

DISCUSSION PAPER SERIES

DP14521

SEARCH, INFORMATION AND PRICES

Stephen Morris, Dirk Bergemann and Benjamin A
Brooks

INDUSTRIAL ORGANIZATION



SEARCH, INFORMATION AND PRICES

Stephen Morris, Dirk Bergemann and Benjamin A Brooks

Discussion Paper DP14521
Published 21 March 2020
Submitted 19 March 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Stephen Morris, Dirk Bergemann and Benjamin A Brooks

SEARCH, INFORMATION AND PRICES

Abstract

Consider a market with many identical firms offering a homogeneous good. A consumer obtains price quotes from a subset of firms and buys from the firm offering the lowest price. The "price count" is the number of firms from which the consumer obtains a quote. For any given ex ante distribution of the price count, we obtain a tight upper bound (under first-order stochastic dominance) on the equilibrium distribution of sale prices. The bound holds across all models of firms' common-prior higher-order beliefs about the price count, including the extreme cases of complete information (firms know the price count exactly) and no information (firms only know the ex ante distribution of the price count). A qualitative implication of our results is that even a small ex ante probability that the price count is one can lead to dramatic increases in the expected price. The bound also applies in a wide class of models where the price count distribution is endogenized, including models of simultaneous and sequential consumer search.

JEL Classification: D41, D42, D43, D83

Keywords: search, price competition, Bertrand Competition, "Law of One Price", Price Count, Price Quote, information structure, Bayes correlated equilibrium

Stephen Morris - semorris@mit.edu
M.I.T. and CEPR

Dirk Bergemann - dirk.bergemann@yale.edu
Yale University and CEPR

Benjamin A Brooks - babrooks@uchicago.edu
University of Chicago

Acknowledgements

We acknowledge financial support through NSF Grants SES 1459899. We thank Glenn Ellison, Ben Golub, Jason Hartline and Phil Reny for helpful conversations.

1 Introduction

1.1 Background and Motivation

When two or more identical firms engage in Bertrand competition, the standard prediction of economic theory is that the price of a homogenous good will be competed down to cost. The model is therefore consistent with the “law of one price.” But as Varian (1980) noted forty years ago, “the law of one price is not a law at all,” and price dispersion even for homogenous goods seems to be an ubiquitous feature of the market economy. Since the time of Varian’s writing, a large literature has developed various equilibrium theories of the price distribution that are consistent with the failure of the law of one price. An essential feature of these theories is that firms face uncertainty about the amount of competition and, in particular, that there is not common knowledge of whether there are at least two firms competing for a given consumer. Many explanations have been proposed for why common knowledge would fail, including unobserved consumer search, advertising, or because of the influence of intermediaries on the flow of information in the market, and we discuss these theories and evidence below.

The success of this literature in providing so many theories of the equilibrium price distribution is also a curse. Which of the various mechanisms driving the price distribution is empirically relevant and should be used for understanding behavior in a particular market? Even if we fix the theoretical model, the sale price distribution that arises in equilibrium often depends on features of preferences which may be difficult or impossible to measure, such as the psychological cost to consumers of obtaining price quotes. Finally, for the sake of tractability, many of these models make strong simplifying assumptions about firms’ knowledge of how much competition they face. In particular, much of the literature assumes that firms only know the *ex ante* distribution of the number of firms competing for a given consumer. This assumption is extreme: while it allows for a non-degenerate price distribution, we will see that it implies that the expected sale price is the same as it would be if firms had complete information about the number of competing firms.¹ Thus, the standard search models explain non-classical price distributions, but they cannot explain non-classical average price levels.

In this paper, we develop predictions about the equilibrium distribution of sale prices that hold across a wide variety of models. A central object in models of the price distribution is the number of price quotes that a consumer receives, which we refer to as the *price count*. Price dispersion is driven by firms’ uncertainty about the consumer’s price count.

¹In fact, it is the same surplus the firm would obtain if they only targeted captive consumers, for whom the firm is a monopolist. See Corollary 3 below.

Rather than modeling the origins of the price count, we instead take as a primitive the distribution of the price count (something that is now being measured in empirical work discussed below). We are agnostic about how much firms know about the price count: there may be complete information about the price count, as in the standard models of monopoly and perfect Bertrand competition; no information as in the simultaneously search models of Varian (1980) and Burdett and Judd (1983); or any other form of common-prior higher order beliefs. For any such model of beliefs and a Bayesian equilibrium, the firms' strategies induce a distribution over the *sale price*, i.e., the lowest price offered by the firms and the price at which trade occurs. It is the sale price distribution that determines consumer and producer welfare.

1.2 Results

Our main result, Theorem 1, is a tight upper bound on the equilibrium sale price distribution, in the sense of first-order stochastic dominance. The bound holds across all common-prior beliefs consistent with the given price count distribution. This bound immediately translates into an upper bound on producer surplus and a lower bound on consumer surplus.

Theorem 1 provides an empirical test for collusion: If the observed sale price distribution is not below the bound corresponding to the observed price count distribution, then prices cannot be explained by Bertrand competition under incomplete information. The theorem 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 $\sqrt{\mu(2-\mu)}$ of monopoly revenue. Thus, a qualitative finding is that if we allow for firms to have partial information about the price count, producer surplus is non-linear in the probability of monopoly, and in fact, marginal revenue in the probability of monopoly is unbounded near $\mu = 0$. Under either no information or complete information, revenue is linear in μ . Thus, a small amount of monopoly power may translate into rents for firms that are much larger than that suggested by either information benchmark.

The bound we construct is based on the following logic: If the distribution of prices 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 non-trivial bounds on how high the price distribution can go, and that the critical equilibrium constraints are those associated with cutting prices. We focus on a particular class of deviations, whereby for some fixed price level, firms deviate by pricing at the minimum of that level and whatever price they would have set in equilibrium. We refer to this as a *uniform*

downward deviation. We show that the requirement that firms not want to uniformly deviate downward can be expressed as a constraint on the sale price distribution. We further show that there is a highest sale price distribution that satisfies all of the uniform downward incentive constraints, which is associated with firms being indifferent to *all* uniform downward deviations.

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 which is a lower bound on the realized price count, and at least one firm observes the true price count. One 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 decrease in the number of other firms they observe, so that the sale is always made by a firm who observes the highest price count.

Theorem 1 takes the price count distribution as given. A critical question is whether the bound still applies when the price count is endogenized. Immediately after proving the Theorem 1, we consider a variety of extensions, suggested by the literature, that microfound the price count distribution. These extensions can be broadly grouped into two categories. 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 equilibrium strategy. In this case, we say that the model has *no feedback*. Any model in which price counts and prices are simultaneously determined, such as Burdett and Judd (1983), falls into this category. Our bounds immediately apply to any model with no feedback, because the critical deviation that pins down our bound, in which firms cut prices without affecting price counts is still available, and generates the same surplus for the deviator as in our baseline model.

The other category of extensions are those in which there is *feedback* from firms' realized prices to price counts and other firms' prices. As a result, a price deviation could result in a different surplus for the deviator than if we held fixed the equilibrium joint distribution of price counts and other firms' prices. The simplest examples of models with feedback are Stackelberg pricing, in which the leading firm's price directly affects the follower's price, and sequential search, such as Stahl (1989, 1996), in which prices affect the price count through consumer's search behavior. Our bounds do not hold for the Stackelberg game, but they do apply in the sequential search model, given some additional assumptions. A key question is whether price cuts cause price counts to increase or decrease. In sequential search, there is a tendency for lower prices to lead to less search, and hence lower price counts. This

effect makes the uniform downward deviations *more* attractive, so that the inequalities that lead to our bound must still hold. We present conditions under which this logic is exactly correct. We also discuss plausible scenarios where consumers learn about market conditions from prices in such a way that lower prices could lead to more search, in which case our bounds might no longer hold.

After endogenizing the price count, we consider various extensions. We use our methodology to analyze what happens when firms have no information beyond the ex ante distribution of the price count. Here we show that there is a unique symmetric equilibrium under which the expected price is the same as under complete information. The induced sale price distribution is less dispersed than under complete information. Thus, there is no tight lower bound on the sale price distribution corresponding to our tight upper bound. We can, however, characterize a tight lower bound on the expected price, which is attained under no information and complete information. Finally, largely for notational convenience, we carry out our analysis for the case where there is single-unit demand and firms are all equally likely to be quoted. However, our upper bound on the sale price generalizes to downward sloping demand and asymmetric probabilities of being quoted, although the bound is not necessarily attained with asymmetric quoting.

1.3 Related Literature

The last fifty years has seen an explosion of theoretical work on understanding the failure of law of one price, by providing micro-foundations for the failure of common knowledge that the price count is at least two. This failure can occur even with free entry and low costs for the consumers to search, and low costs for sellers to publicly post prices or information intermediaries to collect, advertise and distribute prices: see, for example, Stigler (1961), Diamond (1971), Rothschild (1973), Varian (1980), Burdett and Judd (1983), Stahl (1989; 1996), and Baye and Morgan (2001). These models show the possibility of price dispersion, but are not sharply distinguished. Because this literature typically assumes firms do not have information about the price count, the expected price is the same as in the complete information case. This paper offers a novel perspective on this literature, abstracting from the fine details of how the price count and firms' information about the price count is generated. Instead, we take the price count distribution as given and identify novel tight bounds on equilibrium prices that are consistent with that price count distribution.

There is also a substantial literature on empirical price distributions, which has found that price dispersion has persisted and sometimes increased in the the internet age: see Baye, Morgan, and Scholten (2006) for an early survey. The focus of this literature is on the

dispersion of realized prices in the cross-section. By contrast, we make predictions about the sale price distribution, i.e., the distribution of prices at which sales take place.

Our theoretical result links the price count distribution (however determined) to the sale price distribution. What is known about the price count distribution? The early empirical literature has no information on the price count. But Hong and Shum (2006) show how to identify the price count distribution, under the assumption that firms have no information about the price count and are following equilibrium pricing strategies revealed by the empirical price distribution. They then use the price count distribution to identify the search cost under the assumption that the price count distribution is generated by the equilibrium of a simultaneous search model. Our bounds could be used to test the robustness of this identification to the information structure, i.e., we could establish an upper bound on the price count distribution consistent with observed prices and any information structure. More recently, browser history has been used to provide direct evidence on the price count distribution (e.g., De los Santos, Hortacsu, and Wildenbeest, 2012).

Our model has many firms selling to a single consumer with unknown price count. There is a mathematically equivalent interpretation in which there is a continuum of consumers with heterogeneous price counts, and firms price discriminate based on what they know about consumers' types. If uniform pricing were imposed on firms, then it would be as if they had no information and the equilibrium would correspond to the no-information case we also analyze. Armstrong and Vickers (2019) analyze this problem with this interpretation, examining how prices (and so profits) and consumer surplus vary between price discrimination and uniform pricing. Our model is less general in some dimensions, as we focus on a symmetric count distribution and assume unit demand, whereas they allow asymmetric distributions over price quotes and general demand functions. The latter restriction does not alter the analysis of prices, but it means we have less to say about consumer surplus. Our model is more general in allowing (less importantly) many firms throughout and (more importantly) general asymmetric information structures, whereas Armstrong and Vickers (2019) restrict attention to public information (meaning all firms observe the same public signal about the price count). Under a symmetric price count distribution, they showed that—with two firms—price discrimination raises revenue when firms have public information, while we show that it holds true for arbitrary information. And we identify an information structure with differential and private information under the which the sale price distribution is higher than under any public information structure.

While we frame our results in terms of firms offering prices to consumers, the model can also be used to study the wage distribution in a market where several firms make take-it-or-leave-it employment offers to workers (cf. Mortensen, 2003; Rogerson, Shimer, and Wright,

2005). Under this interpretation, our main result relates the distribution of workers' job search activity (analogous to the price count) to the distribution of wages (analogous to sale prices). Theorem 1 could be used to test whether empirical wage distributions is consistent with competitive wage setting, given workers' labor search behavior.

In solving for what happens in all information structures for a fixed static game, we follow a general methodology described in Bergemann and Morris (2013, 2016). Bergemann, Brooks, and Morris (2017) have applied this methodology to first-price auctions where bidders do not necessarily know their values of the object being sold. The Bertrand pricing game that we analyze is strategically equivalent to a particular first-price procurement auction, where firm know their costs (which is equivalent to knowing their values in a standard auction). Restricting attention to information structures where bidders know their values makes the problem harder to solve, by imposing more constraints on higher-order beliefs. Our current results generalize arguments developed in Bergemann, Brooks, and Morris (2017) to account for the additional informational constraints.²

We also report informational robustness results for an extensive form game: the sequential search game. Makris and Renou (2017) extend the methodology of Bergemann and Morris (2013, 2016) to general extensive form games, and illustrate the approach using a pricing game.

This paper goes beyond the methodology described in Bergemann and Morris (2013, 2016) by establishing results that hold across a wide variety of games. Sutton (1991) has argued that it is a major weakness of modern theoretical industrial organization that results seem sensitive to the exact extensive form used to model a given strategic environment and proposes that theorists find results that are robust to the exact extensive form. Segal and Whinston (2003) is one paper that follows up on this concern, by identifying the robust properties of bilateral contracts independent of the exact process of reaching agreement on contracts. Doval and Ely (2020) provide a general characterization of attainable equilibrium outcomes if any extensive form as well as any information structure is possible, for a given game specifying actions and payoffs. This paper considers more structured classes of alternative games in the particular context of price competition.

