RAND Journal of Economics
Vol. 55, No. 4, Winter 2024
pp. 627–657

When is competition price-increasing? The impact of expected competition on prices

Sephorah Mangin*

We examine the effect of expected competition on markups in a random utility model where the number of competing firms may differ across consumers. Firms observe consumers' utility shocks and set prices using personalized pricing. We derive a precise condition under which the expected markup across consumers can be represented by a simple expression involving consumers' expected utility and the expected demand. This delivers a general condition under which greater expected competition is price-increasing. Whether this condition holds depends on the distribution of utility shocks, consumers' outside option, the expected number of competing firms, and the distribution of the number of firms competing for each consumer.

1. Introduction

■ What is the effect of competition on prices? Economists traditionally believe that greater competition tends to reduce prices, but theory suggests that this is not always the case. In early articles by Satterthwaite (1979) and Rosenthal (1980), for example, an increase in the number of firms can sometimes lead to an increase in the equilibrium price. More recently, Chen and Riordan (2008) shows that greater competition can be price-increasing – that is, the duopoly price can sometimes exceed the monopoly price – and Chen and Savage (2011) provides empirical support for this surprising theoretical prediction.

We study the impact of competition on prices in a random utility model where consumer choice is determined by firm-specific i.i.d. utility shocks (Anderson, de Palma, and Thisse, 1992). Gabaix, Laibson, Li, Li, Resnick, and de Vries (2016) considers a broad class of symmetric random utility models (e.g., Perloff and Salop, 1985; Sattinger, 1984; and Hart, 1985) and uses extreme value theory to examine the impact of competition on prices in large markets. The authors prove that, in the limit as the number of firms becomes infinite, the elasticity of markups

* Research School of Economics, Australian National University; sephorah.mangin@anu.edu.au.
I would like to thank the editor, Nicola Persico, and three anonymous referees for extremely valuable suggestions that have significantly improved the article. I thank Louis Becker for outstanding research assistance. I would particularly like to thank my discussants Hongyi Li and Simon Loertscher for valuable comments and suggestions, as well as George Mailath, Guillaume Roger, and Philipp Ushchev for useful discussions. I thank participants at the Melbourne IO and Theory Workshop, Melbourne, 2019; the Australasian Economic Theory Workshop, Sydney, 2019; the Econometric Society World Congress (Virtual), 2020; and the Asia-Pacific Industrial Organization Conference (Virtual), 2021. I am grateful to the Australian National University for providing a grant which funded Louis Becker's position as a research assistant.

converges to the *tail index* γ of the distribution of utility shocks (a measure of tail fatness). If this distribution has positive tail index $\gamma > 0$, competition is asymptotically price-increasing.[1] Although the asymptotic results presented in Gabaix et al. (2016) are insightful and elegant, they only apply to large markets.

This article examines the behavior of markups in an environment where the number of firms competing for any given consumer is finite but random. Different consumers may therefore have a different number of competing firms in their "choice set". In particular, we consider a *general distribution*, P_n, of the number of competing firms $n \in \mathbb{N}$ in a consumer's choice set.

We study the impact of expected competition on prices by studying how an increase in the expected number of firms affects the *expected markup*, that is, the average markup across all consumers who make a purchase. In contrast to Gabaix et al. (2016), we are interested in the behavior of the average markup across "local" markets that are potentially of *any* size – large, small, or even very small (including local monopoly or duopoly).

There are three different reasons why we model competition using this approach.

First, as discussed in Armstrong and Vickers (2022), there are many real-world environments in which the exact number of firms a consumer considers for their purchase may vary across consumers. For example, suppose a consumer searches for a product or service on the internet. There may be a large number of firms that could potentially compete for that particular consumer, but there are also various *frictions* that may influence the degree of realized competition for that consumer. These frictions may be regarding a firm's ability to serve the consumer, or the consumer's ability to find a suitable firm – which may be due to either limited search by the consumer, or limited advertising or "reach" by the firm. One potential result of these frictions is that the effective number of competing firms may vary across consumers. Given this, market analysts and regulators who are examining the overall effect of competition on markups may need to take this into consideration.

Second, this approach offers significant advantages compared to the same environment with a deterministic number of competing firms which is constant across consumers. First of all, our environment features much greater generality. We can nest the standard setting where the number of firms $n \geq 2$ is deterministic as a special case where P_n is degenerate. At the same time, our results are arguably simpler due to the fact that the expected number of firms is a continuous variable, which enables differentiability of key expressions.

Third, this approach opens up a range of new results that would not be possible if the number of firms was deterministic. In particular, we find that the behavior of markups depends not only on the distribution of utility shocks but also on the *distribution* P_n of the number of competing firms in a consumer's choice set.

In our environment, firms set prices *after* observing the realizations of both the number of competing firms and consumers' utility shocks (or valuations). That is, firms' price setting occurs without any uncertainty regarding either the number of competitors or consumers' preferences. Consumers choose whether or not to purchase a single unit of a good, and which good to purchase, after observing both prices and utility shocks. Following Rhodes and Zhou (2024), we call this *personalized pricing* (although it is referred to as *limit pricing* in Gabaix et al., 2016). When there are at least two firms, the equilibrium markup (i.e., price minus marginal cost) is equal to the difference between the highest and second-highest utility shock. When there is exactly one firm, the markup is equal to the difference between that firm's utility shock and the consumer's outside option.

Gabaix et al. (2016) shows that the equilibrium markups for all of the random utility models they consider are asymptotically proportional to the personalized pricing markup. For example, in the Perloff and Salop (1985) model, which is similar except that firms set prices *before* ob-

[1] If the distribution has negative tail index $\gamma < 0$, on the other hand, competition is asymptotically price-decreasing. If $\gamma = 0$, markups are relatively insensitive to changes in the degree of competition as the number of firms becomes large. See Gabaix et al. (2016) for details.

serving the realizations of consumers' utility shocks, the equilibrium markup is asymptotically proportional to the personalized pricing markup. Although these results only hold asymptotically, they suggests a common pricing logic underlying this class of models.

There are four main reasons why we choose to focus on personalized pricing.

First, the key results in this article can be applied to auctions. This is because the expected personalized pricing markup is identical to the winning bidder's expected surplus in a second-price auction with a random number of bidders. Our results apply more generally to any type of auction where the revenue equivalence theorem applies.

Second, personalized pricing is relevant not only for formal auctions. It is arguably better suited than the Perloff-Salop model to *any* environment where prices are individually tailored to each consumer. For example, it may be well-suited to environments involving haggling and negotiation rather than uniform retail prices. At the same time, it may be a good approximation of certain types of customized price-setting for online purchases when firms are able to acquire information about consumers' individual preferences prior to setting prices. Given that firms have access to an increasingly large amount of information about consumer preferences—for example, through data collection via social media—this type of personalized pricing is highly relevant today, as discussed in Rhodes and Zhou (2024).

