure $(\bar{T}, \bar{\pi})$, and the resulting ex ante sales price distribution is $\bar{S}$.

To prove proposition 3, we first verify that $\bar{\pi}$ as defined in equation (13) is in fact a conditional probability distribution. We then show that this information structure and these strategies induce the upper-bound ex ante sales price distribution $\bar{S}$. This is essentially an application of the binomial theorem: when the price count is k , the number of firms that observe a signal of k is a truncated binomial, where at least one firm must observe k . We then compute the expectation of $(G(x|k))^l$ over the number of firms l that observe a signal k , which is exactly $\bar{S}(x|k)$.

Finally, we show that the strategies in equation (14) are an equilibrium. This is established by separately considering price increases and price cuts. For price increases, it is shown that a firm with signal $t_i = k$ strictly prefers any price in $[x_k, x_{k-1}]$ to any price greater than x_{k-1} . For price cuts, we show that firms are actually indifferent between all prices in $[x_n, x_{k-1}]$. In fact, this is necessary in order to attain the bound on the price distribution: constraint (11) says that for each x , firms on average do not benefit from a uniform price cut to x . The critical distribution $\bar{S}$ satisfies these constraints as equalities, meaning that firms are on average indifferent to uniform price cuts. Of course, in equilibrium, firms cannot want to deviate in any manner, so it must be that all firms are indifferent to uniform price cuts; otherwise, if some firm had a strict preference not to uniformly cut prices, some other firm would have a strict preference to uniformly cut. By a similar logic, if firms do not benefit on average by cutting prices to x from all equilibrium prices above x , they must in fact be indifferent to a price cut from any given price above x . Thus, firms must be indifferent to all price cuts. As we argue in the formal proof, this is precisely the case for the information structure and strategies we constructed.

We can now complete the proof of theorem 1.

Proof of theorem 1.—Suppose that S is the sales price distribution induced by an information structure $\{T, \pi\}$ and equilibrium F consistent with μ . By proposition 1, S must deter uniform price cuts. Proposition 2 then implies that S is first-order stochastically dominated by $\bar{S}$. This proves the first part of the theorem, while the second part follows immediately from proposition 3. QED

C. Comparative Static in Price Count Distribution μ

We now report a simple and intuitive comparative static. Given two price count distributions μ and μ' , we say that μ first-order stochastically dominates μ' if $\sum_{l=1}^k \mu(l) \geq \sum_{l=1}^k \mu'(l)$ for all $k = 1, \dots, n$.

PROPOSITION 4 (Price count and equilibrium sales price distribution). Let μ and μ' be price count distributions, with corresponding maximal sales price distributions $\bar{S}$ and $\bar{S}'$. If μ' first-order stochastically dominates μ , then $\bar{S}$ first-order stochastically dominates $\bar{S}'$.

The proof of proposition 4 in the appendix actually shows an even stronger result: any ex ante sales price distribution that deters uniform price cuts under μ' also deters uniform price cuts under μ . A fortiori, the maximal sales price distribution under μ' also deters uniform price cuts under μ and hence is dominated by the maximal distribution for μ .

We illustrate this with our uniform example in which $v = 1$ and $\mu(k) = 1/n$ for all k . In figure 5, we display the ex ante sales price distribution as we vary the expected number of price quotes (and the maximal number of price quotes). As we increase the number of firms, the maximum sales price distribution decreases in the sense of first-order stochastic dominance.

Proposition 4 implies that holding fixed the probability of a single price count, $\mu(1)$, the upper bound on equilibrium sales price distribution is maximized when $\mu(2) = 1 - \mu(1)$, that is, when the price count is either 1 or 2. In that case, the maximum expected sales price is $v[\mu(1)(2 - \mu(1))]^{1/2}$. In contrast, under either full information or no information (see sec. IV.D),




FIG. 5.—Sales price distributions for different n and uniform price counts. Kinks occur at the cutoffs x_n which are boundaries between the supports of conditional sales price distributions, as depicted in fig. 4.

the expected sales price is $\mu(1)$. Figure 6 contrasts the resulting expected revenue as we vary the probability $\mu(1)$. Thus, in the presence of incomplete information, maximum revenue grows very quickly in the probability $\mu(1)$ of there being a monopoly. In particular, the marginal growth of maximum revenue is unbounded when $\mu(1) \approx 0$. This analysis does, however, show that the expected sales price converges to zero as $\mu(1)$ goes to zero, and we recover the competitive outcome as beliefs converge to common knowledge that there are at least two firms (in the product topology on higher-order beliefs). We formalize this result as the following corollary.

COROLLARY 2 (Competitive limit). Among all price count distributions with probability $\mu(1)$ of a price count of 1, a tight upper bound on the expected sales price is $v[\mu(1)(2 - \mu(1))]^{1/2}$. Marginal maximal revenue with respect to the probability of being a monopolist is $(1 - v)/[\mu(1)(2 - \mu(1))]^{1/2}$, which is unbounded at $\mu(1) = 0$, where the market is fully competitive.

Conversely, holding fixed $\mu(1)$, proposition 4 implies that $\bar{S}$ is minimized when $\mu(n) = 1 - \mu(1)$. In this case, constraint (11) becomes




FIG. 6.—Maximal revenue versus full-/no-information revenue, consistent with a given probability of monopoly. As proposition 5 shows, the straight line is also a lower bound on revenue, given $\mu(1)$, so that all possible revenues are between the two curves.

$$x(1 - \mu(1))(n - 1)S(x|n) \leq \int_{y=x}^v S(y) dy.$$

Since the right-hand side is bounded above by $v - x$, this equation implies that $S(x|n)$ converges to zero pointwise as n goes to infinity. Thus, when the expected number of firms grows large, revenue converges to the full-information benchmark, and firms obtain positive revenue only when they are monopolists. The upshot is that in order to lift prices above the full-information level, it is insufficient for firms to have partial information about whether they are monopolists; it is also necessary for price counts to be bounded.

D. Minimum Expected Price and No Information

We have shown that there is an equilibrium sales price distribution that first-order stochastically dominates any equilibrium sales price distribution arising under any information structure. A fortiori, we have also characterized the highest expected sales price across all equilibria and information structures. The following proposition provides a corresponding characterization of the minimum equilibrium expected sales price.

PROPOSITION 5 (Minimum expected price). The minimum expected sales price across all information structures and equilibria is $v\mu(1)$. The maximum consumer surplus across all information structures and equilibria is $v - \mu(1)v$. Minimum revenue and maximum consumer surplus are attained under full information and no information.

Proof. In any information structure and equilibrium, each firm i can always set a price $p = v$. This strategy guarantees firm i revenue of v when the consumer receives only one price quote, so that each firm's ex ante expected equilibrium revenue is bounded below by $v\mu(1)/n$, so that producer surplus is at least $v\mu(1)$. Clearly, this is producer surplus under full information, in which case firms price at v when the price count is one, and otherwise, Bertrand competition forces the price down to cost, which is zero. For the no-information environment, lemma 2 of Burdett and Judd (1983) established that there is a unique symmetric equilibrium. In this equilibrium, firms use nonatomic mixed strategies with support of the form $[p, v]$. As a result, each firm is indifferent to setting a price of $p = v$, in which case their ex ante profit is $v\mu(1)/n$. Finally, as we observed in the proof of corollary 1, total surplus is always v , so that consumer surplus is maximized when producer surplus is minimized. QED

Note that in many cases, the no-information setting has other asymmetric equilibria, although they must have weakly higher revenue than $v\mu(1)$.¹²

¹² Narasimhan (1988) identifies the unique equilibrium when there are only two firms. Baye, Kovenock, and de Vries (1992) construct all equilibria in the model of Varian (1980), where the price count is either $k = 1$ or $k = N$.

For example, when $n = 3$ and the price count is either $k = 1$ or $k = 2$, there are asymmetric equilibria in which two of the three firms essentially play the $n = 2$ equilibrium and the third firm prices at the monopoly level. Note that in this equilibrium, all firms are indifferent to setting the monopoly price and selling only to captive consumers, so that the expected sales price is the same. We do not know of any asymmetric equilibria with different expected price.

We close this section with a few remarks on the implications of proposition 5. First, a failure of common knowledge of the price count alone cannot explain an increased expected price, unless firms have partial and private information about the price count. In addition, price discrimination weakly increases the expected price relative to uniform pricing (which corresponds to the no-information case). Proposition 4 of Armstrong and Vickers (2019) showed this to be true for any public signal structure. We show that price discrimination weakly increases the expected price relative to the no-information case under any information structure, including private information structures.¹³ Finally, since the equilibrium sales price distributions under full information and no information are distinct, we can see that there is no lower bound on equilibrium sales price distribution analogous to the upper bound in theorem 1.