The rest of this paper proceeds as follows. Section 2 describes our model of Bertrand pricing under incomplete information. Section 3 illustrates our results by way of a simple two-firm example. Section 4 contains our main result, Theorem 1, which describes a tight upper bound on the equilibrium sale price distribution as a function of the price count distribution. In Section 5, we consider various extensions of the basic model in which the price count distribution is endogenized and for which Theorem 1 continues to hold. Section

²The connection with auction theory is discussed in detail in Section 4.6.

6 discusses several further topics and extensions, and Section 7 is a conclusion. Omitted proofs are in the Appendix.

2 Model

A single consumer has a willingness to pay $v > 0$ for a single unit of a homogeneous good. There are N firms, indexed by $i \in \{1, \dots, N\} = \mathcal{N}$, who can produce the good at zero cost.³ The consumer receives price quotes from a subset of those firms. The *price count* is the number of price quotes n that the consumer receives. The price count is random and we write $\mu \in \Delta(\{1, \dots, N\})$ for the ex ante distribution of the price count. Given the price count n , all firms are equally likely to be quoted.⁴ The set of quoted firms is $\tilde{\mathcal{N}} \subseteq \mathcal{N}$, with each such set having equal likelihood of $1/\binom{N}{n}$, conditional on n . We will focus on the single consumer interpretation of our model, but as usual, there is an alternative interpretation in which there is a continuum of consumers and $\mu(n)$ is the proportion of consumers who obtain n price quotes. We will reference this interpretation occasionally in discussing our results.

An *information structure* consists of measurable sets of signals T_i for each firm, and for each $\tilde{\mathcal{N}} \subseteq \mathcal{N}$, a joint probability measure $\pi_{\tilde{\mathcal{N}}}(t_{\tilde{\mathcal{N}}})$ on $\prod_{i \in \tilde{\mathcal{N}}} T_i$. The interpretation is that when firms in $\tilde{\mathcal{N}}$ are quoted, each quoted firm $i \in \tilde{\mathcal{N}}$ receives a *signal* t_i , and with the joint distribution of signals given by $\pi_{\tilde{\mathcal{N}}}$.

Given an information structure, firms choose prices conditional on their signals. We assume, without loss, that firms set prices between 0 and v . It is notationally convenient to describe price distributions in terms of upper cumulative distribution functions (sometimes referred to as decumulative distribution functions): Conditional on observing signal t_i , the likelihood that firm i sets a price p_i *greater* than x is given by

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

We let $F_i(dp_i|t_i)$ denote the measure over firm i 's price, and $F_{\tilde{\mathcal{N}}}(dp|t_{\tilde{\mathcal{N}}})$ denote the independent joint measure over prices of firms in $\tilde{\mathcal{N}}$ given their respective signals $t_{\tilde{\mathcal{N}}}$. When clear from the context, we write $F(dp|t)$ for $F_{\mathcal{N}}(dp|t_{\mathcal{N}})$. We shall use upper cumulative distribution functions for all (price) distributions throughout the paper.

The consumer will buy from one of the firms offering the lowest price, with ties broken

³In Section 6, we report how the analysis extends to a setting with general downward sloping demand for a homogenous good rather than unit demand.

⁴In Section 6, we report how our bound continues to hold, but may not be tight, if some firms are more likely to be quoted than others.

uniformly.⁵ Given a realized tuple of prices $p \in \mathbb{R}^{\tilde{\mathcal{N}}}$, let

$$L_{\tilde{\mathcal{N}}}(p) = \left\{ i \in \tilde{\mathcal{N}} \mid p_i = \min_{j \in \tilde{\mathcal{N}}} p_j \right\}$$

be the set of firms in $\tilde{\mathcal{N}}$ with the lowest price, conditional on firms in $\tilde{\mathcal{N}}$ being quoted at prices p . Then firm i 's revenue from the strategy profile $F = (F_1, \dots, F_N)$ is

$$R_i(F) = \left(\sum_{n=1}^N \mu(n) \frac{1}{\binom{N}{n}} \sum_{\{\tilde{\mathcal{N}} \subseteq \mathcal{N} \mid |\tilde{\mathcal{N}}|=n\}} \int_{T_{\tilde{\mathcal{N}}}} \int_{[0,v]^{\tilde{\mathcal{N}}}} p_i \frac{\mathbb{1}_{i \in L_{\tilde{\mathcal{N}}}(p)}}{|L_{\tilde{\mathcal{N}}}(p)|} F_{\tilde{\mathcal{N}}}(dp|t_{\tilde{\mathcal{N}}}) \pi_{\tilde{\mathcal{N}}}(dt_{\tilde{\mathcal{N}}}) \right),$$

The strategy profile F is an equilibrium if and only if

$$R_i(F) \geq R_i(F'_i, F_{-i})$$

for each i and strategy F'_i .

3 Two Firms

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 μ and two (competitive) quotes with probability $1 - \mu$. Thus the consumer gets a quote from firm 1 *only* with probability $\frac{1}{2}\mu$, firm 2 *only* with probability $\frac{1}{2}\mu$, and both firms with probability $1 - \mu$. In the continuum interpretation of the model, proportion $\frac{1}{2}\mu$ are “captive” consumers of firm 1, proportion $\frac{1}{2}\mu$ are captive of firm 2, and proportion μ of consumers are “contested,” and can choose where to purchase.⁶

3.1 Complete Information

First suppose that there is complete 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 sale price is 1 with probability μ and 0 with probability $1 - \mu$. The equilibrium (*ex ante*) sale price distribution is denoted by $S(x)$, i.e., $S(x)$ is the probability that the lowest price is *at least* x . Thus, $S(x)$ is also an

⁵The uniform tie breaking assumption is for simplicity of exposition, and in no way impacts our results. An asymmetric tie breaking rule will not impact the fundamental inequality (5) that drives our main result.

⁶This terminology in this interpretation follows Armstrong and Vickers (2019).

upper cumulative distribution function. The function $S(x)$ is depicted as the blue curve in Figure 2 for $\mu = \frac{1}{2}$: the probability that the price is at least 0 is 1, and the probability that the price is at least x for any $0 < x \leq 1$ is $1/2$. In the continuum interpretation, this corresponds to the case where the firm can see if a customer is captive or contested and price discriminate accordingly.

3.2 No Information

Now suppose that the firms have no information about the price count or, if they have information, are not able to condition on that information. In the continuum interpretation, this corresponds to the assumption that firms must offer a uniform price and cannot price discriminate. A firm asked to quote a price will therefore assign probability

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

to being the monopolist.

This model has a unique mixed-strategy equilibrium where firms randomize over prices in the interval

$$\left[\frac{\mu}{2 - \mu}, 1 \right].$$

Both firms use the same mixed strategy, wherein the probability of choosing price p_i or above is

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

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

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

Prices outside the support of F_j yield a strictly lower payoff, so that these strategies are an equilibrium. The resulting ex ante sale price distribution is

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

The no-information sale price distribution is the red curve in Figure 2, again for $\mu = \frac{1}{2}$. Because firms are always indifferent between the price p and charging the monopoly price $p = 1$, the expected price will equal the industry revenue (sum of the firms' profits) of μ .

Thus, even though the no-information and complete-information price distributions cannot be ranked in terms of first-order stochastic dominance, they do have the same mean.

In Section 6.2, we return to no-information case and the comparison with the complete information case. We will see more generally that the expected sale price is the same when there is no information as when there is complete information. Armstrong and Vickers (2019) therefore study an asymmetric model in order to obtain distinct expected prices. However, a key message of our paper is that these two extreme cases are not representative. In particular, both full information and no information minimize the expected sale price.

3.3 Partial Information

We now consider one class of parameterized partially informative information structures for the two firm case and solve for equilibria. It will turn out that this class contains an information structure whose equilibrium sale price distribution first-order stochastically dominates not only those arising under either full information or no information but also under any other information structure.

The information structures are as follows. Each quoted firm i receives a signal $t_i \in \{1, 2\}$. If the price count is 1, then the active firm receives a signal 1. If the price count is 2, then with probability $1 - 2\alpha$, both firms receive a signal 2, and with the complementary probability exactly one of the firms receives a signal 1 and the other firm receives a signal 2. Thus, the conditional probability distribution of signals only depends on the price count and are as follows (where a signal 0 indicates that the given firm was not quoted):

	$\pi(t n = 1)$			$\pi(t n = 2)$			
t_1/t_2	0	1	2	t_1/t_2	0	1	2
0	0	$\frac{1}{2}$	0	0	0	0	0
1	$\frac{1}{2}$	0	0	1	0	0	α
2	0	0	0	2	0	α	$1 - 2\alpha$

Figure 1: Signal distribution with two firms.

In effect, a signal 1 means that there is at least one quoted firm, and a signal of 2 means that there are two quoted firms. But when both firms are quoted, at least one firm is always told that both firms were quoted. The parameter $\alpha \in [0, 1/2]$ controls the dispersion in the beliefs of the market participants. If α is close to 0, then with high probability both of the firms believe that they are in a competitive environment. If α is close to $1/2$, then one of the firms believes that it is equally likely that it is in a monopolistic or a competitive environment.

In the continuum interpretation, this corresponds to an information structure where— if a consumer is in fact contested—with probability α , firm 1 is told that the consumer is contested; with probability α , firm 2 is told that the consumer is contested; and with probability $1 - 2\alpha$, neither is told if the consumer is contested. If the consumer is captive, neither is told.

We now describe an equilibrium where the firm which has received signal $t_i = 1$ charges deterministically a high price, $p_i = 1$, whereas the firm that receives the signal $t_i = 2$ will mix according to:

$$F_i(p_i) = \frac{\alpha}{1 - 2\alpha} \frac{1 - p_i}{p_i},$$

with support

$$\left[\frac{\alpha}{1 - \alpha}, 1 \right].$$

We will refer to the firm that receives the signal 2 as “informed,” as such a firm knows the number of equilibrium quotes. Conversely, a firm who receives the signal 1 is “uninformed,” as the firm is uncertain whether it is in a monopoly or in 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 1 and with probability $1 - 2\alpha$ the other firm observed 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 the firm with type $t_i = 2$ is indeed willing to randomize. Similarly, for the uninformed firm, it is either a monopolist with probability $\mu/2$, or it is in 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:

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

and in particular at the lower bound of the support at $p_i = \alpha/(1 - \alpha)$:

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

We can rewrite this inequality to isolate α :

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

We note for future reference that since $\sqrt{\mu(2-\mu)} < 1$, it must be that $\alpha^* < 1/2$, so that (1) 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 (1) holds, and this equation is satisfied for a nontrivial interval of α 's.

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

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

This formula is of course only correct when it is less than 1. For x sufficiently small, the expression blows up and $S(x) = 1$. As for the expected price, that is even easier to calculate: Firms are always indifferent to setting a price 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 must therefore be $\mu + (1-\mu)2\alpha$.

Notice that the sale price distribution is increasing in α for every x , i.e., in the first-order stochastic dominance order. The sale price distribution is therefore maximized at $\alpha = \alpha^*$, and the highest expected sale price is

$$\sqrt{\mu} \left(\sqrt{2-\mu} \right) \quad (2)$$

This is the upper bound on revenue that we referenced in the introduction. Indeed, Theorem 1 below will show that this information structure and equilibrium maximize both the expected sale price and the sale price distribution.

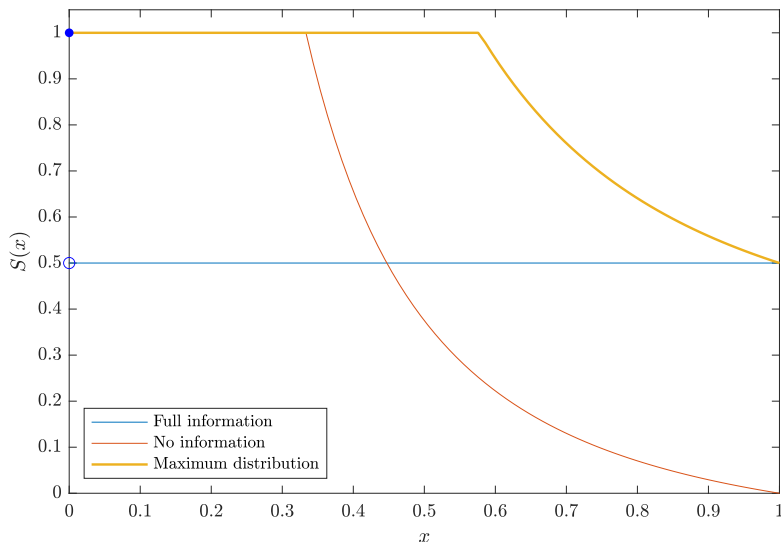


Figure 2: Price distributions when there are two firms.

Continuing our numerical example, when $\mu = \frac{1}{2}$, then

$$\alpha^* = \frac{1}{2} (\sqrt{3} - 1) \approx 0.366$$

and the support of the sale price is $[\alpha/(1 - \alpha), 1] \approx [0.577, 1]$. The resulting expected sale price is

$$\sqrt{\mu} (\sqrt{2 - \mu}) = \frac{1}{2} \sqrt{3} \approx 0.866.$$

The corresponding sale price distribution is the yellow curve in Figure 2.