Third, this type of pricing is also widely used in the literature on trade and macroeconomics. For example, Bernard, Eaton, Jensen, and Kortum (2003) incorporates a variant of this form of pricing (which is often called Bertrand competition) into a model of international trade with imperfect competition and heterogeneous firms; a large literature has followed.

Finally, personalized pricing proves to be highly tractable in our environment where the number of competing firms is random. We provide a general condition on the distribution P_n which enables us to derive a remarkably simple expression for the *expected markup* $\mu(\theta)$ as a function of the *expected number of firms*, θ. The simplicity of this expression is related to the fact that personalized pricing delivers efficient entry of firms in our setting.

The simple expression we obtain under personalized pricing – provided the general condition on P_n holds – relates the expected markup $\mu(\theta)$ to the consumer's *expected utility*, $M(\theta)$, and the *expected demand*, $D(\theta)$. The consumer's *expected utility* $M(\theta)$ is the expected utility a consumer receives from either purchasing a good or taking their outside option. The function $M(.)$ is not a standard utility function but instead represents the consumer's expected utility as a function of the expected number of firms θ. The function $M(.)$ also depends on the distribution of utility shocks and the value of the consumer's outside option. The *expected demand* $D(\theta)$ is the probability that a given firm successfully sells their good.

Our expression for the expected markup is remarkably simple: $\mu(\theta) = M'(\theta)/D(\theta)$. That is, the expected markup $\mu(\theta)$ is equal to the marginal contribution $M'(\theta)$ of an additional firm to the consumer's expected utility, divided by the expected demand $D(\theta)$. An analogous difference equation holds when the number of firms $n \geq 2$ is deterministic, but there are some crucial differences in our setting. First, $\theta \in \mathbb{R}_+$ is the *expected* number of firms, which is continuous (not discrete) and $M(.)$ is differentiable. Second, $M(\theta)$ incorporates the consumer's outside option. Third, we allow the possibility that the number of competing firms is zero, one, two, or more. Finally, in our environment $M(\theta)$ depends not only on the distribution of utility shocks, but also on the distribution of the number of competing firms.

This simple expression for the expected markup is related to our result that personalized pricing delivers the *efficient* level of firm entry (that is, the level of entry that maximizes the social surplus minus entry costs) when the distribution P_n satisfies our general condition. When this condition holds, personalized pricing ensures that firms' expected profits equal their marginal contribution to the social surplus, that is, the difference between the highest and the second-highest utility shock (or the consumer's outside option, if there is only one firm). When this condition fails, this is not always true. This highlights the fact that efficiency of entry depends not only on the type of pricing but also on the distribution of the number of competing firms—a result which only becomes apparent when the number of firms is stochastic.

For any distribution P_n that satisfies our general condition, the connection between the expected markup and expected utility delivers a simple condition under which competition is either *price-increasing* ($\mu'(\theta) > 0$) or *price-decreasing* ($\mu'(\theta) < 0$). In particular, whether or not competition is price-increasing depends on the local curvature of the expected utility function $M(.)$ at θ. This measure of local curvature is $r_M(\theta) = -M''(\theta)\theta/M'(\theta)$, which is the elasticity of $M'(.)$ at θ. Competition is price-increasing if and only if the local curvature $r_M(\theta)$ is strictly less than $\varepsilon_D(\theta)$, the elasticity of expected demand with respect to θ.

Intuitively, as the number of firms rises, the marginal increase $M'(\theta)$ in the consumer's expected utility decreases because $M(.)$ is concave and $M''(\theta) < 0$. However if the rate of decrease in $M'(\theta)$ is sufficiently low, that is, if $M(.)$ is not *too* concave relative to the demand elasticity, then competition is price-increasing and $\mu'(\theta) > 0$. Importantly, this is a local condition. Whether or not it holds depends not only on the properties of the distribution of utility shocks and the value of the consumer's outside option, but also on the expected number of firms θ. In addition, it depends on the distribution P_n of the number of competing firms in a consumer's choice set.

Outline. Section 2 discusses the related literature. Section 3 presents the model. Section 4 derives some preliminary results. Section 5 presents our lead example, the Poisson distribution. Sections 6 and 7 contain our main results. Section 8 discusses the application of our results to auctions. Section 9 contains our asymptotic results. Section 10 presents some examples. The Appendix contains all proofs not found in the main text.

2. Related literature

■ This article builds on an existing literature that considers the possibility of price-increasing competition, starting with the classic early articles of Satterthwaite (1979) and Rosenthal (1980), both of which describe environments in which an increase in the number of firms can lead to an increase in prices.[2] More recently, Chen and Riordan (2008) shows that, in an environment featuring perfect information and pure strategies, the symmetric duopoly price is higher than the single-product monopoly price when consumers' utility shocks are independent and the distribution has a decreasing hazard rate. Although our approach is different, our results are complementary to those found in Chen and Riordan (2008).

This article is also complementary to Chen and Riordan (2007). The authors prove that an increase in the number of firms can lead to higher equilibrium prices in the spokes model of nonlocalized spatial competition. Part of the reason why competition may be price-increasing in Chen and Riordan (2007) is that firms can sell to consumers in different submarkets—some of which are (effectively) duopolistic and some of which are (effectively) monopolistic. A higher number of firms increases the proportion of submarkets that are duopolistic, which can affect the overall demand elasticity. As a result, the equilibrium price can be higher under certain conditions. In our environment, firms sell in a single market where they may be either one firm, two firms, or more than two firms competing for the consumer. In our setting, however, this is because the number of competing firms is *random*.

The distribution P_n of the number of competing firms, which is central in our article, is related to similar distributions that appear in some recent articles: the distribution of the price count (i.e., the number of firms from which a consumer obtains a quote) in Bergemann, Brooks, and Morris (2021), and the distribution of the number of firms in consumers' consideration sets (i.e., the set of firms a consumer considers for their purchase) in Armstrong and Vickers (2022). In both articles, goods are homogeneous and consumers purchase from the firm offering the lowest price, whereas in our environment consumers receive random utility shocks, firms set

[2] In Satterthwaite (1979), the environment features imperfect consumer information, whereas in Rosenthal (1989) the result is obtained using mixed-strategy pricing.

prices after observing both the shocks and the number of competitors, and consumers purchase from the firm that maximizes their net utility.

As discussed, this article is closely related to Gabaix et al. (2016), although our motivation and focus are different. Our approach is complementary because in our environment the number of competing firms is *finite* but random. Whereas Gabaix et al. (2016) shows that the tail index γ of the distribution of utility shocks is key to understanding the impact of competition on prices, we prove that it is the local curvature of the consumer's expected utility, as measured by $r_M(\theta) = -M''(\theta)\theta/M'(\theta)$, that is key in our setting. Importantly, this depends not only on the distribution of utility shocks, but also on the value of the consumer's outside option, the expected number of firms, and the distribution of the number of firms.