V. Endogenizing the Price Count

A. *Feedback versus No Feedback*

We have thus far studied equilibrium sales price distributions holding the price count distribution fixed. As discussed in the introduction, there is a plethora of models of how the price count distribution is determined. An important question is whether our bounds still apply when we endogenize price counts, for example, with a dynamic model of consumer search. In this section, we explore this issue in detail by considering various ways of endogenizing price counts that have been proposed in the literature. In all of these models, the firms' prices and the price count are jointly determined in equilibrium. There is an important distinction, however, as to how a firm's price affects the price count and other firms' behavior. In many of these models, price counts and prices depend only on beliefs about how firms will price in equilibrium. In particular, there is no feedback from realized prices to price counts and to other firms' prices.¹⁴ When there is no feedback, our analysis in the previous section immediately applies to

¹³ Proposition 4 of Armstrong and Vickers (2019) assumes that there are two firms. They allow a general demand function, while we assume single-unit demand, but our lower bound on the expected price extends easily, as discussed below.

¹⁴ These models, including the seminal contributions of Butters (1977), Varian (1980), Burdett and Judd (1983), and Baye and Morgan (2001), share two significant features: (1) the price counts are determined before or simultaneous with the prices, and (2) the prices are determined simultaneously by the firms. As a result, there is no feedback from

whatever price count distribution is realized in equilibrium. The reason is that the critical uniform price cut that drives our bounds is still available to firms. From the perspective of a deviating firm, changing their price has no effect on the price count or the prices set by other firms. As a result, the equilibrium sales price distribution must still satisfy the critical inequality (9), and proposition 2 shows that the equilibrium distribution is bounded by $\bar{S}$.

There are other models, however, that exhibit feedback, meaning that a firm's realized price can directly affect price counts and/or other firms' prices. Whether our bounds hold in such models depends on the particular form of feedback. To illustrate the possibilities, we give one simple example in which our bounds are violated and another rich example for which our bounds apply. For the former, consider a two-firm, full-information Stackelberg game. This game has an equilibrium in which, on the equilibrium path, both firms price at or above the monopoly level and the second firm makes the sale. This is supported by off-path play in which the following firm undercuts any price set by the leader. This equilibrium obviously violates our bounds. The reason is that when firms move sequentially and observe one another's prices, the strategic response can lower the payoff from price cuts, so that the equilibrium sales price distribution need not satisfy inequality (9).

Another prominent example of feedback is sequential consumer search, in which a consumer iteratively solicits price quotes and, after observing the quoted price, decides whether to purchase or continue searching. Firms' realized prices feed back directly into the price count through the consumer's decision of when to stop. While observable prices can make price cuts less attractive, sequential search generally has the opposite effect and makes price cuts more attractive. The reason is that price cuts tend to make consumers stop searching sooner, so that a deviating firm faces less competition relative to the benchmark with fixed price counts. Hence, constraint (9) will still hold in equilibrium, and the rest of our bounding argument goes through.

The remainder of this section formally develops the analogue of theorem 1 for sequential search. In particular, theorem 2 shows that for a fairly large class of sequential search models, where we vary both firms' information and the consumers' search costs, the distribution $\bar{S}$ is a tight upper bound on the equilibrium ex ante sales price distribution.

B. A Model of Sequential Search

Time is discrete. At each period, a consumer decides whether to purchase at the lowest price found thus far or continue searching. If they

firms' pricing decisions to the price count or to other firms' prices. Numerous other examples in this class are described in a survey by Baye, Morgan, and Scholten (2006).

choose to search, a new firm is drawn without replacement and quotes a price. The latent order in which firms will be searched is denoted by a permutation $\xi: \mathcal{N} \rightarrow \mathcal{N}$, where $\mathcal{Z}$ denotes the set of permutations. The interpretation is that if the consumer searches at least k firms, then firm $\xi(k)$ provides the k th quote. All orders are equally likely.

As in our baseline model, the consumer has value v for a single unit, which can be produced at zero cost by each of the firms. In addition, the consumer has a type θ in a measurable set Θ , which is distributed according to $\eta \in \Delta(\Theta)$. The consumer chooses a number $k \geq 1$ of firms to search. If a consumer searches k firms, then they pay a cost $c(k, \theta)$. The parameter θ allows for heterogeneity in search costs among consumers. We make the simplifying assumption that for all θ , k , and k' , $c(k, \theta) \neq c(k', \theta)$.¹⁵ If the consumer purchases at price p after visiting k firms, their payoff is $v - p - c(k, \theta)$. The payoff to the firm that makes the sale is p , and other firms' payoffs are zero. As before, ties for lowest price at the time the consumer stops searching are broken uniformly.

We continue to model the firms' beliefs using an information structure. Each firm has a set of signals T_i . For this section, we assume for simplicity that the signal sets are finite. Conditional on θ and ξ , there is a joint distribution over signals denoted by $\pi(t|\theta, \xi)$. We further assume that after searching k firms, the consumer sees the history $(\theta, \{(\xi(l), t_{\xi(l)}, p_{\xi(l)})\}_{l=1}^k)$. The set of such histories of length k is denoted H_k , and the set of all histories is H . These sets are endowed with their natural product measurable structure. Thus, firms have partial information about both the consumer's type and the order in which firms are searched. The consumer, on the other hand, knows their own type and the identities, signals, and quoted prices of the firms that were searched.

The strategy of firm i is a pricing kernel

$$F_i: T_i \rightarrow \Delta([0, v]).$$

As before, $F_i(\cdot|t_i)$ is an upper cumulative function. The strategy of the consumer is a measurable function $\sigma: H \rightarrow [0, 1]$, where $\sigma(h)$ is the probability that the consumer continues searching at history h . With the complementary probability, the consumer buys from one of the firms with the lowest price quoted thus far, breaking ties uniformly. We further impose that for $h \in H_n$, $\sigma(h) = 0$; that is, the consumer must buy after searching all of the firms. Note that a strategy profile for the firms and the consumer induces a distribution over the number of firms the consumer searches, that is, the price count, as well as a distribution over the sales price.

¹⁵ Our results can be readily generalized to allow for any cost function, if we assume that the consumer breaks ties in favor of searching more firms if doing so does not increase the search cost.

We analyze perfect Bayesian equilibria (Fudenberg and Tirole 1991). That is, we analyze Nash equilibria where players are also sequentially rational off the equilibrium path, relative to beliefs off the equilibrium path that are consistent with Bayes rule where possible. In particular, we require that the consumer's strategy continues to be optimal even if firms deviate in their prices. The remaining and complete description of the extensive-form game is in the appendix.

To summarize, the parameters of the sequential search model are $\{\Theta, \eta, c, T, \pi\}$. This class of sequential search models generalize Stahl (1989, 1996), where the consumer simply has a constant cost per search.¹⁶

C. Extending the Upper Bound to Sequential Search

The following result shows that $\bar{S}$ is a tight upper bound on the equilibrium sales price distribution for the sequential search model just described.

THEOREM 2 (Sequential search and upper bound). Fix a price count distribution $\mu \in \Delta(\{1, \dots, n\})$. For any sequential search model $\{\Theta, \eta, c, T, \pi\}$ and equilibrium (F, σ) such that the equilibrium price count distribution is μ , the induced sales price distribution is first-order stochastically dominated by $\bar{S}$, given by equations (5) and (6). Moreover, there exists a sequential search model and equilibrium that induce μ , and the equilibrium sales price distribution is $\bar{S}$.

The full proof is in the appendix. We here provide a sketch. First, it is straightforward to adapt the construction preceding proposition 3 to sequential search, so that the strategies $\bar{F}$ are an equilibrium and induce the ex ante sales price distribution $\bar{S}$. We simply set $\Theta = \{1, \dots, N\}$, $\eta(\theta) = \mu(k)$, and¹⁷

$$c(k, \theta) = \begin{cases} 0 & \text{if } k = \theta, \\ v + k & \text{otherwise.} \end{cases}$$

With this model, it is a strictly dominant strategy for the consumer of type θ to search θ firms. As a result, the equilibrium price count distribution must be μ , regardless of firms' strategies. In addition, the information is again given by $\bar{T}_i = \{1, \dots, n\}$ and

¹⁶ While our formal model is one of sequential search, our bounds will also apply for any pricing game that ends in sequential search, but where firms and the consumer can take earlier actions that influence the endogenous determination of the price count distribution without adding feedback beyond that in the sequential search model. For example, Ellison and Ellison (2009) and Ellison and Wolitzky (2012) have shown empirically and theoretically that firms have an incentive to increase search costs strategically to raise prices, and our bounds will apply in the case of the latter theoretical model as well.