Theorem 1 below establishes that a generalization of this construction delivers the highest possible sale price distribution for any number of firms and price count distribution. The analyses of this section and the next take the price count distribution as exogenous, as in Varian (1980). We will continue the example in Section 5, where we extend the analysis to endogenous price count distributions.

4 A Tight Upper Bound on the Sale Price Distribution

We now present our general results. Rather than studying equilibria directly, we will begin by studying the ex ante sale price distributions can arise in equilibrium. We will first establish an upper bound on the equilibrium sale price distribution, in the sense of first-order stochastic dominance. We then construct an information structure and equilibrium that exactly attain that bound. The bound immediately translates into an upper bound on

equilibrium revenue and a lower bound on consumer surplus.⁷ Finally, we show that the maximal sale price distribution increases as the price count distributions decreases, both in the first-order stochastic dominance order. 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.

4.1 A Constraint on the Sale Price Distribution

We first establish notation for the sale price distribution induced by an information structure (T, π) and strategy profile F . Let $S_i(x|n)$ 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 n firms quoting prices. This is

$$S_i(x|n) = \frac{1}{\binom{N}{n}} \sum_{\{\tilde{\mathcal{N}} \subseteq \mathcal{N} | i \in \tilde{\mathcal{N}}, |\tilde{\mathcal{N}}|=n\}} \int_{T_{\tilde{\mathcal{N}}}} \int_{p \in [x, v]^{\tilde{\mathcal{N}}}} \frac{\mathbb{I}_{i \in L_{\tilde{\mathcal{N}}}(p)}}{|L_{\tilde{\mathcal{N}}}(p)|} F_{\tilde{\mathcal{N}}}(dp|t_{\tilde{\mathcal{N}}}) \pi_{\tilde{\mathcal{N}}}(dt).$$

Also let

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

denote the *conditional sale price distribution*, given a price count of n . Finally, let

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

denote the *ex ante sale price distribution*.

We will argue that information contained in the sale price distributions $S_i(x|n)$ are rich enough objects to compute equilibrium revenue. In particular, firm i 's equilibrium revenue is

$$\sum_{n=1}^N \mu(n) \int_{x=0}^v x S_i(dx|n), \tag{3}$$

and total revenue is

$$\int_{x=0}^v x S(dx). \tag{4}$$

Our first result is the following integral inequality on the sale price distribution:

⁷We report a lower bound on revenue in Section 6.

Proposition 1 (Upper Bound on Sale Price Distribution).

In any equilibrium, the induced sale price distributions must satisfy, for all $x \in [0, v]$,

$$x \sum_{n=1}^N \mu(n) n S(x|n) \leq \int_{y=x}^{\infty} y S(dy). \quad (5)$$

The formal proof of this result is in the Appendix. To develop some intuition for the above inequality, consider the case where there is zero ex ante probability that x is the sale price, i.e., $S(\cdot)$ 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 deviation down to x* . We claim that the surplus resulting from such a deviation (where we average across which firm is the deviator) is precisely

$$\sum_{n=1}^N \mu(n) \left(\frac{n}{N} x S(x|n) + \frac{1}{N} \int_{y=0}^x y S(dy|n) \right). \quad (6)$$

The expression (6) can be understood as follows: Conditional on the price count being n , there is a n/N chance that the firm we picked to deviate is observed by the consumer. Conditional on being quoted, there is a probability $S(x|n)$ chance that the equilibrium sale 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 sale price would have been less than x , either (i) the deviating firm would have set a price less than x in which case they do not change their price, or (ii) another firm has the lowest price which is less than x . As a result, the deviation does not affect which firm has the lowest sale price. Thus, the deviating firm's surplus is simply what they would have received in equilibrium, i.e., there is a $1/n$ likelihood of having the lowest price, which is distributed according to $S(dy|n)$. If no firm wants to deviate in this manner, then it must be that the average surplus from this deviation across firms, given in (6), is less than the firms' average equilibrium revenue, which is $1/N$ of (4). Multiplying this inequality by N yields (5).

The inequality (5) is central to our subsequent analysis. For future reference, we can integrate (5) by parts and rearrange it into the following form

$$x \sum_{n=1}^N \mu(n) (n-1) S(x|n) \leq \int_{y=x}^v S(y) dy. \quad (7)$$

Here we have used the facts that S is an upper cumulative, so $S(dx) = -dS(x)/dx$ and

that $S(x) = 0$ for $x > v$.

4.2 The Maximal Price Distribution

To simplify terminology, we will say that an ex ante sale price distribution $S(\cdot)$ *deters uniform downward deviations* if there exist conditional sale price distributions $\{S(\cdot|n)\}$ that induce $S(\cdot)$ and satisfy (5). We now show that there is an ex ante sale price distribution $\bar{S}(\cdot)$ that deters uniform downward deviations and first-order stochastically dominates every other ex ante sale price distribution that deters uniform downward deviations. A fortiori, $\bar{S}(\cdot)$ also first-order stochastically dominates every equilibrium ex ante sale price distribution. The next section will construct a model of information and an equilibrium which generate $\bar{S}(\cdot)$, thus showing that it is a tight upper bound on the equilibrium sale price distribution.

The distribution $\bar{S}(\cdot)$ is induced by conditional sale price distributions $\{\bar{S}(\cdot|n)\}$ that have the following *ordered supports* property: There is a decreasing sequence of cutoff prices $v = x_0 = x_1 > x_2 > \dots > x_N$ such that the support of $\bar{S}(\cdot|n)$ is $[x_n, x_{n-1}]$. The cutoffs, and the functional form of $\bar{S}(\cdot|n)$, are chosen to satisfy (5) with equality at all $x \in [x_N, x_0]$. In particular, $\bar{S}(\cdot|1)$ puts probability one on v , and for $n > 1$,

$$\bar{S}(x|n) = \frac{\left(\frac{x_n}{x}\right)^{\frac{n}{n-1}} - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}}{1 - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}}. \quad (8)$$

Note that $\bar{S}(x_{n-1}|n) = 0$ and $\bar{S}(x_n|n) = 1$. Finally, the cutoffs x_n for $n > 1$ are

$$x_n = v \left(\prod_{m=1}^n \left(\frac{Q_{m-1}}{Q_m} \right)^{\frac{m-1}{m}} \right), \quad (9)$$

where

$$Q_m = \sum_{l=1}^m l \mu(l). \quad (10)$$

To visualize this construction, consider the case where $v = 1$ and there is a uniform distribution on the price count, so that $\mu(n) = 1/N$ for all n , i.e., the price count is uniformly distributed. We studied this example in the case where $N = 2$ in the previous section. Figure 3 plots $\bar{S}(x|n)$ in the case where $N = 5$. The resulting ex ante sale price distribution $\bar{S}(\cdot)$ places a mass point of size $\mu(1)$ on v , and for $x \in [x_n, x_{n-1}]$ for $n > 1$,

$$\bar{S}(x) = \mu(n) \bar{S}(x|n) + \sum_{m=1}^{n-1} \mu(m). \quad (11)$$

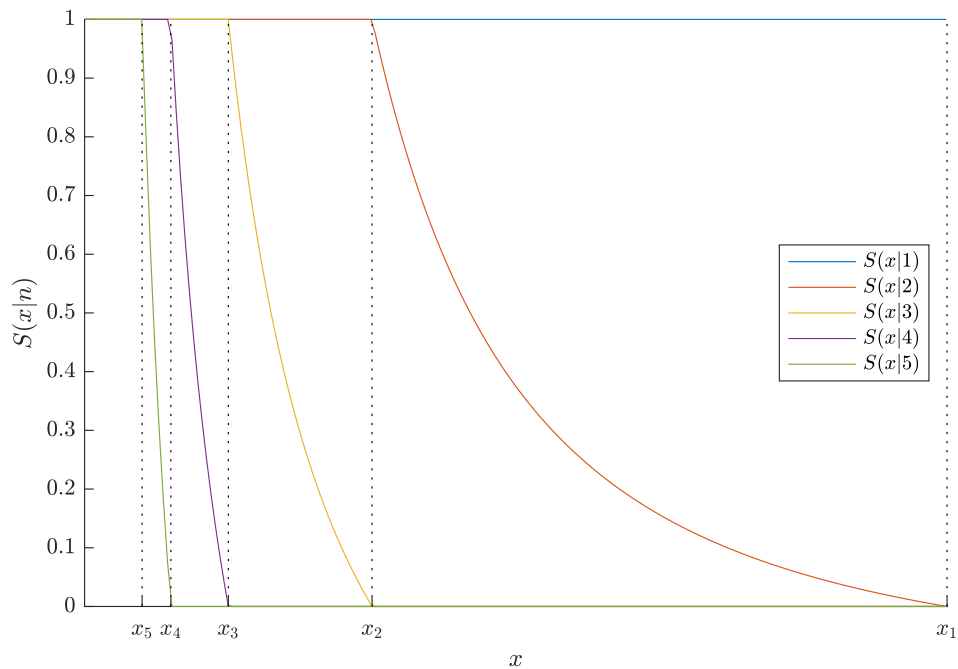


Figure 3: Conditional sale price distributions for $N = 5$ and uniformly distributed price counts.

In Figure 4 in Section 4.5, we plot the ex ante sale price distribution for various values of N including $N = 2$ and $N = 5$.

We now show the following:

Proposition 2 (\bar{S} Deters Uniform Downward Deviations).

The distributions $\{\bar{S}(\cdot|n)\}_{n=1}^N$ satisfy (5) for all $x \in [0, v]$. As a result, \bar{S} deters uniform downward deviations.

Proof of Proposition 2. We verify (5) using the reformulation in (7). Note that when $x = v$, both sides are equal to zero, so that the constraint is satisfied. Thus, it is sufficient to verify that both sides of (7) have the same derivative. When $x \in [x_n, x_{n-1}]$, the derivative of the

left-hand side is

$$\begin{aligned}
& \sum_{m=1}^n \mu(m)(m-1)\bar{S}(x|m) + x\mu(n)(n-1)\frac{d}{dx}\bar{S}(x|n) \\
&= \sum_{m=1}^{n-1} \mu(m)(m-1) + \mu(n)(n-1) \left[\frac{\left(\frac{x_n}{x}\right)^{\frac{n}{n-1}} - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}}{1 - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}} - \frac{n}{n-1} \frac{\left(\frac{x_n}{x}\right)^{\frac{n}{n-1}}}{1 - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}} \right] \\
&= Q_{n-1} - \sum_{m=1}^{n-1} \mu(m) + \mu(n) \left[(n-1) \frac{\left(\frac{x_n}{x}\right)^{\frac{n}{n-1}} - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}}{1 - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}} - n \frac{\left(\frac{x_n}{x}\right)^{\frac{n}{n-1}}}{1 - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}} \right] \\
&= Q_{n-1} - \sum_{m=1}^{n-1} \mu(m) - \mu(n) \left[\frac{\left(\frac{x_n}{x}\right)^{\frac{n}{n-1}} - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}}{1 - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}} + n \frac{\left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}}{1 - \left(\frac{x_n}{x_{n-1}}\right)^{\frac{n}{n-1}}} \right] \\
&= Q_{n-1} - \sum_{m=1}^{n-1} \mu(m) - \mu(n) \left[\bar{S}(x|n) + n \frac{Q_{n-1}}{Q_n - Q_{n-1}} \right] \\
&= Q_{n-1} - \sum_{m=1}^{n-1} \mu(m) - \mu(n) \left[\bar{S}(x|n) + n \frac{Q_{n-1}}{n\mu(n)} \right] = -\bar{S}(x),
\end{aligned}$$

which is precisely the derivative of the right-hand side of (7). In this derivation, we have used, in sequence, the definition of $\bar{S}(x|n)$, the definition of S_{n-1} , the definitions of x_n and x_{n-1} , and finally the definitions of Q_n and Q_{n-1} . Finally, the right-hand side has a derivative of zero when $x < x_N$, whereas the left-hand side is decreasing, so that the constraint is strictly satisfied for $x \in [0, x_N]$. \square

Next, given ex ante sale price distributions S and S' , we say that S *first-order stochastically dominates* S' if for all $x \in [0, v]$, $S(x) \geq S'(x)$. We now argue that not only does \bar{S} deter uniform downward deviations, but it also first-order stochastically dominates every other distribution satisfying the same:

Proposition 3 (First-Order Stochastic Dominance).

If the ex ante sale price distribution S deters uniform downward deviations, then \bar{S} first-order stochastically dominates S .

The proof of this result is in the Appendix. Here we sketch the three main steps. First, we argue that when maximizing the ex ante sale price distribution, it is without loss to consider distributions that have *ordered supports*, meaning that the supports of $S(\cdot|n)$ are intervals of the form $[y_n, y_{n-1}]$, where $\{y_n\}_{n=1}^N$ is an increasing sequence. In other words, the sale 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 n and x are negatively correlated, which leaves the right-hand side of (5) unchanged but decreases the left-hand side, thereby relaxing the constraint. Second, we argue that it is without loss to consider distributions for which (5) holds as an equality. If not, it is possible to increase push up the sale price distribution everywhere, while still satisfying (5). Third, we show that the ordered supports property, together with (5) as an equality, reduce to a first-order differential equation whose unique solution is the distribution \bar{S} .