This article is also related to Weyl and Fabinger (2013) and Quint (2014). Both articles focus on different questions, but point out that the comparative statics of pricing behavior depends on log-concavity of the demand function. Weyl and Fabinger (2013) and Quint (2014) examine different environments, but both articles highlight the fact that greater competition decreases markups if the density of the distribution of utility shocks is log-concave, but increases markups if this density is log-convex.[3] Our results are different for three main reasons. First, log-concavity or log-convexity is a *global* criterion: markups are either increasing or decreasing for *all* $n \geq 2$. By contrast, our criterion for markups to be increasing (decreasing) is *local* because it depends on the expected number of firms θ (in addition to the distribution of utility shocks). Second, our criterion also depends on the distribution P_n of the number of competing firms. Third, we do not restrict attention to cases where there are two or more competing firms. This means that local monopoly (i.e., only one competing firm) is a possible outcome. As a result, the consumer's outside option is also important.

For example, consider a Pareto distribution of utility shocks (which has a log-convex density). The deterministic expected markup is strictly increasing in the number of firms. However, in general, greater competition is *not* always price-increasing. For both our Poisson and geometric examples, the expected markup varies *nonmonotonically* with the expected number of firms when the consumer's outside option is below the minimum utility shock. At first, when the expected number of firms is relatively low, competition is price-decreasing. Later, when the expected number of firms is sufficiently high, competition is price-increasing. Importantly, we find that the distribution of the number of competing firms can affect both the number of firms at which competition switches from being price-decreasing to price-increasing and the rate of convergence of the markup elasticity to its asymptotic value.

Although our results apply to more general distributions of the number of competing firms, our lead example of the Poisson distribution is related to Platt (2017), which simplifies and extends limit results in an auction environment to finitely many bidders by assuming the number of bidders is Poisson distributed. This assumption is common in both the theoretical literature on large auction markets, such as Satterthwaite and Shneyerov (2007), and the empirical literature involving estimation of auctions with a stochastic number of bidders, such as Coey, Larsen, and Platt (2020). Our article is also related to the empirical search cost literature, including Allen, Clark, and Houde (2019) and Salz (2022), in which the number of competing firms varies across consumers due to the presence of search frictions.[4]

Some of the techniques used in our article are also closely related to the competing auctions literature. In competing auctions, a large number of sellers compete to attract buyers by posting auctions with reserve prices (Peters and Severinov, 1997).[5] Our result that personalized

[3] In an earlier article, Anderson, De Palma, and Nesterov (1995) consider the Perloff-Salop model and show that a sufficient condition for markups to be weakly decreasing in the number of firms is log-concavity.

[4] In Allen et al. (2019), which studies the Canadian mortgage market, searching consumers run an auction between their home firm and rival firms. In Salz (2022), which studies the trade waste industry in New York, the firm's problem in the search market is effectively a first-price auction with unknown number of bidders.

[5] More recent articles using competing auctions include Albrecht, Gautier, and Vroman (2012, 2014), Kim and Kircher (2015), Lester, Visschers, and Wolthoff (2015) and Mangin (2017).

pricing delivers an *efficient* level of firm entry mirrors the well-known result that buyer entry is constrained efficient in competing auctions environments where sellers' reserve prices are equal to their own valuations (Peters and Severinov, 1997; Albrecht et al., 2012). Although most of the competing auctions literature focuses on the Poisson distribution, Eeckhout and Kircher (2010) and Lester et al. (2015) consider more general meeting technologies.

3. Model

■ Consider a single product market with a single consumer. The number of competing firms $n \in \mathbb{N} = \{0, 1, 2, \ldots\}$ is random. The expected number of competing firms is $\mathbb{E}(n) = \theta \in \Theta$ where $\Theta \subseteq (0, \infty)$. For now, we take the expected number of firms θ to be exogenous in order to focus our attention on the effect on markups of varying θ.[6]

For any $n \in \mathbb{N}$, let $P_n(\theta)$ denote the probability that there are n firms competing for a consumer. We have $P_n : \Theta \to [0, 1]$ and $\sum P_n(\theta) = 1$, and we assume $\sum nP_n(\theta) = \theta$ and $\sum n^2 P_n(\theta) < \infty$, which implies finite variance. Although we focus on a single consumer for our exposition, the distribution P_n can also be interpreted as the distribution *across consumers* when there is a continuum of consumers. Under this interpretation, $P_n(\theta)$ is the proportion of consumers facing n competing firms.

It is sometimes useful to consider the distribution P_n from the perspective of firms. For any $n \in \mathbb{N}\backslash\{0\}$, let $Q_n(\theta)$ denote the probability that a firm faces $n - 1$ competitors. We have $Q_n : \Theta \to [0, 1]$ and $\sum_1 Q_n(\theta) = 1$.[7] The distribution Q_n is implied by the distribution P_n, which is exogenous. A useful identity, which allows us to determine Q_n from P_n, is

$$\theta Q_n(\theta) = nP_n(\theta). \tag{1}$$

This identity must hold for all $n \in \mathbb{N}$ to ensure the distributions P_n and Q_n are consistent. For discussion of this identity, see Eeckhout and Kircher (2010) and Lester, Visschers, and Wolthoff (2015).[8]

Firms are *ex ante* identical. Each firm can produce one unit of the good at marginal cost $c \geq 0$. After the number of competing firms n is realized, the consumer draws an i.i.d. utility shock x_i for each firm $i \in \{1, 2, \ldots, n\}$ from an exogenous distribution with cdf G.

The following assumption is maintained throughout the article.

Condition 1. The distribution of utility shocks has a continuous, twice-differentiable cdf G with pdf $g = G' > 0$, support $[x_0, \bar{x}) \subseteq \mathbb{R}_+$ where $\bar{x} \in \mathbb{R}_+ \cup \{+\infty\}$, and a finite mean.

To keep the environment as general as possible, we allow the possibility that the consumer can obtain the good elsewhere at cost c and receive utility $z \in [0, x_0]$. In this case, the consumer receives net utility $z - c$. We refer to this possibility as the consumer's *outside option*. We assume the consumer always chooses to purchase the good when indifferent.

Firms set prices simultaneously, *after* observing both the consumer's firm-specific utility shocks x_i for each firm $i \in \{1, 2, \ldots, n\}$ and the number of competing firms. After observing prices, the consumer decides whether to purchase one unit of the good and, if he chooses to purchase, selects the firm i that maximizes his net utility, $x_i - p_i$. Profits for the successful firm i are given by the markup, which is defined as $\mu_i \equiv p_i - c$.

---

[6] In Section 6, we endogenize the expected number of firms θ.

[7] For notational simplicity, we let $\sum a_n$ denote the summation $\sum_{n=0}^{\infty} a_n$, and we let $\sum_k a_n$ denote the summation $\sum_{n=k}^{\infty} a_n$ (whenever there is no possibility of confusion).

[8] Consider a continuum of consumers of measure L and a continuum of firms of measure V. Suppose the expected number of firms competing for a consumer is $\theta = V/L$. The measure of firms with $n - 1$ competitors, $VQ_n(\theta)$, equals n times the measure of consumers facing n competing firms, that is, $nLP_n(\theta)$. This implies (1).