¹⁷ The purpose of the $+k$ term is merely to satisfy our genericity assumption that $c(k, \theta) \neq c(k', \theta)$ for all k, k' , and θ .

$$\pi(t|\theta, \xi) = \bar{\pi}(t|\{i|\xi(i) \leq \theta\}).$$

By the same argument as for proposition 3, we conclude that firms' strategies are an equilibrium that induce the sales price distribution $\bar{S}$.

The rest of the proof generalizes proposition 1 by showing that the equilibrium sales price distribution in any sequential search model must deter uniform price cuts. Recall that in the proof of proposition 1, we showed that the expected payoff from a uniform price cut (averaged across firms) is given by expression (10). This expression presumes that price counts are unaffected by the deviation. With sequential search, it turns out that this is a lower bound on the average payoff from the uniform price cut. The reason is that the price cut may cause the consumer to search less and hence may result in a higher probability of the deviator making a sale. This step requires careful argument. Because the consumer sees all the signals of previously searched firms, the consumer's beliefs about future prices do not depend on past prices. As a result, consumers will adopt simple cutoff strategies in equilibrium, whereby they purchase the good as soon as the lowest price encountered is below a cutoff that depends on previously searched firms' signals and the consumer's own type. Once we know that consumers use cutoff strategies, it is easy to see that price cuts will lead to lower price counts and a weakly higher payoff for the deviating firm, relative to a benchmark in which price counts are held fixed.

As an illustration, consider a version of Stahl's model in which there are two firms. With probability $1/2$, the consumer is a "shopper" who observes both firms' prices for free, but with probability $1/2$, the consumer is a nonshopper who observes one quote and can then choose (after observing the price) to pay $c > 0$ to observe the second price. In equilibrium, the nonshopper gets only one quote. Firms follow a mixed strategy

$$F(p) = \frac{r - p}{2p},$$

with support $[r/3, r]$, where

$$r \triangleq \frac{c}{1 - (1/2) \ln(3)/2}.$$

In figure 7, we add the equilibrium sales price distribution for the sequential search model for $c = 1/2$ (for which $r \approx 0.74$) to the uniform price count distribution in section III. Sales price distributions for this example under various informational assumptions are depicted in figure 3. We can see that sequential search lowers the sales price distribution relative to simultaneous search with the same equilibrium price count distribution.

Thus, we conclude that our bounds will extend to a nontrivial class of models with sequential search. Importantly, we have assumed that each




FIG. 7.—Equilibrium sales price distribution under sequential search.

firm does not observe other firms' prices before setting their own price. We have also assumed that firms have no information that is unobservable to the consumer, so that consumers do not learn from realized prices about firms' beliefs. These assumptions cannot be easily dispensed with. Generalizing our bounds to allow for richer learning by firms and consumers is an important direction for future work.

VI. Further Topics

We now consider three additional topics. We first discuss how our results can be generalized to allow for downward-sloping demand. We then discuss how our results can be extended when firms have different probabilities of being quoted. Finally, we discuss the relationship between the pricing game analyzed here and the first-price auction.

A. Beyond Single-Unit Demand

We derived our results on the maximal sales price distribution under the assumption of single-unit demand. This assumption is easily relaxed at the cost of extra notation, with our tight upper bound on the sales price distribution continuing to hold.

Suppose that the consumer has multiunit demand. We maintain the assumption that firms produce homogeneous goods and quote a single price for all units, so that the consumer purchases only from a low-price firm. If the lowest price is p , the consumer purchases $D(p)$ units. With this additional structure the uniform downward incentive constraint (9) becomes

$$xD(x) \sum_{k=1}^n \mu(k) kS(x|k) \leq \int_{y=x}^v yD(y)S(dy).$$

We assume without loss that firms use only prices in the set

$$P^* = \{p \mid \nexists p' < p \text{ such that } p'D(p') \geq pD(p)\}.$$

Using a price p outside of P^* is weakly dominated, in that there is another price that induces weakly more revenue and is lower than p , so that it is more likely to be the lowest price and attract consumers. Firms never price above the monopoly price

$$p^M = \arg \max_{p \in P^*} pD(p).$$

Under the further assumption that $D(p)$ is continuous, the set of possible revenue levels $\{pD(p) \mid p \in P^*\}$ is convex and in fact is the interval $[0, p^M D(p^M)]$. We can then treat the associated revenue levels as the prices in the baseline model, with $p^M D(p^M)$ being the analogue of the consumer's value v . All the derivations from section IV go through as before to obtain an equivalent result for theorem 1, which would state that the equilibrium distribution of $\min pD(\min p)$ is bounded above by $\tilde{S}$, where $\tilde{S}(x) = \bar{S}(y)$, where $y \in P^*$ is such that $yD(y) = x$.

The single-unit-demand assumption does deliver the result that the allocation is always efficient, and therefore the sum of producer surplus (the price) and consumer surplus is always v . This allows us to report straightforward implications of our pricing results for consumer surplus. It might be possible to derive implications for consumer surplus under downward-sloping demand using the additional structure assumed in Armstrong and Vickers (2019) in their analysis of the two-firm case.

More generally, theorem 1 can be adapted to other settings in which firms sell an abstract set of goods and offer menus of good/price bundles and consumers purchase the bundle offering them the most surplus. Under fairly general conditions, such a model is strategically equivalent to the Bertrand pricing game with unit demand, as has been noted by Armstrong and Vickers (2001).

B. Beyond a Symmetric Quote Distribution

We derived our results on the maximal sales price distribution under the assumption that all firms are equally likely to be quoted, as in classic search

models, such as Varian (1980) and Burdett and Judd (1983). A number of asymmetric generalizations of the incomplete-information Bertrand game have since been analyzed: see Armstrong and Vickers (2020) for a recent analysis of this problem (in the no-information case) as well as a review of this literature. Their findings suggest that even under no information, the equilibrium analysis with asymmetric distributions can be quite complicated. We argue here that our construction continues to yield an upper bound on the sales price distribution even if quotation probabilities differ across firms, although the bound may no longer be tight.

Recall that μ denotes the ex ante distribution of the price count and that $\nu(K|k)$ is the distribution over the set of quoted firms conditional on the price count. We assumed that, conditional on the price count k , all firms are equally likely to be quoted. Suppose that we instead allowed an arbitrary conditional distribution $\nu(K|k)$, so that firms need not be equally likely to be quoted. What would happen to our results? Suppose that the sales price distribution S can be attained in an equilibrium under some information structure, given ν . Now, imagine that we generated a new distribution ν' by permuting the identities of the bidders. Then clearly, there is an information structure and equilibrium under which S is the sales price distribution when the distribution of the set of quoted firms is ν' , where we just push the permutation of firms' identities through the information and strategies. Finally, consider the distribution $\tilde{\nu}$, which is obtained from ν by taking an average over all permutations of bidders' identities. Then again, S is an equilibrium sales price distribution of this model. But $\tilde{\nu}$ has the property that firms are equally likely to be quoted, so that S is less than the maximal sales price distribution $\bar{S}$ corresponding to μ .

Thus, we conclude that $\bar{S}$ is still an upper bound on the equilibrium sales price distribution, even with asymmetric quotation probabilities. Importantly, the bound may no longer be tight, and the construction in proposition 3 uses the fact that firms are equally likely to be quoted. Note that the bound is tight if we actually started with the "symmetrized" distribution $\tilde{\nu}$. But while the previous paragraph shows that if S is attainable under ν then it is also attainable under $\tilde{\nu}$, the converse does not hold. The reason is that if bidders learned their permuted identities, then this constitutes a lower bound on their information that the construction of proposition 3 does not satisfy.¹⁸

C. Connection to First-Price Auctions

The pricing game that we analyze here is strategically related to a first-price auction where each bidder has either a low or a high value for a good and each bidder knows their private value but is uncertain about

¹⁸ Bergemann, Brooks, and Morris (2013) report some preliminary results on tight bounds for the asymmetric quote distribution case when there are two firms.

the values of the other bidders. In the equilibrium of the first-price auction, low-value bidders will always bid the low value, and high-value bidders follow mixed strategies that depend on their beliefs about the number of other bidders with high values. The pricing game studied in this paper can be viewed as a procurement auction, where bidders quote prices at which they are willing to sell and the auctioneer buys at the lowest price. Quoted firms are analogous to low-cost bidders, while nonquoted firms are like high-cost bidders.

Fang and Morris (2006) analyzed the two-bidder first-price auction in which bidders have known binary private values and also observe additional information about the other bidder's value. Fang and Morris (2006) restricted attention to conditionally independent binary noisy signals about the opponent's value and noted that the expected price is necessarily higher with partial information than with either no information or full information, a result that this paper generalizes. Åzakis and Vida (2015) allowed many conditionally independent signals and also considered the possibility of correlated signals, and they noted that the critical information structure we identify in section III gives rise to a higher expected price than any conditionally independent signal structure. Our unpublished working paper, Bergemann, Brooks, and Morris (2013), initiated the study of what can be said in first-price auctions under all information structures and showed that the highest expected price in the two-bidder binary-value case is attained in the critical information structure identified in section III. Bergemann, Brooks, and Morris (2013) also provided results on two-bidder auctions with binary private values in asymmetric environments that are not reported in this paper.