4.3 Attaining the Bound

We now construct an information structure and equilibrium that attain the upper bound distribution of sale prices. Each firm receives signals in $\bar{T}_i = \{1, \dots, N\}$. First, for $n > 1$ and $n \geq k > 1$, we define

$$\alpha(k|n) = \frac{Q_{k-1}x_{k-1} \left(\left(\frac{x_{k-1}}{x_k} \right)^{\frac{1}{k-1}} - 1 \right)}{Q_{n-1}x_{n-1} \left(\frac{x_{n-1}}{x_n} \right)^{\frac{1}{n-1}}} \quad (12)$$

and

$$\alpha(1|n) = \frac{Q_1v}{Q_{n-1}x_{n-1} \left(\frac{x_{n-1}}{x_n} \right)^{\frac{1}{n-1}}}$$

and

$$\beta_n = 1 - (1 - \alpha(n|n))^n.$$

The signals are then generated according to the following distribution:

$$\bar{\pi}_{\tilde{\mathcal{N}}}(t_{\tilde{\mathcal{N}}}) = \begin{cases} \frac{1}{\beta_{|\tilde{\mathcal{N}}|}} \prod_i \alpha(t_i | |\tilde{\mathcal{N}}|) & \text{if } \left| \{i \in \tilde{\mathcal{N}} | t_i = |\tilde{\mathcal{N}}|\} \right| > 0; \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Thus, the signals are generated by taking independent draws from $\alpha(\cdot | |\tilde{\mathcal{N}}|)$ and throwing out realizations where all firms draw numbers less than $|\tilde{\mathcal{N}}|$, the realized price count. An interpretation of the signal is that it is a noisy observation of the price count. At least one firm observes the true price count, while the others observe numbers that are weakly lower. A key feature is that firms receive a signal that gives a hard lower bound on the price count. This is natural if we assume that the consumer collects prices sequentially and firms see a subset of the firms from whom the consumer has collected prices so far. However, it is important that at least one firm sees *all* of the firms that have been quoted.

Finally, firms use the pricing strategy

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

We are now ready to state the following result:

Proposition 4 (Maximal Information and Equilibrium).

The strategies \bar{F} defined by (14) are an equilibrium for the information structure $(\bar{T}, \bar{\pi})$, and these strategies induce \bar{S} as the equilibrium ex ante sale price distribution.

The complete proof is in the Appendix. We first verify that $\bar{\pi}$ as defined in (13) is in fact a probability distribution. We then show that this information structure and strategies induce the upper bound ex ante sale price distribution \bar{S} . This is essentially an application of the binomial theorem: When the price count is n , the number of firms that observe a signal of n is a truncated binomial, where at least one firm must observe n . We then compute the expectation of $(G(x|n))^k$ over the number of firms k that observe a signal n , which is exactly $\bar{S}(x|n)$.

Finally, we show that the strategies in (14) are an equilibrium. This is established by separately considering upward and downward deviations. For upward deviations, it is shown that a firm with signal $s_i = n$ strictly prefers to price in $[x_n, x_{n-1}]$ than any price greater than x_{n-1} . For downward deviations, we show that firms are *indifferent* between all prices in $[x_N, x_{n-1}]$. In fact, this is necessary in order to attain the bound on the price distribution: The constraint (7) says that for each x , firms on average do not want to uniformly deviate down to x . The critical distribution \bar{S} satisfies these constraints as equalities, meaning that firms are on average indifferent to uniform downward deviations. Of course, in equilibrium, firms cannot want to deviate in any manner, so it must be that all firms are indifferent to uniform downward deviations; otherwise, if some firm had a strict preference not to deviate, some other firm would have a strict preference to deviate. By a similar logic, if firms do not benefit on average by deviating down to x from all equilibrium prices above x , they must in fact be indifferent to deviating down from any given price above x . Thus, firms must be indifferent to *all* downward deviations. As we argue in the formal proof, this is precisely the case for the information structure and strategies we constructed.

4.4 Main Result

We now state our main result:

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 sale prices must be first-order stochastically dominated by \bar{S} given by (8)–(11). Moreover, there exists an information structure and equilibrium consistent with μ for which \bar{S} is the equilibrium sale 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 $\bar{R} = \int_{[0,v]} x \bar{S}(dx)$. 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 expectation of the sale price, and since \bar{S} is an equilibrium sale price distribution and first-order stochastically dominates every equilibrium sale price distribution, maximum expected producer surplus is the expected sale price under \bar{S} . Since the good is always sold, total surplus is always v , and hence minimum consumer surplus is $v - \bar{R}$. \square

4.5 Comparative Statics and an Elementary Bound on the Expected Sale Price

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 $\sum_{k=1}^n \mu(k) \geq \sum_{k=1}^n \mu'(k)$ for all n .

Proposition 5 (Price Count and Equilibrium Sale Price Distribution).

Let μ and μ' be price count distributions, with corresponding upper bounds \bar{S} and \bar{S}' . If μ' first-order stochastically dominates μ , then \bar{S} first-order stochastically dominates \bar{S}' .

Thus, as the price counts increases, the environment becomes more competitive, and the maximal sale price distribution shifts down. Holding fixed $\mu(1)$, the sale price distribution is therefore maximized when the remaining mass in the price count distribution is on $n = 2$. We therefore have an elementary bound on the expected sale price:

$$v \sqrt{\mu(1)(2 - \mu(1))}$$

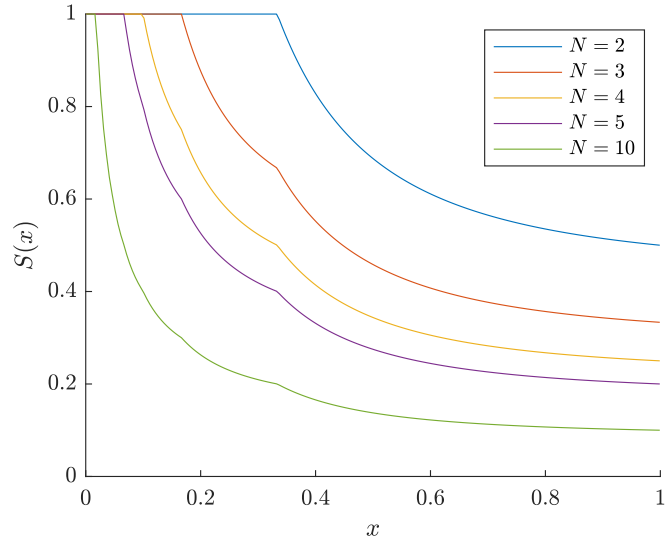


Figure 4: Sale price distributions for different N and uniform price counts. Kinks occur at the cutoffs x_n which are boundaries between the supports of conditional sale price distributions, as depicted in Figure 3.

We illustrate this with our uniform example with $v = 1$ and $\mu(n) = 1/N$ for all n . In Figure 4 we display the ex ante sale 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 sale price distribution decreases in the sense of first-order stochastic dominance.

Proof of Proposition 5. We will show the result for the case when μ is obtained from μ' by shifting a mass of ϵ from $n + 1$ to n , i.e.,

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

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 (5) for the price

count distribution μ' . Let us define

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

The induced ex ante distribution is precisely S' , so that the right-hand side of (5) is unchanged. But the left-hand side is now

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

Thus, the left-hand side has decreased, so that $\{S(\cdot|k)\}_{k=1}^N$ satisfy (5) for all x . Thus, any ex ante sale price distribution that deters uniform downward deviations 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}' . \square

Proposition 5 implies that holding fixed the probability of a single price count, $\mu(1)$, the upper bound on equilibrium sale price distribution is maximized when $\mu(2) = 1 - \mu(1)$, i.e., the price count is either one or two. In that case, the maximum expected sale price is

$$v \sqrt{\mu(1)(2 - \mu(1))}.$$

In contrast, under either complete information or no information (see Section 6.2 below), the expected sale price is $\mu(1)$. Figure 5 below contrasts the resulting expected revenue as we vary the probability $\mu(1)$. Thus, in the presence of incomplete information, maximum revenue is growing very quickly with 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 sale price converges to zero as $\mu(1)$ goes to zero. Thus, 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:

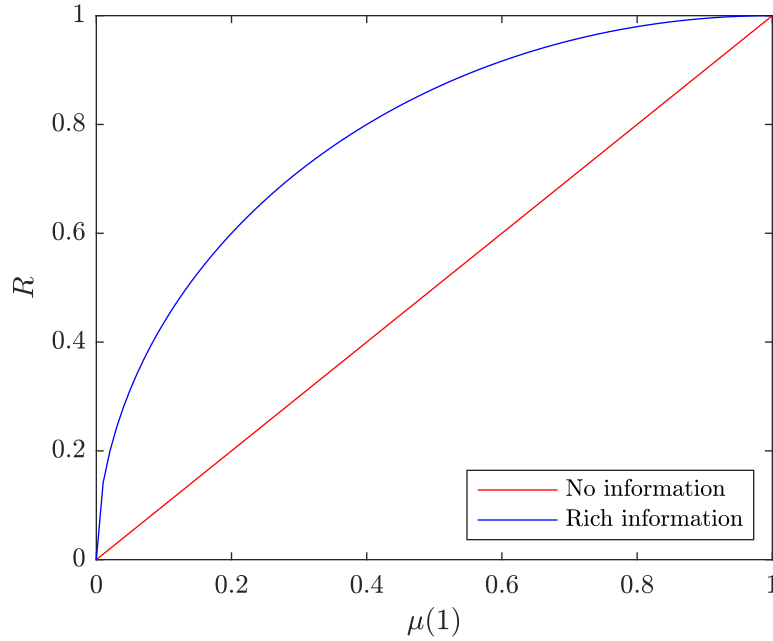


Figure 5: Maximum revenue for a given probability of monopoly.

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 sale price is $v\sqrt{\mu(1)(2-\mu(1))}$. Thus, marginal revenue with respect to the probability of being a monopolist is, $(1-v)/\sqrt{\mu(1)(2-\mu(1))}$, which is unbounded at $\mu(1) = 0$ where the market is fully competitive.

Conversely, holding fixed $\mu(1)$, Proposition 5 implies that \bar{S} is minimized when $\mu(N) = 1 - \mu(1)$. In this case, (7) becomes

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

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 complete information benchmark, and firms only obtain positive revenue when they are monopolists. The upshot is that in order to lift prices above the complete information level, it is insufficient for firms to have partial information about whether they are monopolists; it is also necessary that the price count be relatively small.

4.6 Connection to First-Price Auctions

The pricing game that we analyze here is strategically equivalent to a first-price auction where each bidder has either a low or a high value for a good, and each bidder knows his private value but is uncertain about the value 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 will depend on their beliefs about the number of other bidders with high values. The pricing game is a procurement auction where bidders quote prices at which they are willing to sell and auctioneer buys at lowest price. The strategies of quoted firms correspond high valuation bidders, while non-quoted firms are like low valuation bidders in the first price auction.

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. The two firm special case analyzed in Section 3 is equivalent to this case. Fang and Morris (2006) restricted attention to conditionally independent binary noisy signals about the opponent's value, and noted that that the expected price is necessarily higher with partial information than with either no information or complete information, a result which this paper generalizes. Azacis and Vida (2015) allowed many conditionally independent signals, and also considered the possibility of correlated signals and noted that the critical information structure we identify in Section 3 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 two value case arose in the critical information structure identified in Section 3. Bergemann, Brooks, and Morris (2013) also provides results on two bidder auctions with binary private values in asymmetric environments not reported in this paper.

Bergemann, Brooks, and Morris (2017) gives characterizes what can happen in general first-price auctions under general information structures. Lemma 1 of Bergemann, Brooks, and Morris (2017) established bounds on the equilibrium bids in the first-price auction, when each bidder is assumed to have no information about his value and those of the competing bidders. Proposition 1 in this paper establishes similar bounds when the bidder knows his own value. This is a different, and in general much harder, problem to solve as there are more constraints on bidders' possible higher-order beliefs. We are able to completely solve it in 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.⁸

5 Endogenizing the Price Count

We have thus far studied equilibrium sale price distributions holding fixed the price count distribution. As discussed in the introduction, there is a plethora of explanations of how the price count distribution is determined. An important question is whether our bounds still apply when we endogenize price counts, e.g., with a dynamic model of consumer search. In this section, we explore this issue in detail by considering various microfoundations of 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. Our bounds immediately apply to models with endogenous price counts and no feedback. There are many models, however, that exhibit *feedback*, meaning that a firm's realized price directly affects price counts and/or other firms' prices. Our bounds may or may not extend to models with feedback. We will, however, show that the bounds apply to models with sequential search, as long as prices are not observable to other firms and do not signal market conditions to consumers.

5.1 No Feedback

The price count could be a result of choices made either by consumers, firms, or even other agents who have influence over trade. For example, in Varian (1980), there is an exogenous distribution over the number of price quotations that a consumer observes, with some consumers (“shoppers”) observing all prices and some observing only one. This could be interpreted as arising from heterogeneous costs, with some consumers having positive costs of search and some having zero costs. Burdett and Judd (1983) consider a symmetric model where the consumer chooses a fixed number of quotes to observe and pays a constant cost per quote. Baye and Morgan (2001) consider a model where an intermediary manages a clearinghouse and charge firms and consumers for access. Firms and consumers can always opt out of the clearinghouse and trade in their respective “local” markets. Given the fees

⁸In 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 therein gives an implicit and incomplete characterization of maximal bid distributions.

set by the intermediary, there are mixed strategy equilibria where some firms and consumers pay the fees and trade in the clearinghouse. This endogenously gives rise to a price count distribution and asymmetric information about the price count.