Timing of events:

1. Number of competing firms n is realized
2. Random utility shocks x_i are realized
3. Firms observe shocks
4. Firms set prices simultaneously
5. Consumer observes prices
6. Consumer makes purchase decision
7. Production takes place
8. Firm profits are realized

4. Equilibrium

■ In this section, we present the equilibrium expected markup, the expected demand for a single firm's product, and the consumer's expected utility.

Equilibrium markup. Suppose that the realizations of both the number of competing firms n and the consumer's utility shocks are known by all agents. Define $M_n \equiv \max\{x_1, \ldots, x_n\}$, the highest utility shock, and let S_n denote the second-highest utility shock.

Markup. When there are $n \geq 2$ firms, the one with the highest shock sets a price equal to $p = M_n - S_n + c$, which gives the consumer net utility $M_n - p = S_n - c$. That is, the firm sets a price which is just low enough to keep the second-best firm out of competition for the consumer. In equilibrium, the consumer chooses to purchase the good (because $S_n \geq x_0 \geq z$), and he purchases from the firm with the highest utility shock. The equilibrium markup is equal to the difference between the highest and second-highest shocks, $\mu = M_n - S_n$.

When there is exactly one firm, the firm sets a price $p = x_1 - z + c$, which gives the consumer net utility $x_1 - p = z - c$. That is, the firm sets a price that ensures the consumer is indifferent between purchasing the good and their outside option, which gives net utility $z - c$. In equilibrium, the consumer purchases from the firm and the markup is $\mu = x_1 - z$.

Without loss of generality, we normalize $c = 0$ throughout the remainder of the article.

Expected markup μ_n. We first calculate the expected markup μ_n given that there are n competing firms. Importantly, our setting allows for the possibility that there may be either no firms, one firm, two firms, or more than two firms competing.

When there is exactly one firm competing for a consumer, we call this a "local monopoly". In this case, the the consumer's outside option z is important and the expected markup is

$$\mu_1 = \mathbb{E}_G(x) - z. \tag{2}$$

Given a fixed number $n \geq 2$ of firms, the expected markup is equal to the expected value of the difference between the first and second order statistic, that is, $\mathbb{E}(M_n - S_n)$. It is straightforward to verify that the following holds.[9]

$$\mathbb{E}(M_n - S_n) = n\mathbb{E}(M_n - M_{n-1}). \tag{3}$$

Next, using integration by parts yields[10]

$$\mathbb{E}(M_n - S_n) = \int nG(x)^{n-1}(1 - G(x))dx. \tag{4}$$

[9] Note $\mathbb{E}(M_n) = \int x dH_n(x)$ where $H_n(x) = G(x)^n$, the cdf of the first order statistic, and $\mathbb{E}(S_n) = \int x dH_n^2(x)$ where $H_n^2(x) = G(x)^n + nG(x)^{n-1}(1 - G(x))$, the cdf of the second order statistic.

[10] To simplify notation throughout the article, we use $\int$ to denote $\int_{x_0}^{\bar{x}}$ whenever no confusion is possible.

Therefore, when there are two or more firms, the expected markup given n firms is[11]

$$\mu_n = \int nG(x)^{n-1}(1 - G(x))dx. \tag{5}$$

Expected markup $\mu(\theta)$. We now calculate the *ex ante expected markup* $\mu(\theta)$, that is, the average markup across consumers, given that the number of competing firms n is stochastic and $n \sim P_n(\theta)$, where θ is the expected number of firms. In particular, we define $\mu(\theta)$ as the expected markup conditional on at least one firm, that is, $n \geq 1$, which is given by

$$\mu(\theta) = \frac{\sum_1 P_n(\theta)\mu_n}{1 - P_0(\theta)}. \tag{6}$$

Lemma 1 follows directly from the above expressions.

Lemma 1. For any $\theta \in \Theta$, the equilibrium expected markup is

$$\mu(\theta) = \frac{\sum_1 P_n(\theta) \int nG(x)^{n-1}(1 - G(x))dx + P_1(\theta)(x_0 - z)}{1 - P_0(\theta)}. \tag{7}$$

It is unclear from examining (7) whether $\mu(\theta)$ is increasing or decreasing in the expected number of firms, or *degree of competition*. That is, it is unclear whether competition is *price-increasing* ($\mu'(\theta) > 0$) or *price-decreasing* ($\mu'(\theta) < 0$). In Section 6, however, we will derive a simple expression for the expected markup which delivers a precise condition under which competition is either price-increasing or price-decreasing.

Before presenting our main results, we first derive their necessary components.

Expected demand. Let $D(\theta)$ denote the expected demand faced by single firm, that is, the probability of a sale. The following lemma gives us a general expression for the expected demand $D(\theta)$.

Lemma 2. For any $\theta \in \Theta$, the expected demand is given by

$$D(\theta) = \frac{1 - P_0(\theta)}{\theta}. \tag{8}$$

Proof. For a fixed number of firms $n \geq 1$, the expected demand D_n for a single firm's product in equilibrium is $D_n = \int g(x)G(x)^{n-1}dx$, which is equal to $1/n$.[12] Here, D_n is equal to the probability, $G(x)^{n-1}$, that the firm's utility shock x is higher than that of all $n - 1$ other firms, weighted by the pdf $g(x)$. The expected demand if $n \sim P_n(\theta)$ is $D(\theta) = \sum_1 Q_n(\theta)D_n$. Using the identity $\theta Q_n(\theta) = nP_n(\theta)$, plus $D_n = 1/n$, yields $D(\theta) = \frac{1}{\theta}\sum_1 P_n(\theta)$. □

Consumer's expected utility. We now derive expressions for the distribution of the consumer's utility and the expected value of this distribution, which we call the consumer's *expected utility*. When there is at least one firm, the consumer's utility equals the highest utility shock x among n draws when $n \sim P_n(\theta)$. When there are no firms, the consumer receives utility z, their outside option.

[11] Observe that this is equivalent to the well-known expression, $\mathbb{E}(M_n - S_n) = \mathbb{E}_{H_n}\left(\frac{1-G(x)}{g(x)}\right)$.

[12] Note that this expression is equivalent to the expression for the expected market share in symmetric equilibrium found in equation (10) of Perloff and Salop (1985).

Let $H_n(.)$ denote the cdf of the distribution of the maximum of n draws from $G(x)$, that is, $H_n(x) = G(x)^n$. The cdf of the distribution of the consumer's utility when $n \sim P_n(\theta)$ is denoted by $H(.;\theta)$, which is given by:

$$H(x;\theta) = \begin{cases} \sum P_n(\theta)G(x)^n \text{ if } x \in [z, \bar{x}) \\ 0 \text{ if } x \in [0, z) \end{cases}. \tag{9}$$

The distribution $H(.;\theta)$ has support $[x_0, \bar{x}) \cup \{z\}$. If $P_0(\theta) > 0$, it features a mass point at z because with probability $P_0(\theta)$ there are no firms competing for the consumer and their utility equals their outside option z.