In subsequent work, Bergemann, Brooks, and Morris (2017) characterized what can happen in first-price auctions under general information structures. Lemma 1 of Bergemann, Brooks, and Morris (2017) established bounds on the equilibrium bid distribution in the first-price auction, when bidders are allowed to have arbitrary common prior information about all bidders' values. Proposition 1 in this paper establishes similar bounds when bidders know their own values. This is a different, and in general much harder, problem to solve, as there are more constraints on bidders' higher-order beliefs. We are able to completely solve this case only because there are only two values (in the standard auction interpretation of the problem). The proof of theorem 1—establishing that the bound in proposition 1 is tight and showing that a single sales distribution bounds all possible equilibrium sales distributions—has no analogue in our earlier published work.¹⁹

¹⁹ In Bergemann, Brooks, and Morris (2015), an early version of Bergemann, Brooks, and Morris (2017), we reported further initial steps for the known private-value environment with binary values and many players. In particular, theorem 11 there gives an implicit and incomplete characterization of maximal bid distributions.

VII. Conclusion

We have revisited the standard model of price dispersion in homogenous-goods markets, in which firms randomize over prices because of a failure of common knowledge of whether the consumer has quoted at least two prices. The novelty of our analysis is that rather than trying to fully explain the origins of the price count, we simply take it as a primitive, and from it we derive a tight upper bound on the equilibrium distribution of sales prices. The bound holds across a rich family of models that endogenize the price count and for all common-prior beliefs that firms might have about the price count.

A primary application of the bound is to test whether prices in a given market can be rationalized by competitive behavior, given the distribution of the number of prices quoted by consumers. This test does not require the analyst to know what motivated the observed price count, such as consumers' perceived costs of searching for price quotes.

An important direction for future research is to further relax our modeling assumptions by allowing for more complicated forms of feedback from prices to price counts and partial observability of prices by other firms. Moreover, we have assumed that firms produce homogenous goods and have symmetric and publicly known costs of production. These assumptions played an important role in our analysis. Generalizing the theory to allow for heterogeneous goods and private information about costs are important goals for making this theory more empirically relevant.

Appendix

Proofs

A1. Proofs for Section IV

A1.1. Proof of Proposition 1

Fix x . Let $\{\epsilon_l\}_{l=0}^\infty$ be a sequence of positive numbers, converging to zero, such that S does not have an atom at $x - \epsilon_l$ for all l . Such a sequence exists because S has at most countably many atoms. Suppose that firm i deviates in the following manner: whenever firm i would have set a price $p_i > x - \epsilon_l$, it sets a price of $x - \epsilon_l$ instead. This deviation affects the outcome only when the lowest price would have been greater than $x - \epsilon_l$, and in particular, the deviator's surplus is

$$\sum_{k=1}^N \mu(k) \sum_{K \subseteq N} \nu(K|k) \int_{t \in T_k} \int_{p \in [0, v]^K} \left((x - \epsilon_l) \mathbb{I}_{\min p > x - \epsilon_l} + \min p \mathbb{I}_{\min p \leq x - \epsilon_l} \frac{\mathbb{I}_{j \in K(p)}}{|K(p)|} \right) F_k(dp|t) \pi(dt|K).$$

Note that this expression must be less than firm i 's equilibrium surplus, given in expression (7). As l goes to infinity, the dominated-convergence theorem implies that the deviator's surplus converges to

$$\sum_{k=1}^n \mu(k) \sum_{K \subseteq N} \nu(K|k) \int_{t \in T_k} \int_{p \in [0, v]^K} \left(x \mathbb{I}_{\min p \geq x} \mathbb{I}_{i \in K} + \min p \mathbb{I}_{\min p < x} \frac{\mathbb{I}_{i \in K(p)}}{|K(p)|} \right) F_K(dp|t) \pi(dt|K),$$

which is necessarily also less than (7). Summing the deviation surplus across i , we obtain

$$\begin{aligned} & \sum_{k=1}^n \mu(k) \sum_{K \subseteq N} \nu(K|k) \int_{t \in T_k} \int_{p \in [0, v]^K} (kx \mathbb{I}_{\min p \geq x} + \mathbb{I}_{\min p < x}) F_K(dp|t) \pi(dt|K) \\ &= \sum_{k=1}^n \mu(k) \left(kx \int_{t \in T_k} \int_{p \in [x, v]^K} \sum_{K \subseteq N} \nu(K|k) F_K(dp|t) \pi(dt|K) \right. \\ & \quad \left. + \int_{t \in T_k} \int_{p \in [0, x]^K} \min p \mathbb{I}_{\min p < x} \sum_{K \subseteq N} \nu(K|k) F_K(dp|t) \pi(dt|K) \right) \\ &= \sum_{k=1}^n \mu(k) \left(kx S(x|k) + \int_{y=0}^x y \int_{t \in T_k} \int_{\{p | \min p = y\}} \sum_{K \subseteq N} \nu(K|k) F_K(dp|t) \pi(dt|K) \right) \\ &= \sum_{k=1}^n \mu(k) \left(kx S(x|k) + \int_{y=0}^x y S(dy|k) \right). \end{aligned}$$

This must be less than the sum of the firm's equilibrium revenues, which is exactly inequality (9). QED

A1.2. Proof of Proposition 2

Fix $x \in [0, v]$. Consider the problem of maximizing $S(x)$ over all $\{S(\cdot|k)\}_{k=1}^n$ that satisfy inequality (9) and where the functions $S(\cdot|k)$ are measurable functions that map $[0, v]$ into $[0, 1]$. Note that the set of conditional distributions is compact in the weak* topology (which is the topology of pointwise convergence on $\{S(\cdot|k)\}_{k=1}^n$), constraint (11) is closed, and the objective $S(x)$ is continuous, so that an optimal conditional distribution exist. We show here that $\bar{S}(x)$ is the optimal value. This is established in three steps.

Step 1.—When maximizing $S(x)$, it is without loss to restrict attention to $\{S(\cdot|k)\}_{k=1}^n$ that satisfy the following ordered-supports property:

$$S(y|k) < 1 \Rightarrow S(y|k') = 0 \quad \forall k' > k.$$

Indeed, given any $\{S(\cdot|k)\}$ and associated ex ante distribution $S(\cdot)$, we can define a new $\{\tilde{S}(\cdot|k)\}$ with the same ex ante distribution, but where there is negative assortative matching between k and x . In particular, noting that $S(1) \leq \mu(1)$ from constraint (11), for each k , we define $\tilde{x}_k$ as the infimum y such that $S(y) \geq \sum_{m=1}^k \mu(m)$. We then set $\tilde{S}(y|k) = (S(y) - \sum_{m=1}^{k-1} \mu(m)) / \mu(k)$ on $[\tilde{x}_{k-1}, \tilde{x}_k]$. For each y , this correlation structure minimizes

$$\sum_{k=1}^n \mu(k) (k-1) S(y|k),$$

and hence the left-hand side of (11). For future reference, note that for conditional distributions satisfying ordered supports, (11) is, for all $y < v$ and $k > 1$, equivalent to

$$S(y|k) \leq \max \left\{ 0, \frac{1}{\mu(k)(k-1)} \left(\frac{1}{y} \int_{z=y}^v S(z) dz - \sum_{m=1}^{k-1} \mu(m)(m-1) \right) \right\}. \quad (A1)$$

Step 2.—Among distributions with the ordered-supports property, it is obviously without loss to set $S(v|1) = 1$ (since $S(v|1)$ is unconstrained except for $S(v|1) \leq 1$). Now, if a solution does not satisfy condition (A1) as an equality when $S(y|k) < 1$, we can define a new solution, which is

$$\tilde{S}(y|k) = \max \left\{ 0, \min \left\{ 1, \frac{1}{\mu(k)(k-1)} \left(\frac{1}{y} \int_{z=y}^v S(z) dz + \sum_{m=1}^{k-1} \mu(m)(m-1) \right) \right\} \right\},$$

which satisfies ordered supports, necessarily satisfies $\tilde{S}(y|k) \geq S(y|k)$ (strictly whenever [11] is strict), and therefore induces a higher ex ante distribution $\tilde{S}$. Thus, it is without loss to restrict attention to solutions for which (11) holds as an equality whenever $S(y|k) < 1$.