These three models have the feature that price counts are determined before or simultaneous with the determination of prices. In addition, the firms' prices are all determined simultaneously. As a result, there is no feedback from 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). When there is no feedback, our analysis in the previous sections applies to whatever price count distribution is realized in equilibrium. The reason is that the critical uniform downward deviation 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, and the resulting surplus is still given by (6). As a result, the equilibrium sale price distribution must still satisfy the critical inequality (5), and Proposition 2 shows that the equilibrium distribution is bounded by \bar{S} .

To illustrate, one way to microfound the price count distribution from the example of Section 3 is with the simultaneous search model of Burdett and Judd (1983): Suppose there are two firms and a consumer can decide up front whether to observe one price for free or two prices at a cost, which is either zero or positive with equal probabilities. As long as the high cost is above a threshold (which turns out to be $(\ln 3 - 1)/2$), the high-cost consumers will choose to search a single firm. The rest of the equilibrium is as constructed in Section 3 and, in particular, the equilibrium sale price distribution is below our bound.

We will not need to provide a formal extension of our model with endogenous price counts and no feedback. In any such model, a price count distribution and firms' information about the price count will be part of the description of equilibrium, depending on variables like the consumer's cost of obtaining more price quotes. Theorem 1 bounds the sale price distribution as a function of the endogenous price count distribution, independent of the perhaps endogenous information structure and without reference to variables of the no feedback model like the cost of more quotes.

5.2 Feedback and Sequential Search

5.2.1 Preliminary Observations

When firms' prices can directly affect price counts or other firms' prices, our bounds may not hold. A trivial example of this is when there are two firms who price sequentially: If the following firm observes the price set by the leader, then there is an equilibrium in which the follower always has the lowest price, which is equal to the monopoly price. Thus, the

sale price distribution places probability one on the monopoly price even though there is common knowledge that there are two price quotes, which obviously violates our bounds.

Even when firms cannot condition on other firms' prices, there may still be feedback in which prices affect price counts. The leading example of this is sequential search. For example, in Stahl (1989), some consumers are "shoppers" and quote prices from all N firms, whereas other consumers are "searchers" and pay a cost $c > 0$ for each quote they receive. The critical feature that distinguishes this model from that of Burdett and Judd (1983) is that search is sequential: after seeing a price quote, the consumer can decide whether to buy at the lowest price found thus far, or to pay the cost and obtain another quote. In equilibrium, the searchers have a reservation price r , and they search as long as the lowest price found thus far is greater than r . In fact, firms never set prices above r in equilibrium, so that searchers only obtain a single price quote, and the resulting equilibrium price count distribution is the one in which shoppers receive N quotes and searchers receive a single quote. If a firm were to deviate to a price above r , however, they would induce searchers to obtain another quote, thus changing the price count from one to two.

Notice that in Stahl's model, firms can influence price counts, but only by deviating to higher prices. If firms deviate to lower prices, searchers will still not seek another quote. As a result, the uniform downward deviation results in exactly the same payoff as we computed in Section 4, and our bound still applies to the equilibrium sale price distribution.

This is a fairly general feature of models with sequential search: when firms reduce prices, this leads consumers to stop searching sooner. This in turn has the effect of reducing competition and making downward deviations more attractive. The net result is to tighten the incentive constraint associated with uniform downward deviations, so that our bounds still apply. In the remainder of this section, we will formalize this intuition in a class of sequential search models that generalize Stahl (1989, 1996), where the consumer makes a sequential choice whether to search.⁹

This class incorporates some important examples from the literature. Even so, there are scenarios in which lower prices might lead to more search. We have already discussed why our bounds would fail if firms learned directly about previous price quotes, as in a complete information Stackelberg game. Even if firms do not learn about previous price quotes, lower prices may still lead to higher price counts. For example, firms might have private

⁹While 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.

information about market conditions or other firms' beliefs. These factors may influence the prices set by as-yet unsearched firms. It is conceivable that firms will set low prices when they expect other firms to price low. Knowing this, a consumer may infer from low prices that searching more firms will result in lower prices and rationally respond by quoting more firms. These channels are excluded in the extension that we now pursue.

5.2.2 Sequential Search

We now formally extend our model by endogenizing price counts through consumer 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 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 Ξ denotes the set of permutations. The interpretation is that if the consumer searches at least n firms, then firm $\xi(n)$ provides the n 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 $\eta \in \Delta(\Theta)$. If a consumer searches n firms, then they pay a cost $c(n, \theta)$. We add the parameter θ to allow for heterogeneity in search costs among consumers, as in Stahl (1996). We make the simplifying assumption that for all θ, n , and n' , $c(n, \theta) \neq c(n', \theta)$.¹⁰ If the consumer purchases at price p after visiting n firms, the payoff is $v - p - c(n, \theta)$. The resulting payoff to the firm that makes the sale is p .

As before, we 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 $\pi(t|\theta, \xi)$. Note that firms may or may not know the order in which they were searched. We further assume that after searching n firms, the consumer sees the history $(\theta, \{(\xi(k), t_{\xi(k)}, p_{\xi(k)})\}_{k=1}^n)$. In other words, in addition to knowing their own type, the consumer also knows all of the identities, signals, and quoted prices of the firms that were searched.¹¹ The set of such histories of length n is denoted H_n , and the set of all histories is H . These sets are endowed with their natural product measurable structure.

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. With the

¹⁰Our 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 it does not increase the search cost.

¹¹This assumption means that consumers do not learn about firms' signals or the order from firms' realized prices. Such learning could lead to a violation of our bounds. See the discussion after Theorem 2.

complementary probability, the consumer buys from one of the firms with the lowest price quoted thus far, breaking ties at random. We further impose that for $h \in H_N$, $\sigma(h) = 0$, i.e., the consumer must buy after searching all of the firms.

Let us define

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

to be the probability that the consumer buys from the n th firm, 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

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

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

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

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

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

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

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

and summing up over all firms:

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

and all price counts:

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

We will analyze perfect Bayesian equilibria (Fudenberg and Tirole (1991)). That is, we will 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 will require that the consumer's strategy continues to be optimal even if firms' deviate in their prices. This is formalized as follows.

Fix a number $n \in \{0, \dots, N-1\}$ and a history $h \in H_n$. Given a history

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

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

$$\begin{aligned} \Xi(h) &= \{\xi' \in \Xi | \xi'(k) = \xi(k) \forall k = 1, \dots, n\}; \\ T(h) &= \{t' \in T | t'_{\xi(k)} = t_{\xi(k)} \forall k = 1, \dots, n\}; \\ P(h) &= \{p' \in [0, v]^N | p'_{\xi(k)} = p_{\xi(k)} \forall k = 1, \dots, n\}. \end{aligned}$$

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

$$\begin{aligned} U(F, \sigma, h) &= \frac{1}{|\Xi(h)|} \sum_{\xi \in \Xi(h)} \frac{1}{\sum_{t' \in T(h)} \pi(t'|\theta, \xi)} \\ &\times \sum_{t \in T(h)} \int_{p \in P(h)} U(\sigma, \theta, \xi, t, p) F_{\xi(k+1), \dots, \xi(N)}(dp_{\xi(k+1), \dots, \xi(N)}|t) \pi(t|\theta, \xi). \end{aligned}$$

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 $n = 0, \dots, N - 1$, $h \in H_n$, 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.)¹² 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.

To summarize, the parameters of the sequential search model are $\{\Theta, \eta, c, T, \pi\}$. We take as given the price count distribution μ . Under the hypothesis that μ is induced by an equilibrium (F, σ) of some sequential search model $\{\Theta, \eta, c, T, \pi\}$, we characterize bounds on the corresponding sale price distribution $S(\cdot)$.

5.2.3 Analysis

We now establish that our bounds apply to any equilibrium sale price distribution for the previously described sequential search model. At a history $h \in H_n$, let

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

denote the lowest price quoted thus far. We further define \tilde{H}_n as the set of histories of length n excluding prices, i.e., the set whose elements are of the form $(\theta, \{(\xi(k), t_{\xi(k)})\}_{k=1}^n)$. The union of the sets \tilde{H}_n across $n < N$ is denoted \tilde{H} . Our first result is the following:

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 $p(h) > r(\tilde{h})$.

The proof is in the Appendix. Similar reserve price characterizations are used in the sequential search literature. Our richer modeling of asymmetric information means we cannot directly apply existing arguments. We first argue that the consumer's problem is a standard dynamic programming problem, which has a value function $V(h)$ that satisfies a Bellman equation. Moreover, it turns out that V only depends on the non-price history \tilde{h} and the lowest price $p(h)$. The reason is that: (i) past prices are uninformative about the distribution of future prices, conditional on the non-price history \tilde{h} , and (ii) the value of purchasing at a previously quoted price only depends on $p(h)$. Thus, we may without loss denote the value function by $V(\tilde{h}, p)$. The last step of the argument is that $V(\tilde{h}, p)$ is convex and decreasing in p , with a slope greater than -1 . Combined with the fact that the search cost is

¹²This 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.

strictly increasing, this implies that the expected continuation value crosses the value from stopping at most once, and from below. The point where these payoffs cross is $r(\tilde{h})$.

Given that the consumer's equilibrium behavior is characterized by cutoffs $r(\tilde{h})$, we now argue that the constraint (5) must be satisfied by the equilibrium sale price distribution. The following is just a re-statement of the conclusion of Proposition 1 for the sequential search model.

Proposition 7 (Sequential Search and Sale 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 sale price distribution $S(\cdot)$ deters uniform downward deviations.

Note that 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.¹³ The Proposition is vacuous in this case.

In broad strokes, the argument for Proposition 7 is as follows. The uniform downward incentive constraint (5) says that firms should not want to uniformly deviate downward, under the premise that the price count distribution is μ (and all firms are equally likely to be quoted) both before and after the deviation. With sequential search, a downward deviation may cause the consumer to search differently. But Proposition 6 implies that the change is in the direction of stopping earlier, after encountering the lower price. On the whole, this makes the deviating firm better off than when price counts do not respond. In fact, we show that the deviating firm's surplus from a price cut is higher, even if we condition on (θ, ξ, t, p_{-i}) . This then implies the weaker conclusion that the ex ante payoff from a uniform downward deviation is higher than (6), so that (5) still holds.

Thus, if a sequential search model and equilibrium induce the price count distribution μ , the accompanying conditional sale price distributions $\{S_i(\cdot|n)\}$ must satisfy (5). Proposition 3 then implies that the induced ex ante distribution of sale prices must be first-order stochastically dominated by \bar{S} .

Finally, we observe that the critical information structure from Section 4.3 can be trivially adapted to sequential search, so that the strategies \bar{F} are an equilibrium and induce the ex ante price distribution \bar{S} . We simply set $\Theta = \{1, \dots, N\}$, $\eta(\theta) = \mu(n)$, and

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

¹³Janssen, Moraga-González, and Wildenbeest (2005) incorporate this possibility into the model of Stahl (1989).

(The purpose of the $+n$ term is merely to satisfy our genericity assumption that $c(n, \theta) \neq c(n', \theta)$ for all n, n' , and θ .) With this model, it is a strictly dominant strategy for the consumer of type θ to search θ firms. As a result, the price quote distribution is exactly μ , regardless of how firms set prices. In addition, the information is again given by $\bar{T}_i = \{1, \dots, N\}$ and

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

From Section 4.3, it is clear that the respective strategies of firms are still an equilibrium that induce the sale price distribution \bar{S} . This completes the proof of the following result:

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) that induce the price count distribution μ , the induced sale price distribution is first-order stochastically dominated by \bar{S} given by (8)–(11). Moreover, there exists a sequential search model and equilibrium that induce μ and the sale price distribution \bar{S} .

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

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

with support

$$\left[\frac{r}{3}, r\right], \text{ with } r = \frac{c}{1 - \frac{1}{2} \ln(3)}.$$

To verify this, observe that the expected price in the second search is $\frac{1}{2}r \ln(3)$, so the gain from searching is $r(1 - \frac{1}{2} \ln(3))$.

This equilibrium gives rise to the uniform price count distribution for the two firm case discussed in Section 3. In Figure 6, we add the equilibrium sale price distribution for the sequential search model for $c = \frac{1}{2}$ and so $r \approx 0.74$ to Figure 2. We can see how sequential search lowers the sale price distribution below that for simultaneous search giving rise to the same price count distribution. Thus, we conclude that our bounds will extend to a non-trivial class of models with sequential search. Importantly, we have assumed that firms do not observe one another’s prices. We have also assumed that firms have no information that is unobservable to the consumer, so that consumers do not learn from realized prices. Generalizing our bounds to allow for richer learning by firms and consumers is an important topic for future work.