Let $M(\theta)$ denote the *expected utility* of the consumer, that is, $M(\theta) \equiv \mathbb{E}_H(x)$.

Lemma 3. For any $\theta \in \Theta$, the consumer's expected utility is

$$M(\theta) = \int xh(x;\theta)dx + P_0(\theta)z. \tag{10}$$

Proof. Starting with (9) and the fact that $M(\theta) \equiv \mathbb{E}_H(x)$, we obtain (10), where $h(.;\theta)$ is the pdf given by $h(x;\theta) = \frac{d}{dx}\sum P_n(\theta)G(x)^n$ for $x \in [x_0, \bar{x})$. □

5. Lead example: Poisson distribution

■ In this section, we discuss our lead example for the distribution P_n of the number of competing firms. To enable direct comparison, we first discuss a standard environment where the number of competing firms is deterministic and constant across consumers.

Example: Deterministic number of firms. Consider a standard environment where the number n of competing firms is deterministic. This can be nested as a special case of our general framework which allows n to be random. Let $\theta \in \mathbb{N}$ where $\theta \geq 2$. Suppose that $P_n(\theta) = 1$ if $n = \theta$ and $P_n(\theta) = 0$ otherwise. For any $n \in \mathbb{N}$, we have $P_n : \Theta \to [0, 1]$ where $\Theta = \{2, 3, 4, \ldots\}$ and $\sum P_n(\theta) = 1$ where $\mathbb{E}(n) = \theta$.

For any $\theta \in \Theta$, Lemma 1 implies the equilibrium expected markup is

$$\mu(\theta) = \int \theta G(x)^{\theta-1}(1 - G(x))dx, \tag{11}$$

and Lemma 2 implies the expected demand is given by

$$D(\theta) = \frac{1}{\theta}. \tag{12}$$

We have $H(x;\theta) = G(x)^\theta$ and Lemma 3 implies the consumer's expected utility is

$$M(\theta) = \int \theta G(x)^{\theta-1}xg(x)dx. \tag{13}$$

To understand better this expression for the expected markup, (3) implies that[13]

$$\mu(\theta) = \frac{M(\theta) - M(\theta - 1)}{D(\theta)}. \tag{14}$$

This equation is intuitive: it says the expected markup $\mu(\theta)$ is equal to the marginal contribution, $M(\theta) - M(\theta - 1)$, of an additional firm to the consumer's expected utility, divided by the expected demand $D(\theta)$. We might wonder, does this intuitive expression generalize to settings in which the number of competing firms is random?

[13] I thank an anonymous referee for pointing out this difference equation for the deterministic example.

Example: Poisson distribution. We now consider our lead example for the distribution P_n of the number of competing firms. Specifically, assume that the distribution P_n is Poisson.[14] For any $n \in \mathbb{N}$, the probability there are n competing firms is $P_n(\theta) = \frac{\theta^n e^{-\theta}}{n!}$. We have $P_n : \Theta \to [0, 1]$ where $\Theta = (0, \infty)$, and $\sum P_n(\theta) = 1$ where $\mathbb{E}(n) = \theta$. Lemmas 1, 2, and 3 all apply.

Expected markup. Substituting $P_n(\theta) = \frac{\theta^n e^{-\theta}}{n!}$ into (7) from Lemma 1 yields[15]

$$\mu(\theta) = \frac{\int \theta e^{-\theta(1-G(x))}(1 - G(x))dx + \theta e^{-\theta}(x_0 - z)}{1 - e^{-\theta}}. \tag{15}$$

Expected demand. Starting with $P_0(\theta) = e^{-\theta}$ and using (8) from Lemma 2, we have

$$D(\theta) = \frac{1 - e^{-\theta}}{\theta}. \tag{16}$$

Consumer's expected utility. First, by substituting in $P_n(\theta) = \frac{\theta^n e^{-\theta}}{n!}$, we obtain

$$\sum P_n(\theta)G(x)^n = e^{-\theta(1-G(x))}. \tag{17}$$

Starting with (10) and using (17), plus $h(x; \theta) = \frac{d}{dx} \sum P_n(\theta)G(x)^n$ for $x \in [x_0, \bar{x})$,

$$M(\theta) = \int \theta e^{-\theta(1-G(x))} x g(x)dx + e^{-\theta} z. \tag{18}$$

Next, applying Leibniz' integral rule and using integration by parts, we obtain

$$M'(\theta) = \int e^{-\theta(1-G(x))}(1 - G(x))dx + e^{-\theta}(x_0 - z). \tag{19}$$

Simple expression for expected markup. For this example, we can obtain a simple expression that relates the expected markup $\mu(\theta)$ to the consumer's expected utility $M(\theta)$ and the expected demand $D(\theta)$. Dividing expression (19) for $M'(\theta)$ by expression (16) for expected demand, we obtain expression (15) for the expected markup. Therefore, we have

$$\mu(\theta) = \frac{M'(\theta)}{D(\theta)}. \tag{20}$$

Expression (20) says that the expected markup $\mu(\theta)$, which is the expected value of the difference between the highest and second-highest utility shock, is equal to the marginal contribution $M'(\theta)$ of an increase in the expected number of firms θ to the consumer's expected utility, divided by the expected demand.

This expression is a direct analogy of expression (14), which holds in the standard environment where the number of firms n is deterministic (and greater than or equal to two). Importantly, however, there are some crucial differences.

First, $\theta \in \mathbb{R}_+$ is the *expected* number of firms, which is continuous (not discrete) and $M(.)$ is differentiable. Second, $M(\theta)$ incorporates the consumer's outside option. Third, we allow the possibility that the number of competing firms is zero, one, two, or more. Finally, $M(\theta)$ depends not only on the distribution of utility shocks, but also on the *distribution of the number of competing firms*. Given this, we might wonder, does expression (20) hold only for the Poisson distribution, or does it hold more generally?

[14] With free entry of firms, the Poisson distribution could be endogenized, that is, by considering mixed strategies of potential entrants and then taking the limit as the number of potential entrants becomes large.

[15] Here we use the fact that $\sum_{n=1}^{\infty} \frac{(\theta G(x))^{n-1}}{(n-1)!} = \sum_{n=0}^{\infty} \frac{(\theta G(x))^n}{n!} = e^{\theta G(x)}$.

6. Simple expression for expected markup

■ In this section, we present a general condition on the distribution P_n which ensures that the simple expression (20) for the expected markup holds more generally. This condition is, in fact, equivalent to a very natural condition called *invariance* in Lester, Visschers, and Wolthoff (2015). In Appendix A, we provide an intuitive description of invariance and we prove that this property is equivalent to condition (21) in Condition 2.

Condition 2. The distribution P_n is twice-differentiable and it satisfies:

$$-\theta P_n'(\theta) = (n+1)P_{n+1}(\theta) - nP_n(\theta), \tag{21}$$

or, equivalently,

$$-P_n'(\theta) = Q_{n+1}(\theta) - Q_n(\theta), \tag{22}$$

for all $n \in \mathbb{N}$ such that $P_n(\theta) > 0$ and all $\theta \in \Theta$.