Step 3.—We now show that the ordered-supports property and condition (A1) holding as an equality uniquely define the distributions $\{\bar{S}(\cdot|k)\}$. It is immediate that $S(y|k)$ will have a support that is an interval $[y_k, y_{k-1}]$, with $y_0 = y_1 = v$, and it is strictly increasing on its support. In addition, since the right-hand side of (A1) is continuous, we conclude that the only mass point of S is at v . Now, suppose inductively that we have defined $S(y|m)$ and y_m for $m < k$. Then $S(y)$ must satisfy the boundary conditions $S(y_m) = \sum_{l=1}^m \mu(l)$ for all $m < k$. On $[y_k, y_{k-1}]$, (A1) holds as an equality, and moreover

$$S(y) = \mu(k)S(y|k) + \sum_{m=1}^{k-1} \mu(m).$$

As a result, (A1) with equality rearranges to

$$y(k-1)S(y) - \int_{z=y}^v S(z) dz = y \sum_{m=1}^{k-1} \mu(m)(k-m).$$

Multiplying both sides by $y^{-(k-2)/(k-1)}/(k-1)$, we obtain

$$y^{1/(k-1)}S(y) - \frac{y^{-(k-2)/(k-1)}}{k-1} \int_{z=y}^v S(z) dz = \frac{y^{1/(k-1)}}{k-1} \sum_{m=1}^{k-1} \mu(m)(k-m).$$

Integrating both sides, we obtain

$$-y^{1/(k-1)} \int_{z=y}^v S(z) dz = C_k + \frac{y^{k/(k-1)}}{k} \sum_{m=1}^{k-1} \mu(m)(k-m),$$

where C_k is a constant of integration. Thus,

$$-\int_{z=y}^v S(z) dz = y^{-1/(k-1)} C_k + y \left(\sum_{m=1}^{k-1} \mu(m) - \frac{1}{k} Q_{k-1} \right),$$

where Q_{k-1} is defined above in equation (4). Differentiating both sides again, we obtain

$$S(y) = -\frac{C_k}{k-1} y^{-k/(k-1)} + \sum_{m=1}^{k-1} \mu(m) - \frac{1}{k} Q_{k-1}.$$

The boundary condition $S(y_{k-1}) = \sum_{m=1}^{k-1} \mu(m)$ then implies that

$$\begin{aligned} \sum_{m=1}^{k-1} \mu(m) &= -\frac{C_k}{k-1} (y_{k-1})^{-k/(k-1)} + \sum_{m=1}^{k-1} \mu(m) - \frac{1}{k} Q_{k-1} \\ \Leftrightarrow C_k &= -\frac{k-1}{k} y^{k/(k-1)} Q_{k-1}. \end{aligned}$$

As a result,

$$S(y) = \frac{1}{k} \left[\left(\frac{y_{k-1}}{y} \right)^{k/(k-1)} - 1 \right] Q_{k-1} + \sum_{m=1}^{k-1} \mu(m).$$

The next boundary condition, $S(y_k) = \sum_{m=1}^k \mu(m)$, is equivalent to

$$\begin{aligned} k\mu(k) &= \left(\frac{y_{k-1}}{y_k} \right)^{k/(k-1)} Q_{k-1} - Q_{k-1} \\ \Leftrightarrow Q_k &= \left(\frac{y_{k-1}}{y_k} \right)^{k/(k-1)} Q_{k-1} \\ \Leftrightarrow y_k &= y_{k-1} \left(\frac{Q_{k-1}}{Q_k} \right)^{(k-1)/k}. \end{aligned}$$

Together with the initial condition $y_1 = v$, this implies that $y_k = x_k$, the boundaries that define $\bar{S}$. Finally, it must be that for $y \in [y_k, y_{k-1}] = [x_k, x_{k-1}]$,

$$\begin{aligned} S(y|k) &= \frac{S(y) - S(y_{k-1})}{\mu(k)} \\ &= \frac{\frac{Q_{k-1}}{k} (y_{k-1}/y)^{k/(k-1)} - (y_{k-1}/y_{k-1})^{k/(k-1)}}{\mu(k)} \\ &= Q_k \frac{(y_k/y)^{k/(k-1)} - (y_k/y_{k-1})^{k/(k-1)}}{k\mu(k)} \\ &= \frac{(y_k/y)^{k/(k-1)} - (y_k/y_{k-1})^{k/(k-1)}}{1 - (Q_{k-1}/Q_k)}, \end{aligned}$$

which is precisely $\bar{S}(y|k)$. QED

A1.3. Proof of Proposition 3

We first verify that the information structure is well defined, that is, that $\bar{\pi}$ is a conditional probability distribution. Clearly, $\alpha(l|k) \geq 0$. Also, using the formula for x_{l-1} , we can rewrite the numerator in equation (12) as

$$\begin{aligned} Q_{l-1} x_{l-1} \left(\frac{x_{l-1}}{x_l} \right)^{1/(l-1)} - Q_{l-1} x_{l-1} \left(\frac{x_{l-1}}{x_{l-1}} \right)^{1/(l-2)} \\ = Q_{l-1} x_{l-1} \left(\frac{x_{l-1}}{x_l} \right)^{1/(l-1)} - Q_{l-2} x_{l-2} \left(\frac{x_{l-2}}{x_{l-1}} \right)^{1/(l-2)}. \end{aligned}$$

The sum of these terms across l is precisely the denominator in the definition of $\alpha(l|k)$. Together with the remarks after the definition of $\bar{\pi}$, this proves that the information structure is well defined.

We next verify that this information structure and these strategies induce $\bar{S}$. When the price count is k , the probability that the highest price is at least x is

$$\begin{aligned} & \frac{1}{\beta_k} \sum_{l=1}^k \binom{k}{l} (G(x|k) \alpha(k|k))^l (1 - \alpha(k|k))^{k-l} \\ &= \frac{1}{\beta_k} [(1 - \alpha(k|k) + \alpha(k|k) G(x|k))^k - (1 - \alpha(k|k))]. \end{aligned}$$

This follows from the binomial theorem: note that

$$\alpha(k|k) = 1 - \left(\frac{x_k}{x_{k-1}} \right)^{1/(k-1)},$$

and so

$$\beta_k = 1 - \left(\frac{x_k}{x_{k-1}} \right)^{k/(k-1)}.$$

As a result, the conditional probability that the lowest price is at least x reduces to

$$\begin{aligned} & \frac{1}{\beta_k} \left\{ \left(\frac{x_k}{x_{k-1}} \right)^{1/(k-1)} + \left[\left(\frac{x_k}{x} \right)^{1/(k-1)} - \left(\frac{x_k}{x_{k-1}} \right)^{1/(k-1)} \right] - \left(\frac{x_k}{x_{k-1}} \right)^{1/(k-1)} \right\} \\ &= \frac{1}{\beta_k} \left[\left(\frac{x_k}{x} \right)^{1/(k-1)} - \left(\frac{x_k}{x_{k-1}} \right)^{1/(k-1)} \right], \end{aligned}$$

which is $\bar{S}(x)$, as desired.

Finally, we show that these strategies are an equilibrium. We first consider a firm i who receives a signal k and sets a price $p_i \geq x_k$. Then p_i could be the lowest price only if the price count is $k = k$; for the price count must be at least k , and if it were strictly greater, then some firm would have a signal greater than k and be pricing strictly less than x_k . Note also that conditional on getting a signal k , the other firms' signals are conditionally independent draws from $\{1, \dots, k\}$ according to probabilities α , so that the others' prices are conditionally independent draws from

$$\hat{G}(p_j|k) \triangleq \sum_{k'=1}^k \alpha(k'|k) G(p_j|k').$$

Thus, if $p_j \in [x_{k'}, x_{k'-1}]$, then this reduces to

$$\begin{aligned} \hat{G}(p_j|k) &= \alpha(k'|k) G(p_j|k') + \sum_{l=1}^{k'-1} \alpha(l|k) \\ &= \frac{Q_{k'-1} x_{k'-1} [(x_{k'-1}/x_{k'})^{1/(k'-1)} - 1]}{Q_{k-1} x_{k-1} (x_{k-1}/x_k)^{1/(k-1)}} \frac{(x_{k'-1}/p_j)^{1/(k'-1)} - 1}{(x_{k'-1}/x_{k'})^{1/(k'-1)} - 1} + \frac{Q_{k'-1} x_{k'-1}}{Q_{k-1} (x_{k-1}/x_k)^{1/(k-1)} x_{k-1}} \\ &= \frac{Q_{k'-1} x_{k'-1} (x_{k'-1}/p_j)^{1/(k'-1)}}{Q_{k-1} x_{k-1} (x_{k-1}/x_k)^{1/(k-1)}}. \end{aligned}$$

As a result, expected revenue from offering $p_i \in [x_{k'}, x_{k'-1}]$ for $k' \leq k$ is

$$p_i(\hat{G}(p_j|k))^{k-1} \propto (p_i)^{1-(k-1)/(k-1)},$$

which is constant in p_i when $k' = k$ and decreasing in p_i for $k' < k$.