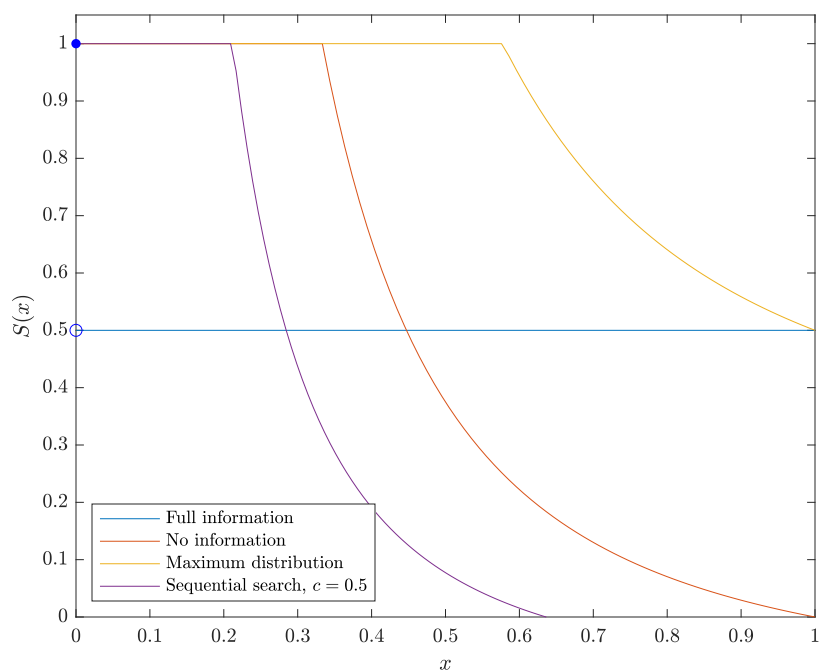


Figure 6: Equilibrium sale price distribution under sequential search.

6 Further Topics

We now consider several further topics and extensions of our results. We first characterize a tight lower bound on the expected sale price. We then analyze the benchmark where firms have no information about the price count. We then discuss how our results can be generalized to allow for downward sloping demand and asymmetric price count distributions.

6.1 Minimum Expected Price

We have shown that there is an equilibrium sale price distribution that first-order stochastically dominates any sale price distribution arising under any information structure. And we have noted that the expected price under this distribution is thus the highest possible equilibrium expected price.

In the next section, we will establish there is not an analogous lower bound on the distribution of prices, i.e., a sales price distribution which is first-order stochastically dominated by any other equilibrium sales price distribution. But we first characterize the lowest possible equilibrium price, and show that it is equal to the complete information expected price which is simply $\mu(1)v$.

Proposition 8 (Minimum Expected Price).

The minimum expected price across all information structures and equilibria is $\mu(1)v$. This is also minimum producer surplus. Maximum consumer surplus across all information structures and equilibria is $v - \mu(1)v$. Minimum producer surplus and maximum consumer surplus are attained under complete information.

Proof of Proposition 8. 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 equilibrium revenue is bounded below by $\mu(1)v/N$, and producer surplus is at least $\mu(1)v$. Clearly, this is producer surplus under complete 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. 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. \square

6.2 No Information

We now consider a variant of our model in which firms have no information beyond the ex ante distribution of the price count. There are at least three reasons for studying this model. First, it enables us to more directly compare our findings with previous results in the literature on consumer search, where the no-information assumption is standard. Second, as discussed in Section 3, the no-information case also corresponds to what would happen if firms had information about the price count but were not allowed to use it to price discriminate. Third, this allows us to establish the non-existence of a lower bound on the sale price distribution.

No information corresponds to the special information structure with $|T_i| = 1$ for all i . It follows that then

$$F_i(p_i|t_i) = F_i(p_i).$$

Under the further assumption that firms use symmetric strategies $F_i(p_i) = F(p_i)$, we can adapt arguments from Section 4 and establish a no-information bound on the sale price distribution. In this case,

$$S(x|n) = \prod_{i=1}^n (F(x))^n.$$

The inequality (5) then reduces to, for all $x \in [0, v]$,

$$x \sum_{n=1}^N \mu(n)(n-1)(F(x))^n \leq \sum_{n=1}^N \mu(n) \int_{y=x}^v (F(y))^n dy. \quad (15)$$

We now have the following result.

Proposition 9 (Equilibrium Sale Price Distribution with No Information).

There exists a highest symmetric price distribution \widehat{F} that satisfies the uniform downward incentive constraints under no information. This distribution satisfies $\widehat{F}(v) = 0$ and all of the constraints as equalities wherever $\widehat{F}(x) < 1$. Moreover, the distribution \widehat{F} is a symmetric Nash equilibrium.

Note that the bound on each firm's individual no-information equilibrium price distribution immediately implies a corresponding bound on the no-information equilibrium sale price distribution. While this result is markedly similar to Theorem 1, the proof is actually much simpler. Let \mathcal{F} be the set of all distributions that satisfy (15). The proof in the Appendix defines \widehat{F} to be the pointwise supremum of functions in \mathcal{F} . A straightforward argument shows \widehat{F} is also an upper cumulative distribution function that satisfies (15). We then use a similar argument as in Step 2 of Proposition 3 to argue that (15) must hold as an equality on the support of \widehat{F} . Finally, a direct calculation of each firm's surplus shows that when other firms use \widehat{F} , other firms are indifferent between all prices in the support of \widehat{F} , so that this distribution is a symmetric equilibrium.

Let us denote by \widehat{S} the resulting equilibrium sale price distribution when firms use the symmetric strategies \widehat{F} . We claim that the expected sale price is $\mu(1)v$, the lower bound from Proposition 8. The reason is that v is in the support of \widehat{F} , and there is no atom at v . As a result, each firm must be indifferent to setting a price of v , in which case they only make a sale when facing a captive consumer, and earn producer surplus $\mu(1)v/N$. This must be each firm's equilibrium surplus as well, so that the expected sale price is $\mu(1)v$.

In addition, since Proposition 9 shows that \widehat{S} is an upper bound on the no-information equilibrium sale price distribution, any other sale price distribution S that is first-order stochastically dominated by \widehat{S} would have to have an expected sale price that is strictly less than $\mu(1)v$. By Proposition 8, this cannot happen in equilibrium, so that in fact \widehat{S} is the *unique* symmetric equilibrium sale price distribution under no information:

Theorem 3 (Unique Sale Price Distribution).

When firms have no information about the price count, the unique symmetric equilibrium is \widehat{F} , the ex ante sale price distribution is \widehat{S} , and the expected sale price is $\mu(1)v$.

We see that the expected price is equal to the expected price over complete information, generalizing what we observed in the two firm example of Section 3. It is also the lower bound on the expected price, as shown in the previous section.

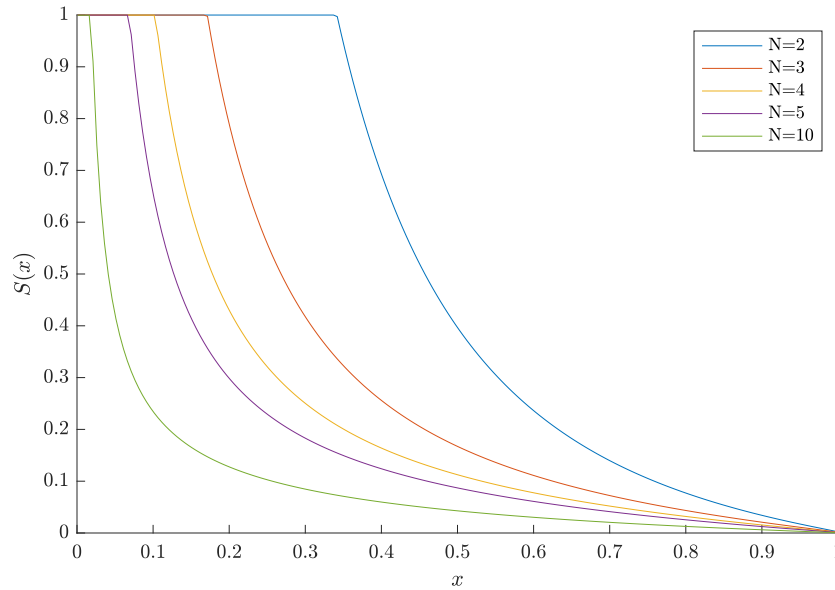


Figure 7: Equilibrium sale price distribution with no information.

We note that the no-information model will, in many cases, have other asymmetric equilibria.¹⁴ For example, when $N = 3$ and the price count is either $n = 1$ or $n = 3$, 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 in this equilibrium, all firms are indifferent to setting the monopoly price and only selling to captive consumers, so that the expected sale price is the same. We do not know of any asymmetric equilibria with different expected price.

In Figure 7 we illustrate how the symmetric equilibrium strategy changes as we increase the number of firms for the uniform example. The equilibrium distribution moves to the left as the expected number of price quotes increases. Figure 7 shows how the symmetric equilibrium with no information translates into the equilibrium sale price distribution, and Figure 8 compares it with our general upper bound in the $N = 5$ case. We see that general information structures support higher prices, in particular because the monopoly price 1 is charged with positive probability in the information structure supporting our upper bound on prices.

The combination of Propositions 8 and 9 have a number of important implications (under the maintained symmetric price count distribution).

First, a failure of common knowledge of the price count cannot explain an increased expected price, unless firms have partial information about the price count.

¹⁴Narasimhan (1988) identifies the unique equilibrium when there are only two firms.

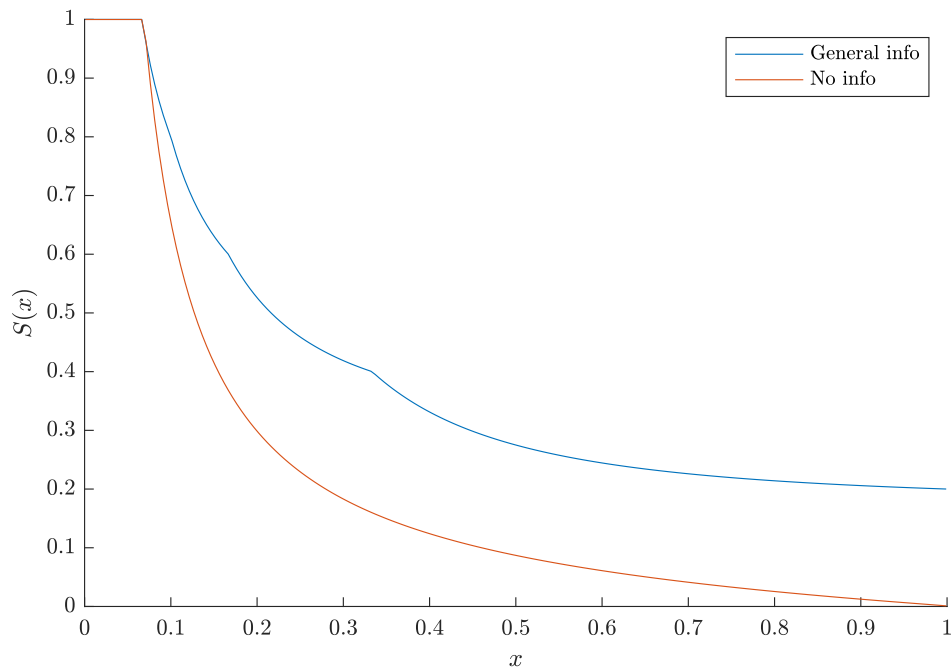


Figure 8: Comparison of sale price distributions under no information and general information.

Second, 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 the expected price is also higher under any information structure, including asymmetric information.¹⁵

Third, there is not a lower bound on equilibrium sales price distribution analogous to our main result, i.e., there is no sales price distribution which is first-order stochastically dominated by any other equilibrium sales price distribution. This can be seen by comparing the equilibrium sales price distributions under no information and complete information.

6.3 Beyond Single-Unit Demand

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

Suppose that the consumer has multi-unit demand with a downward sloping demand

¹⁵Proposition 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. Our analysis of the no-information case focuses on symmetric equilibria, but one can independently show that the expected surplus will be $\mu(1)v$ in any no-information equilibrium.

curve $D(p)$, maintaining the assumption that firms produce homogeneous goods and quote a single price for all units, so that the consumer only purchases from a low-price firm. In that case, the uniform downward incentive constraint (5) becomes:

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

We assume without loss that firms only use “weakly undominated” prices, i.e., the set

$$P^* = \{p \mid \nexists p' < p \text{ s.t. } p'D(p') \geq pD(p)\}.$$

Using a price p outside of P^* is weakly dominated in the sense 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 4 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 . Armstrong and Vickers (2019) analyze the more subtle implications for consumer surplus with downward sloping demand in the two firm case.

More generally, Theorem 1 can be adapted to other settings when firms compete by demanding a certain level of profit and a consumer accepts the contract offering the most surplus. For example, Lester et al. (2019) consider a setting where firms offer menus of price-quantity pairs to a consumer of unknown type, but the consumer’s preferences are “rank-preserving” in equilibrium, so the preferred contract is the same for all consumers. Lester et al. (2019) analyze what happens for a given price count distribution and no information, while Theorem 1 could be used to identify the highest distribution of profits across information structures.

6.4 Beyond a Symmetric Quote Distribution

We derived our results on the maximal sale price distribution under the assumption of single unit demand and a symmetric distribution over which firms are quoted. Both assumptions are easily relaxed at the cost of a lot of extra notation, with our upper bound on the sale price distribution continuing to hold. Here we discuss the straightforward extensions, noting some results which do rely on single unit demand and symmetry.