Notice that the equivalence of (21) and (22) follows from the identity $\theta Q_n(\theta) = nP_n(\theta)$.

Examples. It is straightforward to verify that the Poisson distribution satisfies Condition 2. More generally, this condition holds for a broader class of distributions. In Appendix A, we show that any distribution in the negative binomial family of distributions satisfies Condition 2. In particular, this family includes the geometric distribution as a special case and the Poisson distribution as a limiting case. The binomial distribution also satisfies Condition 2. In fact, the entire family of mixed Poisson distributions satisfies this condition.[16]

Proposition 1. If P_n satisfies Condition 2, the expected markup is given by

$$\mu(\theta) = \frac{M'(\theta)}{D(\theta)}. \tag{23}$$

Proof. Combining Lemma 1 and Lemma 2 yields

$$\mu(\theta) = \frac{\frac{1}{\theta}\sum_1 P_n(\theta)\int nG(x)^{n-1}(1-G(x))dx + \frac{P_1(\theta)}{\theta}(x_0 - z)}{D(\theta)}. \tag{24}$$

Using the identity $\theta Q_n(\theta) = nP_n(\theta)$ and rearranging, this is equivalent to

$$\mu(\theta) = \frac{\int\sum_1 Q_n(\theta)G(x)^{n-1}(1-G(x))dx + \frac{P_1(\theta)}{\theta}(x_0 - z)}{D(\theta)}. \tag{25}$$

If P_n satisfies Condition 2, Lemma L3 in Appendix A says that

$$M'(\theta) = -\int\sum P_n'(\theta)G(x)^n dx - P_0'(\theta)(x_0 - z). \tag{26}$$

Also, if P_n satisfies Condition 2 then $-\theta P_0'(\theta) = P_1(\theta)$. So, $\mu(\theta) = M'(\theta)/D(\theta)$ if and only if

$$-\int\sum P_n'(\theta)G(x)^n dx = \int\sum_1 Q_n(\theta)G(x)^{n-1}(1-G(x))dx. \tag{27}$$

Rearranging the right hand side, using the fact that $Q_0(\theta) = 0$, this is equivalent to

$$-\int\sum P_n'(\theta)G(x)^n dx = \int\sum(Q_{n+1}(\theta) - Q_n(\theta))G(x)^n dx. \tag{28}$$

[16] A mixed Poisson distribution is a Poisson distribution $P_n(\lambda)$ with parameter λ, where λ is itself a positive random variable. It can be shown that any mixed Poisson distribution has a representation that satisfies invariance (as defined in Appendix A) and it therefore satisfies Condition 2.

If P_n satisfies Condition 2, then (22) implies (28) and thus (23) is proven. □

To understand the intuition behind our result better, suppose that $x_0 = z$. A necessary and sufficient condition for Proposition 1 is given by (27). The left side of (27) can be rewritten as $\frac{d}{d\theta}\sum P_n(\theta)\mathbb{E}(M_n)$ using Lemma L3 in Appendix A, and the right side of (27) can be rewritten as $\sum_1 Q_n(\theta)\mathbb{E}(M_n - M_{n-1})$ using (3) and (4), therefore

$$\frac{d}{d\theta}\sum P_n(\theta)\mathbb{E}(M_n) = \sum\nolimits_1 Q_n(\theta)\mathbb{E}(M_n - M_{n-1}). \tag{29}$$

This form is more intuitive. The left term of (29) is the marginal increase in $M(\theta)$ from the consumer's perspective, and the right term is the expected value of the difference $M_n - M_{n-1}$ from the firms' perspective. By (3), the right term of (29) is equal to the average difference between the highest and second-highest utility shock (divided by the number of competing firms), which is what firms expect to be paid under personalized pricing.

When the number of firms is deterministic, that is, the distribution P_n is degenerate, the left sum in (29) is not differentiable. However, the discrete analogue of the left term is the difference, $M(\theta) - M(\theta - 1)$, which is always equal to the right term of (29).

When the number of competing firms is random, the average difference $M_n - M_{n-1}$ is *not* necessarily equal to the marginal increase in $M(\theta)$. To see this, condition (29) is equivalent to (27), which is equivalent to (28). Clearly, condition (28) holds if P_n satisfies Condition 2, but it may or may not hold for arbitrary distributions P_n.

Another way to understand the intuition behind the simple expression presented in Proposition 1 is to think about this result in terms of firms' expected profits. Let $\Pi(\theta)$ denote the *ex ante* expected payoff for an entering firm. Lemma 4 provides an expression for $\Pi(\theta)$.[17]

Lemma 4. The *ex ante* expected payoff for a firm is equal to

$$\Pi(\theta) = D(\theta)\mu(\theta), \tag{30}$$

where $D(\theta)$ is expected demand and $\mu(\theta)$ is the expected markup.

Proof. The expected payoff for a firm is given by $\Pi(\theta) = \sum_1 Q_n(\theta)D_n\mu_n$. Substituting in the identity $\theta Q_n(\theta) = nP_n(\theta)$ and using $D_n = 1/n$, we obtain $\Pi(\theta) = \frac{1}{\theta}\sum_1 P_n(\theta)\mu_n$. Finally, using expression (6) and (8) yields (30). □

Corollary 1 follows immediately from Proposition 1 and Lemma 4.

Corollary 1. If P_n satisfies Condition 2, we have $\Pi(\theta) = M'(\theta)$.

Corollary 1 says that, if the distribution P_n satisfies Condition 2, then firms' expected payoff $\Pi(\theta)$ is equal to the marginal increase $M'(\theta)$ in the consumer's expected utility that results from an increase in the expected number of firms. This is the key to understanding the simple expression (23) in Proposition 1 and it is closely related to the question of whether the expected number of competing firms would be *efficient* if there was entry of firms.

Efficient entry of firms. Suppose the expected number of firms θ is not exogenous but is instead determined by a zero profit condition. If firms pay an entry cost $k > 0$, the zero profit condition says that any equilibrium $\theta^* \in \Theta$ satisfies $\Pi(\theta) = k$. We can interpret k as representing any fixed costs related to entry as a firm (in contrast to c, the marginal cost of producing one unit).

[17] Note: this does *not* say $\mathbb{E}(D_n)\mathbb{E}(\mu_n) = \mathbb{E}(D_n\mu_n)$ because $D(\theta)$ is the expected demand from a *firm's* perspective, but $\mu(\theta)$ is the expected markup from the *consumer's* perspective (conditional on $n \geq 1$).

Suppose also that a social planner were to choose the expected number of firms θ^P that maximizes the expected social surplus minus entry costs, $\Omega(\theta) \equiv M(\theta) - c - k\theta$. The first-order condition for the planner's problem says that any $\theta^P \in \Theta$ satisfies $M'(\theta) = k$.[18]

In this section, we ask the following questions. Under which conditions does there exist a unique equilibrium θ^* and a unique social planner's solution θ^P? Under which conditions is firm entry efficient under personalized pricing, that is, $\theta^P = \theta^*$?