The last step is to verify that firm i does not want to cut prices to $p_i \in [x_{k'}, x_{k'-1}]$ with $k' > k$. Note that conditional on being quoted, the conditional likelihood that the price count is l is

$$\begin{aligned} \frac{\sum_{\{K \subseteq N | i \in K\}} \nu(K|l)\mu(l)}{\sum_{l=1}^n \sum_{\{K \subseteq N | i \in K\}} \nu(K|l)\mu(l)} &= \frac{(l/n)\mu(l)}{\sum_{l=1}^n (l'/n)\mu(l')} \\ &= \frac{l\mu(l)}{\sum_{l=1}^n l'\mu(l')}. \end{aligned}$$

Now, a firm makes a sale in that event only if the equilibrium sales price is at least p_i and they are quoted. The ex ante likelihood of this happening and firm i getting a signal k is proportional to

$$\begin{aligned} D(p_i, k) &\triangleq \frac{\alpha(k|k')}{\beta_{k'}} \mu(k') k' \sum_{l=1}^{k'-1} \binom{k'-1}{l} (\alpha(k'|k') G(p_i|k'))^l (1 - \alpha(k'|k'))^{k'-1-l} \\ &\quad + \sum_{l=k+1}^{k'-1} \frac{\alpha(k|l) [1 - (1 - \alpha(l|l))^{l-1}]}{\beta_l} \mu(l) l + \frac{\alpha(k|k)}{\beta_k} \mu(k) k. \end{aligned}$$

This expression requires some explanation. It is a sum of probabilities of different price counts, times the probability that firm i receives a signal k , and times the probability of making a sale with a price of p_i conditional on the number of firms. Importantly, the signal distribution depends only on the price count and not on the particular set K of quoted firms. The first line gives the probability that firm i gets a signal k when there are $k' > k$ firms and the sales price is at least p_i . Note that the number of firms other than i with a signal of k' is binomially distributed, conditional on that number being at least 1. Conditional on there being l firms with a signal of k' , the likelihood of p_i being the lowest price is $(G(p_i|k'))^l$. The second term is the likelihood that the number of firms l is between k' and k , firm i gets a signal of k , and at least one of the other firms gets a signal of k . The final term is the likelihood that the number of firms is k and firm i gets a signal of k (in this last event, the signals of the other firms are unrestricted).

We can simplify terms in $D(p_i, k)$ as follows. Using the binomial theorem,

$$\begin{aligned} &\sum_{l=1}^{k'-1} \binom{k'-1}{l} (\alpha(k'|k') G(p_i|k'))^l (1 - \alpha(k'|k'))^{k'-1-l} \\ &= (1 - \alpha(k'|k') + \alpha(k'|k') G(p_i|k'))^{k'-1} - (1 - \alpha(k'|k'))^{k'-1} \\ &= \frac{x_{k'}}{p_i} - \frac{x_{k'}}{x_{k'-1}}. \end{aligned}$$

Let us define

$$A(k) \triangleq Q_{k-1} x_{k-1} \left[\left(\frac{x_{k-1}}{x_k} \right)^{1/(k-1)} - 1 \right].$$

Then

$$\begin{aligned} \frac{\alpha(k|l)\mu(l)l}{\beta_l} &= \frac{A(k)}{Q_{l-1} x_{l-1} (x_{l-1}/x_l)^{1/(l-1)}} \frac{Q_l - Q_{l-1}}{[1 - (x_{l-1}/x_l)^{1/(l-1)}]} \\ &= \frac{A(k)}{x_{l-1} (x_{l-1}/x_l)^{1/(l-1)} [1 - (x_{l-1}/x_l)^{1/(l-1)}]} \left(\frac{Q_l}{Q_{l-1}} - 1 \right) \\ &= \frac{A(k)}{x_{l-1} (x_{l-1}/x_l)^{1/(l-1)} [1 - (x_{l-1}/x_l)^{1/(l-1)}]} \left[\left(\frac{x_l}{x_{l-1}} \right)^{1/(l-1)} - 1 \right] \\ &= \frac{A(k)}{x_l}. \end{aligned}$$

Finally,

$$1 - (1 - \alpha(k|k))^{k-1} = 1 - \frac{x_k}{x_{k-1}}.$$

Substituting in these expressions, we can rewrite $D(p_i, k)$ as

$$\begin{aligned} D(p_i, k) &= A(k) \left[\frac{1}{x_k} \left(\frac{x_k}{p_i} - \frac{x_k}{x_{k-1}} \right) + \sum_{l=k+1}^{k'-1} \frac{1}{x_l} \left(1 - \frac{x_l}{x_{l-1}} \right) + \frac{1}{x_l} \right] \\ &= \frac{A(k)}{p_i}. \end{aligned}$$

Thus, the payoff from a price cut to $p_i \in [x_k, x_{k-1}]$ for $k' > k$ is $p_i D(p_i, k) = A(k)$, thus verifying that the proposed strategies are an equilibrium. QED

A1.4. Proof of Proposition 4

We show the result for the case when μ is obtained from μ' by shifting a mass of ϵ from $k+1$ to k , that is,

$$\mu(k) = \begin{cases} \mu'(k) + \epsilon & \text{if } k = k, \\ \mu'(k+1) - \epsilon & \text{if } k = k+1, \\ \mu'(k) & \text{otherwise.} \end{cases}$$

Any μ that is first-order stochastically dominated by μ' can be obtained via a finite sequence of such shifts, so that this special case implies the general result in the statement of the proposition.

To that end, let $\{S'(\cdot|k)\}_{k=1}^n$ be conditional distributions that satisfy inequality (9) for the price count distribution μ' . Let us define

$$S(x|l) = \begin{cases} \frac{\mu(k) - \epsilon}{\mu(k)} S'(x|k) + \frac{\epsilon}{\mu(k)} S'(x|k+1) & \text{if } l = k, \\ S'(x|l) & \text{otherwise.} \end{cases}$$

The induced ex ante distribution is precisely S , so that the right-hand side of (9) is unchanged. But the left-hand side is now

$$\begin{aligned} x \sum_{l=1}^n \mu(l) l S(x|l) &= x \left[\sum_{l=1}^n \mu'(l) l S'(x|l) - \epsilon(k+1) S'(x|k+1) + \epsilon k S'(x|k+1) \right] \\ &= x \left(\sum_{l=1}^n \mu'(l) l S'(x|l) - \epsilon S'(x|k+1) \right) \\ &\leq x \sum_{l=1}^n \mu'(l) l S'(x|l). \end{aligned}$$

The left-hand side has decreased, so that $\{S(\cdot|k)\}_{k=1}^n$ satisfies inequality (9) for all x . Thus, any ex ante sales price distribution that deters uniform price cuts for μ' also deters downward uniform deviations for μ . A fortiori, the bounding distribution $\bar{S}$ for μ must first-order stochastically dominate $\bar{S}'$. QED

A2. Proofs for Section 5

Before proving theorem 2, we first complete the description of the extensive-form game from section V.B. Let us define

$$b_k(\sigma, \theta, \xi, t, p) = \left(\prod_{l < k} \sigma \left(\theta, \{\xi(j), t_{\xi(j)}, p_{\xi(j)}\}_{j=1}^l \right) \right) \left(1 - \sigma \left(\theta, \{\xi(j), t_{\xi(j)}, p_{\xi(j)}\}_{j=1}^k \right) \right)$$

to be the probability that the consumer buys after searching k firms, when using the strategy σ , conditional on the realized type, prices, signals, and order. Then firm i 's expected payoff conditional on (θ, p, t, ξ) , is

$$R_i(\sigma, \theta, p, t, \xi) = p_i \sum_{k=1}^n \frac{\mathbb{I}_{i \in K(p_{\xi(1)}, \dots, p_{\xi(k)})}}{|K(p_{\xi(1)}, \dots, p_{\xi(k)})|} b_k(\sigma, \theta, \xi, t, p).$$

This is the price set by firm i times the probability that at the time the consumer stops, firm i has been searched, has a low price, and wins any tie breaks. Given the strategy profile (F, σ) , firm i 's payoff is then

$$R_i(F, \sigma) = \int_{\Theta} \frac{1}{n!} \sum_{\xi \in \Xi} \sum_{t \in T} \int_{[0, v]^n} R_i(\sigma, \theta, p, t, \xi) F(dp|t) \pi(t|\theta, \xi) \eta(d\theta).$$