We assumed an ex ante distribution of the price count $\mu \in \Delta(\{1, \dots, N\})$ and that, conditional on the price count n , all firms are equally likely to be quoted. Suppose we assumed instead that there was an arbitrary distribution $\psi \in \Delta(2^{\{1, \dots, N\}}/\emptyset)$ over non-empty sets of firms quoted. Thus, firms are treated asymmetrically based on their labels. What would happen to our results?

First observe that ψ will induce an ex ante price count distribution

$$\mu(n) = \sum_{\{\tilde{\mathcal{N}} \subseteq \mathcal{N} \mid |\tilde{\mathcal{N}}|=n\}} \psi(\tilde{\mathcal{N}})$$

Now imagine that firms initially were not informed of their labels, and thought that they were equally likely to be assigned any label. Such firms would share the ex ante belief that the price count distribution was μ , and our upper bound on the sales distribution would apply to models in which firms have additional information about the price count or their labels.

However, if firms know their labels, then they are endowed with more information than the ex ante price count distribution, and we can no longer ensure that there is an information structure that attains the upper bound.

7 Conclusion

We have revisited the standard model of price dispersion, 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 micro-founding the price count, we simply take it as a primitive, and from it we derive a tight upper bound on the equilibrium distribution of sale prices. The bound holds across a rich family of models that micro-found 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 pricing, 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, as long as price cuts are unobservable to other firms and there is no feedback from prices to price counts, or under the assumption that price cuts would lead to lower price counts. An important direction for future work is to relax these assumptions further, by allowing for more complicated forms of feedback from prices to price counts and to allow for partial observability of prices by other firms.

References

- ARMSTRONG, M. AND J. VICKERS (2019): “Discriminating against Captive Consumers,” *American Economic Review: Insights*, 1, 257–272.
- AZACIS, H. AND P. VIDA (2015): “Collusive Communication Schemes in a First-Price Auction,” *Economic Theory*, 58, 125–160.
- BAYE, M. AND J. MORGAN (2001): “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets,” *American Economic Review*, 91, 454–474.
- BAYE, M., J. MORGAN, AND P. SCHOLTEN (2006): “Information, Search and Price Dispersion,” in *Economics and Information Systems*, ed. by T. Hendershott, Emerald Group Publishing, 323–376.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2013): “Extremal Information Structures in First Price Auctions,” Cowles Foundation for Research Discussion Paper 1926, Yale University.
- (2015): “First Price Auctions with General Information Structures: Implications for Bidding and Revenue,” Tech. Rep. 2018, Cowles Foundation for Research in Economics, Yale University.
- (2017): “First Price Auctions with General Information Structures: Implications for Bidding and Revenue,” *Econometrica*, 85, 107–143.
- BERGEMANN, D. AND S. MORRIS (2013): “Robust Predictions in Games with Incomplete Information,” *Econometrica*, 81, 1251–1308.
- (2016): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522.
- BURDETT, K. AND K. JUDD (1983): “Equilibrium Price Dispersion,” *Econometrica*, 51, 955–969.
- DE LOS SANTOS, A., A. HORTACSU, AND M. WILDENBEEST (2012): “Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior,” *American Economic Review*, 102, 2955–2980.
- DIAMOND, P. (1971): “A Model of Price Adjustment,” *Journal of Economic Theory*, 3, 156–168.

- DOVAL, L. AND J. ELY (2020): “Sequential Information Design,” Tech. rep., California Institute of Technology and Northwestern University.
- ELLISON, G. AND S. ELLISON (2009): “Search, Obfuscation and Price Elasticities on the Internet,” *Econometrica*, 77, 427–452.
- ELLISON, G. AND A. WOLITZKY (2012): “A search cost model of obfuscation,” *The RAND Journal of Economics*, 43, 417–441.
- FANG, H. AND S. MORRIS (2006): “Multidimensional Private Value Auctions,” *Journal of Economic Theory*, 126, 1–30.
- FUDENBERG, D. AND J. TIROLE (1991): *Game Theory*, Cambridge: MIT Press.
- HONG, H. AND M. SHUM (2006): “Using Price Distributions to Estimate Search Costs,” *RAND Journal of Economics*, 37, 257–275.
- JANSSEN, M. C., J. L. MORAGA-GONZÁLEZ, AND M. R. WILDENBEEST (2005): “Truly costly sequential search and oligopolistic pricing,” *International Journal of Industrial Organization*, 23, 451–466.
- LESTER, B., A. SHOURIKEH, V. VENKATESWARAN, AND A. ZETLIN-JONES (2019): “Screening and Adverse Selection in Frictional Markets,” *Journal of Political Economy*, 127, 0–0.
- MAKRIS, M. AND L. RENO (2017): “Information Design in Multi-Stage Games,” Tech. rep., University of Southampton and Queen Maru University.
- MORTENSEN, D. (2003): *Wage dispersion: why are similar workers paid differently?*, MIT Press.
- NARASIMHAN, C. (1988): “Competitive Promotional Strategies,” *Journal of Business*, 61, 4.
- ROGERSON, R., R. SHIMER, AND R. WRIGHT (2005): “Search-theoretic models of the labor market: A survey,” *Journal of economic literature*, 43, 959–988.
- ROTHSCHILD, M. (1973): “Models of Market Organization with Imperfect Information: A Survey,” *Journal of Political Economy*, 81, 1283–1308.
- SEGAL, I. AND M. WHINSTON (2003): “Robust Predictions for Bilateral Contracting with Externalities,” *Econometrica*, 71, 757–791.

- STAHL, D. (1989): "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review*, 79, 700–712.
- (1996): "Oligopolistic Pricing with Heterogeneous Consumer Search," *International Journal of Industrial Organization*, 14, 243–268.
- STIGLER, G. (1961): "The Economics of Information," *Journal of Political Economy*, 69, 213–225.
- SUTTON, J. (1991): *Sunk Costs and Market Structure*, MIT Press.
- VARIAN, H. (1980): "A Model of Sales," *American Economic Review*, 70, 651–659.

A Proofs

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 only affects the outcome when the lowest price would have been greater than $x - \epsilon_l$, and in particular, the deviator's surplus is

$$\sum_{n=1}^N \mu(n) \frac{1}{\binom{N}{n}} \sum_{\{\tilde{\mathcal{N}} \subseteq \mathcal{N} | i \in \tilde{\mathcal{N}}, |\tilde{\mathcal{N}}|=n\}} \int_T \int_{p \in \mathbb{R}^{\tilde{\mathcal{N}}}} \left((x - \epsilon_l) \mathbb{I}_{\min p > x - \epsilon_l} \mathbb{I}_{i \in \tilde{\mathcal{N}}} + \min p \mathbb{I}_{\min p \leq x - \epsilon_l} \frac{\mathbb{I}_{i \in L_{\tilde{\mathcal{N}}}(p)}}{|L_{\tilde{\mathcal{N}}}(p)|} \right) F_{\tilde{\mathcal{N}}}(dp|t_{\tilde{\mathcal{N}}}) \pi_{\tilde{\mathcal{N}}}$$

Note that this expression must be less than firm i 's equilibrium surplus, given in (3). As l goes to infinity, countable additivity of F implies that the deviator's surplus converges to

$$\sum_{n=1}^N \mu(n) \frac{1}{\binom{N}{n}} \sum_{\{\tilde{\mathcal{N}} \subseteq \mathcal{N} | i \in \tilde{\mathcal{N}}, |\tilde{\mathcal{N}}|=n\}} \int_T \int_{p \in \mathbb{R}^{\tilde{\mathcal{N}}}} \left(x \mathbb{I}_{\min p \geq x} \mathbb{I}_{i \in \tilde{\mathcal{N}}} + \mathbb{I}_{\min p < x} \frac{\mathbb{I}_{i \in L_{\tilde{\mathcal{N}}}(p)}}{|L_{\tilde{\mathcal{N}}}(p)|} \right) F_{\tilde{\mathcal{N}}}(dp|t_{\tilde{\mathcal{N}}}) \pi_{\tilde{\mathcal{N}}}(dt),$$

which is necessarily also less than (3). Summing the deviation surplus across i , we obtain

$$\begin{aligned} & \sum_{n=1}^N \mu(n) \frac{1}{\binom{N}{n}} \sum_{\{\tilde{\mathcal{N}} \subseteq \mathcal{N} | |\tilde{\mathcal{N}}|=n\}} \int_T \int_{p \in \mathbb{R}^{\tilde{\mathcal{N}}}} (nx \mathbb{I}_{\min p \geq x} + \mathbb{I}_{\min p < x}) F_{\tilde{\mathcal{N}}}(dp|t_{\tilde{\mathcal{N}}}) \pi_{\tilde{\mathcal{N}}}(dt) \\ &= \sum_{n=1}^N \left(nxS(x|n) + \int_{[0,x)} yS(dy|n) \right). \end{aligned}$$

This must be less than the sum of the firms equilibrium revenues, which is exactly the inequality (5). \square

Proof of Proposition 3. Fix $x \in [0, v]$. Consider the problem of maximizing $S(x)$ over all $\{S(\cdot|n)\}_{n=1}^N$ that satisfy (5), and where the functions $S(\cdot|n)$ are measurable functions which 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 $\{S(\cdot|n)\}_{n=1}^N$), (7) is closed, and the objective $S(x)$ is continuous, so that an optimal conditional distribution exist. We will show 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 $\{S(\cdot|n)\}_{n=1}^N$

that satisfy the following *ordered supports* property:

$$S(y|n) < 1 \implies S(y|n') = 0 \quad \forall n' > n. \quad (16)$$

Indeed, given any $\{S(\cdot|n)\}$ and associated ex ante distribution $S(\cdot)$, we can define a new $\{\tilde{S}(\cdot|n)\}$ with the same ex ante distribution, but where there is negative assortative matching between n and x . In particular, noting that $S(1) \leq \mu(1)$ from (7), for each n , we define \tilde{x}_n according to the minimum y such that $S(y) \geq \sum_{m=1}^n \mu(m)$. We then set $\tilde{S}(y|n) = (S(y) - S(\tilde{x}_{n-1})) / \mu(n)$. For each y , this correlation structure minimizes

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

and hence the left-hand side of (7). For future reference, note that for conditional distributions satisfying ordered supports, (7) is equivalent to for all $y < v$ and $n > 1$,

$$S(y|n) \leq \max \left\{ 0, \frac{1}{\mu(n)(n-1)} \left(\frac{1}{y} \int_{z=y}^v S(z) dz - \sum_{m=1}^{n-1} \mu(m)(m-1) \right) \right\}. \quad (17)$$

Step 2: Among distributions with the ordered supports property, it is obviously without loss to set $S(v|1) = \mu(1)$ (since $S(v|1)$ is unconstrained). Now, if a solution does not satisfy (17) as an equality when $S(y|n) < 1$, we can define a new solution, which is

$$\tilde{S}(y|n) = \max \left\{ 0, \min \left\{ 1, \frac{1}{\mu(n)(n-1)} \left(\frac{1}{y} \int_{z=y}^v S(z) dz + \sum_{m=1}^{n-1} \mu(m)(m-1) \right) \right\} \right\}, \quad (18)$$

which satisfies ordered supports and necessarily satisfies $\tilde{S}(y|n) \geq S(y|n)$ (strictly whenever (7) is strict), and therefore induces a higher ex ante distribution \tilde{S} . Thus, it is without loss to restrict attention to solutions for which (7) holds as an equality whenever $S(y|n) < 1$.