First, we describe an additional assumption which ensures the function $M(.)$ has the properties in Lemma 5. The negative binomial family of distributions satisfies Condition 3.[19]

Condition 3. For any $\theta \in \Theta = (0, \infty)$ and $y \in [0, 1)$, the distribution P_n satisfies

1. $\frac{d}{d\theta} \sum P_n(\theta)y^n < 0$ and $\frac{d^2}{d\theta^2} \sum P_n(\theta)y^n > 0$;
2. $\lim_{\theta\to\infty} \frac{d}{d\theta} \sum P_n(\theta)y^n = 0$ and $\lim_{\theta\to 0} Q_1(\theta) = 1$.

Lemma 5. If P_n satisfies Conditions 2 and 3, then $M(.)$ has the following properties:

1. For any $\theta \in \Theta$, we have $M'(\theta) > 0$ and $M''(\theta) < 0$.
2. We have $\lim_{\theta\to 0} M'(\theta) = \mathbb{E}_G(x) - z$ and $\lim_{\theta\to\infty} M'(\theta) = 0$.

Proposition 2 provides sufficient conditions under which there exists a unique equilibrium θ^* and a unique social planner's solution θ^P, and firm entry is efficient.

Proposition 2. With free entry of firms, if P_n satisfies Conditions 2 and 3, then if $k < \mathbb{E}_G(x) - z$,

1. There exists a unique equilibrium expected number of firms $\theta^* \in \Theta$.
2. There exists a unique socially optimal $\theta^P \in \Theta$.
3. Firm entry is efficient: $\theta^P = \theta^*$.

Proof. With free entry of firms, the zero profit condition says $\Pi(\theta) = k$, which is equivalent to $M'(\theta) = k$ if P_n satisfies Condition 2 by Corollary 1. If Condition 3 holds, we have $M''(\theta) < 0$ by Lemma 5. Also, Lemma 5 says that $\lim_{\theta\to 0} M'(\theta) = \mathbb{E}_G(x) - z$ and $\lim_{\theta\to\infty} M'(\theta) = 0$ if P_n satisfies Condition 3. Therefore, there exists a unique solution $\theta^* \in \Theta$ provided that $k < \mathbb{E}_G(x) - z$.[20]

The first-order condition for the planner's problem says that the planner's choice $\theta^P \in \Theta$ satisfies the same equation, $M'(\theta) = k$. Therefore, if Condition 3 holds, there exists a unique solution $\theta^P \in \Theta$ if $k < \mathbb{E}_G(x) - z$. Clearly, $\theta^* = \theta^P$. □

In the next section, we exploit the simplicity of the expression in Proposition 1 to derive a condition that is both necessary and sufficient for competition to be price-increasing.

7. When is competition price-increasing?

■ In this section, we use the simple expression for the expected markup in Proposition 1 to obtain a simple expression for the markup elasticity. This delivers a general condition under which competition is price-increasing. We show that the local curvature of the consumer's expected utility $M(\theta)$ is key to understanding the impact of expected competition on prices.

[18] For simplicity, we consider a single consumer and the *expected* number of firms θ. However, we could consider an environment with a large number of consumers L and then determine the equilibrium number of entering firms V. Letting $\theta \equiv V/L$, the equilibrium θ^* and the planner's choice θ^P would be the same.

[19] Note that for $y = 0$, Condition 3 implies $P_0'(\theta) < 0$, $P_0''(\theta) > 0$, and $\lim_{\theta\to\infty} P_0'(\theta) = 0$ using $0^0 = 1$.

[20] Anderson et al. (1995) consider the Perloff-Salop model and show that log-concavity is a sufficient condition for the existence of equilibrium when there is free entry of firms. In our environment, this assumption is not required because $\Pi'(\theta) = M''(\theta) < 0$, even in cases where the expected markup is increasing in θ.

Before presenting Proposition 3, we provide some preliminary definitions. The *demand elasticity* $\varepsilon_D(\theta)$ is the elasticity of the "demand" function $D(.)$, given by $\varepsilon_D(\theta) \equiv \frac{-D'(\theta)\theta}{D(\theta)}$. The *markup elasticity* $\varepsilon_\mu(\theta)$ is the elasticity of the expected markup $\mu(.)$, defined by $\varepsilon_\mu(\theta) \equiv \frac{\mu'(\theta)\theta}{\mu(\theta)}$.

Proposition 3 also features a measure of the *local* curvature, or degree of concavity, of the function $M(.)$. This measure of curvature is defined as follows:

$$r_M(\theta) \equiv \frac{-M''(\theta)\theta}{M'(\theta)}. \tag{31}$$

Formally, this is essentially the Arrow-Pratt coefficient of relative risk aversion of the function $M(.)$ at θ. However, it is important to remember that $M(.)$ is a function of the expected number of firms θ, not a standard utility function. We therefore refer to $r_M(\theta)$ as the *elasticity of marginal utility* because it is equal to the elasticity of $M'(.)$ at θ.

Proposition 3 presents a general condition – in terms of the demand elasticity $\varepsilon_D(\theta)$ and the elasticity of marginal utility $r_M(\theta)$ – under which competition is price-increasing.

Proposition 3. If P_n satisfies Condition 2, the markup elasticity is

$$\varepsilon_\mu(\theta) = \varepsilon_D(\theta) - r_M(\theta), \tag{32}$$

for any $\theta \in \Theta$. The expected markup $\mu(\theta)$ is strictly increasing in the expected number of firms, that is, $\mu'(\theta) > 0$, and competition is price-increasing at $\theta \in \Theta$, if and only if

$$r_M(\theta) < \varepsilon_D(\theta). \tag{33}$$

Proof. If P_n satisfies Condition 2, $\mu(\theta) = M'(\theta)/D(\theta)$ by Proposition 1. The elasticity of $\mu(\theta)$ equals the elasticity of the numerator $M'(\theta)$ minus the elasticity of the denominator $D(\theta)$,

$$\varepsilon_\mu(\theta) = \frac{M''(\theta)\theta}{M'(\theta)} - \frac{D'(\theta)\theta}{D(\theta)}. \tag{34}$$

Therefore, $\varepsilon_\mu(\theta) = -r_M(\theta) + \varepsilon_D(\theta)$ and we have $\mu'(\theta) > 0$ if and only if $r_M(\theta) < \varepsilon_D(\theta)$. □

The intuition behind this result can be explained in the following way. As the expected number of firms rises, the expected value of consumer's utility $M(\theta)$ increases. However, the marginal increase $M'(\theta)$ in the expected value $M(\theta)$ is decreasing in θ whenever $M''(\theta) < 0$. If the rate of decrease in $M'(\theta)$ is sufficiently low, that is, if $M''(\theta)$ is not *too* negative and $r_M(\theta)$ is not too high relative to the demand elasticity $\varepsilon_D(\theta)$ (i.e., $M(.)$ is not *too* concave), then greater competition is price-increasing, that is, $\mu'(\theta) > 0$.