In addition, let

$$U(\sigma, \theta, \xi, t, p) = \sum_{k=1}^n (v - \min\{p_{\xi(1)}, \dots, p_{\xi(k)}\} - c(k, \theta)) b_k(\sigma, \theta, \xi, t, p)$$

be the payoff to the consumer conditional on $(\sigma, \theta, \xi, t, p)$. The consumer's ex ante equilibrium payoff is

$$U(F, \sigma) = \int_{\Theta} \frac{1}{n!} \sum_{\xi \in \Xi} \sum_{t \in T} \int_{[0, v]^n} U(\sigma, \theta, \xi, t, p) F(dp|t) \pi(t|\theta, \xi) \eta(d\theta).$$

The price count distribution induced by (F, σ) is

$$\mu(k) = \int_{\Theta} \frac{1}{n!} \sum_{\xi \in \Xi} \sum_{t \in T} \int_{[0, v]^n} b_k(\sigma, \theta, \xi, t, p) F(dp|t) \pi(t|\theta, \xi) \eta(d\theta).$$

Finally, the induced sales price distributions for every firm i are

$$S_i(x|k) = \frac{1}{\mu(k)} \int_{\Theta} \frac{1}{n!} \sum_{\{\xi \in \Xi | \xi(i) = k\}} \sum_{t \in T} \int_{[x, v]^n} \frac{\mathbb{I}_{i \in K(p_{\xi(1), \dots, \xi(n)})}}{|K(p_{\xi(1), \dots, \xi(k)})|} b_k(\sigma, \theta, \xi, t, p) F(dp|t) \pi(t|\theta, \xi) \eta(d\theta),$$

and, summing up over all firms,

$$S(x|k) = \sum_{i=1}^n S_i(x|k),$$

and all price counts,

$$S(x) = \sum_{k=1}^n \mu(k) S(x|k).$$

Now, fix a number $k \in \{0, \dots, n-1\}$ and a history $h \in H_k$. Given a history

$$h = \left(\theta, \{\xi(k), t_{\xi(k)}, p_{\xi(k)}\}_{l=1}^k \right),$$

we write $\Xi(h)$, $T(h)$, and $P(h)$ for the orderings, type profiles, and price profiles consistent with history h , respectively:

$$\Xi(h) = \{\xi' \in \Xi | \xi'(l) = \xi(l) \ \forall l = 1, \dots, k\},$$

$$T(h) = \{t' \in T | t'_{\xi(l)} = t_{\xi(l)} \ \forall l = 1, \dots, k\},$$

$$P(h) = \{p' \in [0, v]^n | p'_{\xi(l)} = p_{\xi(l)} \ \forall l = 1, \dots, k\}.$$

We write $U(F, \sigma, h)$ for the consumer's payoff, conditional on history h being reached:

$$U(F, \sigma, h) = \frac{1}{|\Xi(h)|} \sum_{\xi \in \Xi(h)} \frac{1}{\sum_{t \in T(h)} \pi(t|\theta, \xi)} \\ \times \sum_{t \in T(h)} \int_{p \in P(h)} U(\sigma, \theta, \xi, t, p) F_{\xi(k+1), \dots, \xi(n)}(dp_{\xi(k+1), \dots, \xi(n)} | t) \pi(t|\theta, \xi).$$

Now firm strategy F_i is a best response to (F_{-i}, σ) if $R_i(F_i, F_{-i}, \sigma) \geq R_i(F'_i, F_{-i}, \sigma)$ for all F'_i . The consumer's strategy σ is *sequentially rational* with respect to F if $U(F, \sigma, h) \geq U(F, \sigma', h)$ for all $k = 0, \dots, N-1$, $h \in H_k$, and σ' . (Note that the condition on the consumer's strategy implies, in the case where $k = 0$, that the consumer's strategy must be an ex ante best response.)²⁰ The strategy profile (F, σ) is a *perfect Bayesian equilibrium* if F_i is a best response to (F_{-i}, σ) for all i and if σ is sequentially rational.

²⁰ This definition builds in the restriction on out-of-equilibrium beliefs that the consumer uses the same distribution $\pi(\cdot|\theta, \xi)$ over signals and the same conditional distributions $F(\cdot|s)$ over prices as are used on path. We assumed that signal sets are finite in order to simplify the statement of these conditional payoffs.

At a history $h \in H_k$, let

$$p(h) = \min\{p_{\xi(1)}, \dots, p_{\xi(k)}\}$$

denote the lowest price quoted thus far. We further define $\tilde{H}_k$ as the set of histories of length k excluding prices, that is, the set whose elements are of the form $(\theta, \{\xi(k), t_{\xi(l)}\}_{l=1}^k)$. The union of the sets $\tilde{H}_k$ across $k < m$ is denoted $\tilde{H}$.

Our first result is the following proposition.

PROPOSITION 6 (Reservation price). Given the firms' strategies F , there exists a reservation price function $r: \tilde{H} \rightarrow \mathbb{R}$, such that a strategy for the consumer is sequentially rational if and only if $\sigma(h) = 0$ if $p(h) < r(\tilde{h})$ and $\sigma(h) = 1$ if $p(\tilde{h}) > r(\tilde{h})$.

Proof. The result is established in three steps. First, holding fixed F , the consumer's payoff is continuous in σ , so that there is an optimal strategy. Thus, for every history h , there is a value

$$V(h) = \max_{\sigma} U(F, \sigma, h)$$

generated by the consumer's optimal continuation strategy, and the value function must satisfy the following Bellman equation:

$$V(h) = \max\{v - p(h) - c(k, \theta), \mathbb{E}[V((h, \xi(k+1), t_{\xi(k+1)}, p_{\xi(k+1)}))|h]\}.$$

We claim that, in fact, V depends only on $\tilde{h}$ and the lowest price quoted thus far, $p(h)$. The reason is by induction on the length of the history. At histories in H_m , this is obviously true, since $V(h) = v - c(n, \theta) - p(h)$. Now assume that the inductive hypothesis holds for $h \in \tilde{H}_l$ for $l > k$. Then at the history $h \in H_k$, the consumer can either stop and receive a payoff $v - p - c(k, \theta)$ or continue and receive a payoff of

$$\mathbb{E}_{(\xi(k+1), t_{\xi(k+1)}, p_{\xi(k+1)})} [V((\tilde{h}, \xi(k+1), t_{\xi(k+1)}, \min\{p(h), p_{\xi(k+1)}\})|h)],$$

where we have used the fact that the next period's value depends only on the nonprice history and the lowest price. Critically, the distribution of the next firm's price depends only on their identity and signal, and the distribution of the next firm's identity and signal depend only on the current nonprice history $\tilde{h}$ and not on past prices. Thus, it is without loss to condition on $\tilde{h}$ rather than h , so that $V(h)$ depends only on $\tilde{h}$ and $p(h)$.

Second, we argue that $V(\tilde{h}, p)$ is decreasing and convex in p , with a slope at least -1 . It is obvious that $V(\tilde{h}, p)$ is decreasing in p . The reason is that under the optimal strategy σ , each of the terms in $U(F, \sigma, h)$ is decreasing in p . We further claim that the slope of $V(\tilde{h}, p)$ with respect to p is -1 times the probability that p is the lowest price at the time the consumer decides to stop, under the optimal continuation strategy. This is immediate from the fact that p enters the consumer's payoff $U(F, \sigma, h)$ only if the consumer purchases at this price. Thus, the slope is strictly greater than -1 , unless there is probability one that the consumer will purchase at this price. Moreover, we claim that V is convex in p . This is established by induction. Clearly, V is convex in p for histories in $\tilde{H}_n$. Inductively, the expected payoff from continuing to search is

$$\mathbb{E}_{(\xi(k+1), t_{\xi(k+1)}, p_{\xi(k+1)})} [V((\tilde{h}, \xi(k+1), t_{\xi(k+1)}), \min\{p, p_{\xi(k+1)}\}) | \tilde{h}]. \quad (\text{A2})$$

Each term in this expectation is clearly convex, since it is decreasing and convex in $p(h) = p$ when $p(h) < p_{\sigma(k+1)}$ and is constant when the reverse inequality holds. Thus, the payoff from continuing is convex, and the payoff from stopping is linear in p , so that the maximum of these convex functions is also convex. This extends the inductive hypothesis to k .

Now we argue for the existence of reservation prices for $\tilde{h} \in \tilde{H}_k$ with $k < n$. If there is a $k' > k$ with $c(k', \theta) < c(k, \theta)$, then since the consumer can recall past prices, it is strictly optimal to continue searching, and we can set $r(\tilde{h}) = -\infty$. Now suppose that $c(k', \theta) > c(k, \theta)$ for all $k' > k$. Clearly, the payoff from stopping has a slope of -1 in the current lowest price and, since $c(k', \theta) > c(k, \theta)$ for $k' > k$, it is strictly optimal for the consumer to stop if the lowest price is zero. If the consumer is ever indifferent between stopping and continuing at some price p , then it must be because the payoff from stopping and the expected payoff from continuing to search ([A2]) have crossed. But this can happen only if the slope of (A2) is strictly greater than -1 . As the expected payoff from continuing is convex, we conclude that for $p' > p$, the slope of (A2) is also strictly greater than -1 . This implies that the expression (A2) is strictly greater than the payoff from stopping for all $p' > p$. As a result, there is at most one point where the two payoffs cross, which is denoted $r(\tilde{h})$, or if they never cross, we let $r(\tilde{h})$ be any negative number. QED

Given that the consumer's equilibrium behavior is characterized by cutoffs $r(\tilde{h})$, we now argue that constraint (9) must be satisfied by the equilibrium sales price distribution. The following result extends proposition 1 to the sequential search model.