Step 3: We now show that the ordered supports property and (17) holding as an equality uniquely define the distributions $\{\bar{S}(\cdot|n)\}$. It is immediate that $S(y|n)$ will have a support that is an interval $[y_n, y_{n-1}]$, with $y_0 = y_1 = v$, and it is strictly increasing on its support. In addition, since the right-hand side of (17) 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 $m < n$. Then $S(y)$ must satisfy the boundary conditions $S(y_m) = \sum_{l=1}^m \mu(l)$ for all $m < n$. On

$[y_n, y_{n-1}]$, (17) holds as an equality, and moreover

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

As a result, (17) with equality rearranges to

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

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

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

Integrating both sides, we obtain

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

where C_n is a constant of integration. Thus,

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

where Q_{n-1} is defined above in (10). Differentiating both sides again, we obtain

$$S(y) = -\frac{C_n}{n-1} y^{-\frac{n}{n-1}} + \sum_{m=1}^{n-1} \mu(m) - \frac{1}{n} Q_{n-1}.$$

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

$$\begin{aligned} \sum_{m=1}^{n-1} \mu(m) &= -\frac{C_n}{n-1} (y_{n-1})^{-\frac{n}{n-1}} + \sum_{m=1}^{n-1} \mu(m) - \frac{1}{n} Q_{n-1} \\ \iff C_n &= -\frac{n-1}{n} y^{\frac{n}{n-1}} Q_{n-1}. \end{aligned}$$

As a result,

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

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

$$\begin{aligned} n\mu(n) &= \left(\frac{y_{n-1}}{y_n}\right)^{\frac{n}{n-1}} Q_{n-1} - Q_{n-1} \\ \iff Q_n &= \left(\frac{y_{n-1}}{y_n}\right)^{\frac{n}{n-1}} Q_{n-1} \\ \iff y_n &= y_{n-1} \left(\frac{Q_{n-1}}{Q_n}\right)^{\frac{n-1}{n}}. \end{aligned}$$

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

$$\begin{aligned} S(y|n) &= \frac{S(y) - S(y_{n-1})}{\mu(n)} \\ &= \frac{Q_{n-1} \left(\frac{y_{n-1}}{y}\right)^{\frac{n}{n-1}} - \left(\frac{y_{n-1}}{y_{n-1}}\right)^{\frac{n}{n-1}}}{n \mu(n)} \\ &= Q_n \frac{\left(\frac{y_n}{y}\right)^{\frac{n}{n-1}} - \left(\frac{y_n}{y_{n-1}}\right)^{\frac{n}{n-1}}}{n\mu(n)} \\ &= \frac{\left(\frac{y_n}{y}\right)^{\frac{n}{n-1}} - \left(\frac{y_n}{y_{n-1}}\right)^{\frac{n}{n-1}}}{1 - \frac{Q_{n-1}}{Q_n}}, \end{aligned}$$

which is precisely $\bar{S}(y|n)$. □

Proof of Proposition 4. We first verify that the information structure is well-defined, i.e., that $\bar{\pi}$ is a conditional probability distribution. Clearly $\alpha(k|n) \geq 0$. Also, using the formula for x_n , we can rewrite the numerator in (12) as

$$\begin{aligned} &Q_{k-1}x_{k-1} \left(\frac{x_{k-1}}{x_k}\right)^{\frac{1}{k-1}} - Q_{k-1}x_{k-1} \left(\frac{x_{k-1}}{x_{k-1}}\right)^{\frac{1}{k-2}} \\ &= Q_{k-1}x_{k-1} \left(\frac{x_{k-1}}{x_k}\right)^{\frac{1}{k-1}} - Q_{k-2}x_{k-2} \left(\frac{x_{k-2}}{x_{k-1}}\right)^{\frac{1}{k-2}}. \end{aligned}$$

The sum of these terms across k is precisely the denominator in the definition of $\alpha(k|n)$. 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 strategies induce \bar{S} . When the price

count is n , the probability that the highest price is at least x is

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

This follows from the binomial theorem: Note that

$$\alpha(n|n) = 1 - \left(\frac{x_n}{x_{n-1}} \right)^{\frac{1}{n-1}}$$

and so

$$\beta_n = 1 - \left(\frac{x_n}{x_{n-1}} \right)^{\frac{n}{n-1}}.$$

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

$$\begin{aligned} & \frac{1}{\beta_n} \left(\left(\frac{x_n}{x_{n-1}} \right)^{\frac{1}{n-1}} + \left(\left(\frac{x_n}{x} \right)^{\frac{1}{n-1}} - \left(\frac{x_n}{x_{n-1}} \right)^{\frac{1}{n-1}} \right) - \left(\frac{x_n}{x_{n-1}} \right)^{\frac{1}{n-1}} \right) \\ &= \frac{1}{\beta_n} \left(\left(\frac{x_n}{x} \right)^{\frac{1}{n-1}} - \left(\frac{x_n}{x_{n-1}} \right)^{\frac{1}{n-1}} \right), \end{aligned}$$

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 $n = k$; 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 $k = n$, 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

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

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

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

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

$$p_i \left(\widehat{G}(p_j|k) \right)^{k-1} \propto (p_i)^{1 - \frac{k-1}{k'-1}},$$

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

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

This expression deserves 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. Note that the likelihood that the price count is n , conditional on firm i being quoted, is proportional to $\mu(n) n$ (since a firm is more likely to be quoted when more firms are quoted). The first line gives the probability that firm i gets a signal k when there are $k' > k$ firms and the sale 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 their 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 n is between k' and k , firm i gets a signal of k ,

and at least one of the other firms gets a signal of n . 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, we have that

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

Let us denote by

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

Then

$$\begin{aligned} \frac{\alpha(k|n) \mu(n) n}{\beta_n} &= \frac{A(k)}{Q_{n-1} x_{n-1} \left(\frac{x_{n-1}}{x_n} \right)^{\frac{1}{n-1}} \left(1 - \left(\frac{x_{n-1}}{x_n} \right)^{\frac{n}{n-1}} \right)} \frac{Q_n - Q_{n-1}}{\left(1 - \left(\frac{x_{n-1}}{x_n} \right)^{\frac{n}{n-1}} \right)} \\ &= \frac{A(k)}{x_{n-1} \left(\frac{x_{n-1}}{x_n} \right)^{\frac{1}{n-1}} \left(1 - \left(\frac{x_{n-1}}{x_n} \right)^{\frac{n}{n-1}} \right)} \left(\frac{Q_n}{Q_{n-1}} - 1 \right) \\ &= \frac{A(k)}{x_{n-1} \left(\frac{x_{n-1}}{x_n} \right)^{\frac{1}{n-1}} \left(1 - \left(\frac{x_{n-1}}{x_n} \right)^{\frac{n}{n-1}} \right)} \left(\left(\frac{x_n}{x_{n-1}} \right)^{\frac{n}{n-1}} - 1 \right) \\ &= \frac{A(k)}{x_n}. \end{aligned}$$

Finally,

$$1 - (1 - \alpha(n|n))^{n-1} = 1 - \frac{x_n}{x_{n-1}}.$$

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

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

Thus, the payoff from a downward deviation 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. □

Proof of Proposition 6. 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:

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

We claim that in fact V only depends 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_N , this is obviously true, since $V(h) = v - c(N, \theta) - p(h)$. Now assume that the inductive hypothesis holds for $h \in \tilde{H}_k$ for $k > n$. Then at the history $h \in H_n$, the consumer can either stop and receive a payoff $v - p - c(n, \theta)$, or continue and receive a payoff of

$$\mathbb{E}_{(\xi(n+1), t_{\xi(n+1)}, p_{\xi(n+1)})} \left[V \left((\tilde{h}, \xi(n+1), t_{\xi(n+1)}, \min \{p(h), p_{\xi(n+1)}\}) \mid h \right) \right], \quad (19)$$

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

Second, it is obvious that $V(\tilde{h}, p)$ is decreasing in p (since the optimal strategy at a high p must generate a weakly higher payoff if p decreases). 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 only enters the consumer's payoff $U(F, \sigma, h)$ 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

$$\mathbb{E}_{(\xi(n+1), t_{\xi(n+1)}, p_{\xi(n+1)})} \left[V \left((\tilde{h}, \xi(n+1), t_{\xi(n+1)}, \min \{p, p_{\xi(n+1)}\}) \mid \tilde{h} \right) \right], \quad (20)$$

Each term in this expectation is clearly convex, since it is decreasing and convex in $p(h) = p$ when $p(h) < p_{\sigma(n+1)}$, 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 n .

Now we argue for the existence of reservation prices for $\tilde{h} \in \tilde{H}_n$ with $n < N$. If there is an $n' > n$ with $c(n', \theta) < c(n, \theta)$, 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 $c(n', \theta) > c(n, \theta)$ for all $n' > n$. Clearly, the payoff from stopping has a slope of -1 in the current lowest price and, since $c(n', \theta) > c(n, \theta)$ for $n' > n$, 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 (20) have crossed. But this can only happen if the slope of (20) is strictly greater than -1 . As the expected payoff from continuing is convex, we conclude that for $p' > p$, the slope of (20) is also strictly greater than -1 . As a result, (20) 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 by $r(\tilde{h})$, or if they never cross we let $r(\tilde{h})$ be any negative number. \square

Proof of Proposition 7. Let us fix a terminal history $h = (\theta, \xi, t, p)$. Suppose that firm i deviates to $p'_i < p_i$. The resulting payoff is

$$p'_i \sum_{n=\xi^{-1}(i)}^N \frac{\mathbb{I}_{i \in L((p'_i, p_{-i})_{\xi(1), \dots, \xi(n)})}}{|L((p'_i, p_{-i})_{\xi(1), \dots, \xi(n)})|} b_n(\sigma, \theta, \xi, t, (p'_i, p_{-i})). \quad (21)$$

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

$$p'_i \sum_{n=\xi^{-1}(i)}^N \frac{\mathbb{I}_{i \in L((p'_i, p_{-i})_{\xi(1), \dots, \xi(n)})}}{|L((p'_i, p_{-i})_{\xi(1), \dots, \xi(n)})|} b_n(\sigma, \theta, \xi, t, p). \quad (22)$$

(Note that we have dropped terms where the consumer stops searching before reaching firm i .) We claim that (21) is greater than (22). To see why, observe that by Proposition 6, $\sigma(h)$ is weakly increasing in the lowest price. Let

$$B_k(\sigma, \theta, \xi, t, p) = \sum_{n=k}^N b_n(\sigma, \theta, \xi, t, p)$$

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

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

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

$$\frac{\mathbb{I}_{i \in L\left((p'_i, p_{-i})_{\xi(1), \dots, \xi(n)}\right)}}{\left|L\left((p'_i, p_{-i})_{\xi(1), \dots, \xi(n)}\right)\right|}$$

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

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 doesn't respond (so that the price count distribution is the same). Since (6) was computed under the premise that price counts do not respond, it must be that the firm's surplus from a uniform downward deviation is weakly greater than (6). As a result, (5) must still be satisfied by the conditional price distributions that can be generated in equilibrium. \square

Proof of Theorem 3. Let \mathcal{F} be the set of F 's satisfying the uniform downward constraint. Then the pointwise supremum of the F 's, denoted \widehat{F} , is finite and also satisfies the constraints, since for all x ,

$$\begin{aligned} x \sum_{n=1}^N \mu(n)(n-1) (\widehat{F}(x))^n &= \sup_{F \in \mathcal{F}} x \sum_{n=1}^N \mu(n)(n-1) (F(x))^n \\ &\leq \sup_{F \in \mathcal{F}} \sum_{n=1}^N \mu(n) \int_{y=x}^v (F(y))^n dy \\ &= \sum_{n=1}^N \mu(n) \int_{y=x}^v (\widehat{F}(y))^n dy. \end{aligned}$$

Hence, $\widehat{F} \in \mathcal{F}$, and clearly it first-order stochastically dominates every other $F \in \mathcal{F}$.

Next, if \widehat{F} does not satisfy (15) as an equality for all x such that $\widehat{F}(x) < 1$, then we can

define $F(x)$ as the minimum of 1 and the solution z to

$$p \sum_{n=1}^N \mu(n)(n-1)z^n = \sum_{n=1}^N \mu(n) \int_{y=x}^v (\widehat{F}(y))^n dy.$$

Note that the right-hand side of this equation is increasing in x , and the left-hand side is strictly increasing in z . As a result, the new solution F is increasing in x and less than one. Moreover, we must have $F(p) \geq \widehat{F}(p)$, and strictly so when (15) is slack. As a result,

$$x \sum_{n=1}^N \mu(n)(n-1)(F(x))^n \leq \sum_{n=1}^N \mu(n) \int_{y=x}^v (\widehat{F}(y))^n dy \leq \sum_{n=1}^N \mu(n) \int_{y=x}^v (F(y))^n dy$$

so that $F \in \mathcal{F}$. This contradicts the hypothesis that \widehat{F} is the pointwise supremum of \mathcal{F} .

Finally, it remains to show that \widehat{F} is an equilibrium. Let $[\underline{x}, v]$ denote the support of \widehat{F} . Let $R(x)$ be the payoff from a price of x if other firms price according to \widehat{F} . Clearly, \widehat{F} has no atoms at \underline{x} and v , so that $R(x) = x$ for $x \leq \underline{x}$, and $R(x) = 0$ for $x \geq v$. Thus, to show that \widehat{F} is an equilibrium, it is sufficient to show that firms are indifferent between all x such that $\widehat{F}(x) < 1$. In this case,

$$R(x) = x \sum_{n=1}^N \frac{n}{N} \mu(n) (\widehat{F}(x))^{n-1}.$$

Note that since (15) binds on the support of \widehat{F} , it must be that \widehat{F} is differentiable on its support, with

$$\widehat{F}'(x) = \frac{\sum_{n=1}^N \mu(n) n (\widehat{F}(x))^n}{x \sum_{n=1}^N \mu(n) (n-1) n \widehat{F}(x)}.$$

Moreover, as we observed after Proposition 1, the integral form of the uniform downward constraint (7) is equivalent to (5), so that whenever $\widehat{F}(x) < 1$,

$$x \sum_{n=1}^N \mu(n) n (\widehat{F}(x))^n = \sum_{n=1}^N \mu(n) \int_{y=x}^v yn \widehat{F}^{n-1}(y) \widehat{F}'(y) dy.$$

Thus,

$$\begin{aligned} R(x) N \widehat{F}(x) &= \sum_{n=1}^N \mu(n) \int_{y=x}^v yn (\widehat{F}(y))^{n-1} \widehat{F}'(y) dy \\ &= \int_{y=x}^v R(y) \widehat{F}'(y) dy. \end{aligned}$$

As a result,

$$R'(x)N\widehat{F}(x) - NR(x)\widehat{F}'(x) = -R(x)\widehat{F}'(x),$$

so that $R'(x) = 0$. We conclude that firms are indifferent between all prices in the support of \widehat{F} , so that \widehat{F} is an equilibrium. \square