PROPOSITION 7 (Sequential search and sales price distribution). Suppose that the sequential search model $\{\Theta, \eta, c, T, \pi\}$ and equilibrium (F, σ) induce the price count distribution $\mu \in \Delta(\{1, \dots, N\})$. Then the induced ex ante sales price distribution $S(\cdot)$ deters uniform price cuts.

Note that we have assumed that consumers search at least one firm. If costs were sufficiently high, the consumer might never search, and the price count distribution would assign probability 1 to a price count of zero.²¹ The proposition remains vacuously true in this case.

Proof. Fix a terminal history $h = (\theta, \xi, t, p)$. Suppose that firm i deviates to $p'_i < p_i$. The resulting payoff is

$$p'_i \sum_{k=\xi^{-1}(i)}^N \frac{\mathbb{I}_{i \in K((p'_i, p_{-i})_{\xi(1), \dots, \xi(k)})}}{|K((p'_i, p_{-i})_{\xi(1), \dots, \xi(k)})|} b_k(\sigma, \theta, \xi, t, (p'_i, p_{-i})). \quad (\text{A3})$$

On the other hand, if the firm deviated from p_i to p'_i but the consumer did not adjust behavior, the payoff would be

$$p'_i \sum_{k=\xi^{-1}(i)}^N \frac{\mathbb{I}_{i \in K((p'_i, p_{-i})_{\xi(1), \dots, \xi(k)})}}{|K((p'_i, p_{-i})_{\xi(1), \dots, \xi(k)})|} b_k(\sigma, \theta, \xi, t, p). \quad (\text{A4})$$

²¹ Janssen, Moraga-González, and Wildenbeest (2005) incorporate this possibility into the model of Stahl (1989).

(Note that we have dropped terms where the consumer stops searching before reaching firm i .) We claim that (A3) is greater than (A4). To see why, observe that by proposition 6, $\sigma(h)$ is weakly increasing in the lowest price. Let

$$B_k(\sigma, \theta, \xi, t, p) \triangleq \sum_{l=k}^n b_l(\sigma, \theta, \xi, t, p)$$

denote the probability that the consumer searches at least k firms. Then clearly,

$$B_k(\sigma, \theta, \xi, t, p) = \prod_{l \leq k} \sigma\left(\theta, \{\xi(j), t_{\xi(j)}, p_{\xi(j)}\}_{j=1}^l\right),$$

so that $B_k(\sigma, \theta, \xi, t, p) = B_k(\sigma, \theta, \xi, t, (p'_i, p_{-i}))$ for $k < \xi^{-1}(i)$, and $B_k(\sigma, \theta, \xi, t, p) \leq B_k(\sigma, \theta, \xi, t, (p'_i, p_{-i}))$ for $k \geq \xi^{-1}(i)$. Thus, the distribution of the stopping time when the consumer responds is first-order stochastically dominated by the stopping-time distribution when the consumer does not respond. The result then follows from the fact that

$$\frac{\prod_{i \in K} ((p'_i, p_{-i})_{\xi(1), \dots, \xi(k)})}{|K((p'_i, p_{-i})_{\xi(1), \dots, \xi(k)})|}$$

is decreasing in k for $k \geq \xi^{-1}(i)$.

Thus, firm i 's payoff from a price cut is higher when the consumer responds than when the consumer does not respond, conditional on (θ, ξ, t, p) . As a result, the interim payoff from the price cut, taking expectation across (θ, ξ, t, p_{-i}) , is also higher when the consumer's search strategy responds (so that price counts adjust) than when the consumer does not respond (so that the price count distribution is the same). Since expression (10) was computed under the premise that price counts do not respond, it must be that the firm's surplus from a uniform price cut is weakly greater than (10). As a result, inequality (9) must still be satisfied by the conditional price distributions that can be generated in equilibrium. QED

References

- Abrantes-Metz, Rosa M., Luke M. Froeb, John Geweke, and Christopher T. Taylor. 2006. "A Variance Screen for Collusion." *Internat. J. Indus. Org.* 24 (3): 467–86.
- Armstrong, Mark, and John Vickers. 2001. "Competitive Price Discrimination." *RAND J. Econ.* 32 (4): 579–605.
- . 2019. "Discriminating against Captive Consumers." *A.E.R.: Insights* 1 (3): 257–72.
- . 2020. "Patterns of Price Competition and the Structure of Consumer Choice." Working paper, All Souls Coll., Univ. Oxford.
- Āzadis, Helmut, and Péter Vida. 2015. "Collusive Communication Schemes in a First-Price Auction." *Econ. Theory* 58 (1): 125–60.
- Bajari, Patrick, and Lixin Ye. 2003. "Deciding between Competition and Collusion." *Rev. Econ. and Statis.* 85 (4): 971–89.
- Baye, Michael R., Dan Kovenock, and Casper G. de Vries. 1992. "It Takes Two to Tango: Equilibria in a Model of Sales." *Games and Econ. Behavior* 4 (4): 493–510.
- Baye, Michael R., and John Morgan. 2001. "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets." *A.E.R.* 91 (3): 454–74.

- Baye, Michael R., John Morgan, and Patrick Scholten. 2006. "Information, Search and Price Dispersion." In *Economics and Information Systems*, edited by Terrence Hendershott, 323–76. Bingley: Emerald.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2013. "Extremal Information Structures in the First Price Auction." Cowles Found. Discussion Paper no. 1926, Yale Univ.
- . 2015. "First Price Auctions with General Information Structures: Implications for Bidding and Revenue." Cowles Found. Discussion Paper no. 2018, Yale Univ.
- . 2017. "First-Price Auctions with General Information Structures: Implications for Bidding and Revenue." *Econometrica* 85 (1): 107–43.
- . 2020. "Competition and Public Information: A Note." Cowles Found. Discussion Paper no. 2234, Yale Univ.
- Bergemann, Dirk, and Stephen Morris. 2013. "Robust Predictions in Games with Incomplete Information." *Econometrica* 81 (4): 1251–308.
- . 2016. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Econ.* 11 (2): 487–522.
- Burdett, Kenneth, and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51 (4): 955–69.
- Butters, Gerard R. 1977. "Equilibrium Distribution of Sales and Advertising Prices." *Rev. Econ. Studies* 44 (3): 465–91.
- de Clippel, Geoffroy, Kfir Eliaz, and Kareen Rozen. 2014. "Competing for Consumer Inattention." *J.P.E.* 122 (6): 1203–34.
- De los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest. 2012. "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." *A.E.R.* 102 (6): 2955–80.
- Diamond, Peter A. 1971. "A Model of Price Adjustment." *J. Econ. Theory* 3 (2): 156–68.
- Ellison, Glenn, and Sara Fisher Ellison. 2009. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica* 77 (2): 427–52.
- Ellison, Glenn, and Alexander Wolitzky. 2012. "A Search Cost Model of Obfuscation." *RAND J. Econ.* 43 (3): 417–41.
- Fang, Hanming, and Stephen Morris. 2006. "Multidimensional Private Value Auctions." *J. Econ. Theory* 126:1–30.
- Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. Cambridge, MA: MIT Press.
- Hong, Han, and Matthew Shum. 2006. "Using Price Distributions to Estimate Search Costs." *RAND J. Econ.* 37 (2): 257–75.
- Janssen, Maarten C. W., José Luis Moraga-González, and Matthijs R. Wildenbeest. 2005. "Truly Costly Sequential Search and Oligopolistic Pricing." *Internat. J. Indus. Org.* 23 (5–6): 451–66.
- Myatt, David P., and David Ronayne. 2019. "A Theory of Stable Price Dispersion." Working paper, London Bus. School.
- Narasimhan, Chakravarthi. 1988. "Competitive Promotional Strategies." *J. Bus.* 61 (4): 427–49.
- Stahl, Dale O. 1989. "Oligopolistic Pricing with Sequential Consumer Search." *A.E.R.* 79 (4): 700–712.
- . 1996. "Oligopolistic Pricing with Heterogeneous Consumer Search." *Internat. J. Indus. Org.* 14 (2): 243–68.
- Varian, Hal R. 1980. "A Model of Sales." *A.E.R.* 70 (4): 651–59.