This condition differs from existing results in Weyl and Fabinger (2013) and Quint (2014) that imply competition is price-decreasing when the distribution of utility shocks is log-concave. Importantly, our criterion is *local*, not global. Whether or not condition (33) holds depends crucially on the local curvature of the consumer's expected utility $M(\theta)$ at a particular value of θ. This depends not only on the properties of the distribution of utility shocks G, but also on the expected number of firms θ, the value of the consumer's outside option z, and the distribution of the number of competing firms P_n. Given that (33) is a local condition, markups can vary *nonmonotonically* with the expected number of firms.

Consumer surplus. We know that competition is price-increasing whenever condition (33) holds, but the effect on consumer surplus is unclear. To examine this question, we define the *consumer surplus* by $\Delta(\theta) \equiv M(\theta) - (1 - P_0(\theta))\mu(\theta)$. That is, we measure consumer surplus as the consumer's expected utility $M(\theta)$ minus the expected payment by the consumer (i.e., the probability that a consumer purchases the good from a firm, $1 - P_0(\theta)$, multiplied by the expected markup). We define the *consumer surplus share* by $\Delta_s(\theta) \equiv \Delta(\theta)/M(\theta)$.

Proposition 4 presents a simple expression for both the consumer surplus and the consumer surplus share in terms of the function $M(.)$ when the distribution P_n satisfies Condition 2. Before stating our result, we define the elasticity of $M(.)$ by $\eta_M(\theta) \equiv \frac{M'(\theta)\theta}{M(\theta)}$.

Proposition 4. If P_n satisfies Condition 2, the consumer surplus is given by

$$\Delta(\theta) = M(\theta) - \theta M'(\theta), \tag{35}$$

and the consumer surplus share is given by

$$\Delta_s(\theta) = 1 - \eta_M(\theta). \tag{36}$$

If P_n satisfies Conditions 2 and 3, the consumer surplus is strictly increasing in the expected number of firms, that is, $\Delta'(\theta) > 0$ for any $\theta \in \Theta$.

Proof. Starting with $\Delta(\theta) \equiv M(\theta) - (1 - P_0(\theta))\mu(\theta)$, we can use $\mu(\theta) = M'(\theta)/D(\theta)$ from Proposition 1 if Condition 2 holds, and the fact that $D(\theta) = (1 - P_0(\theta))/\theta$ from Lemma 2, to obtain (35). Dividing (35) by $M(\theta)$ yields (36). Next, differentiating (35) yields $\Delta'(\theta) = -\theta M''(\theta)$, so $\Delta'(\theta) > 0$ if $M''(\theta) < 0$. Applying Lemma 5, which uses Condition 3, we obtain $\Delta'(\theta) > 0$. □

Proposition 4 says that, if the distribution P_n satisfies Conditions 2 and 3, the consumer surplus is always strictly increasing in the expected number of firms. Intuitively, this is because the benefit consumers receive from having higher expected utility when there are more firms more than offsets any possible increase in the expected payment by the consumer, even when $\mu'(\theta) > 0$. This suggests that although greater competition can indeed be price-increasing, consumers are always better off – as measured by the consumer surplus.

It is important to bear in mind that this result hinges on our interpretation of random draws from the distribution G as *utility* shocks. As discussed in Gabaix et al. (2016), this distribution can be interpreted either as reflecting true preferences (which are welfare-relevant) or as representing "noise" such as consumer confusion or mistakes.[21] If we were to instead interpret the shocks as *random errors*, our welfare result would no longer hold.[22]

8. Application to auctions

■ All of our results can be applied directly to auctions where the number of bidders is stochastic. Consider a seller who runs a second-price auction for a single indivisible good. Buyers' valuations are private i.i.d. draws from a distribution G that satisfies Condition 1.

The number of bidders $n \in \mathbb{N}$ is given by a distribution $P_n : \Theta \to [0, 1]$ where $P_n(\theta)$ denotes the probability that a seller's auction has n bidders. The expected number of bidders is θ, which is exogenous. We assume there is no reserve price (i.e., $z = 0$).

When there are $n \geq 2$ bidders, all bidders make a bid equal to their own valuation. The bidder with the highest valuation wins the auction and the expected surplus for the winning bidder is the expected value of the difference between the highest and second-highest valuation, $\mathbb{E}(M_n - S_n)$, where M_n is the highest valuation and S_n is the second-highest valuation. When there is exactly one bidder, he gets the full surplus, $\mathbb{E}_G(x)$.

Let $V_B(\theta)$ denote the expected surplus for the winning bidder and let $V_S(\theta)$ denote the expected surplus for the seller. All of our results regarding the expected markup also hold for the

[21] There is a large literature studying how consumer confusion or errors can arise from various mechanisms such as obfuscation by firms. For example, see Gabaix and Laibson (2006), Spiegler (2006), Ellison and Ellison (2009), and Armstrong and Vickers (2012).

[22] For example, suppose the consumer surplus was $\Delta(\theta) = \bar{x} - \theta M'(\theta)$ where $\bar{x}$ is constant, instead of (35). In this case, $\Delta'(\theta) = -\theta M''(\theta) - M'(\theta)$ and thus $\Delta'(\theta) < 0$ if and only if $r_M(\theta) < 1$. Lemma B3 in Appendix B says that $r_M(\theta) \to 1 - \gamma \in (0, 1]$ as θ goes to infinity, suggesting that $\Delta'(\theta) < 0$ if the tail index $\gamma > 0$.

winning bidder's expected surplus. In particular, if the distribution P_n satisfies Condition 2, we obtain expression (37) by Proposition 1. Similarly, expression (35) in Proposition 4 for the consumer surplus represents the expected surplus for the seller.

Corollary 2. If P_n satisfies Condition 2, the expected surplus for the winning bidder is given by

$$V_B(\theta) = \frac{M'(\theta)}{D(\theta)}, \tag{37}$$

and the expected surplus for the seller is given by

$$V_S(\theta) = M(\theta) - \theta M'(\theta). \tag{38}$$

In this setting, $M(\theta)$ can be interpreted as the total expected surplus of the auction, and $D(\theta)$ can be interpreted as the probability of winning faced by each bidder. Recall that $r_M(\theta)$ is the elasticity of $M'(.)$ and $\varepsilon_D(\theta)$ is the elasticity of $D(.)$ at θ.

Corollary 3. If P_n satisfies Condition 2, the expected surplus for the winning bidder is strictly increasing in the expected number of bidders, that is, $V_B'(\theta) > 0$, if and only if

$$r_M(\theta) < \varepsilon_D(\theta). \tag{39}$$

By the revenue equivalence theorem, these results do not depend on the type of auction but apply more generally to *any* type of auction which satisfies the conditions of the revenue equivalence theorem (e.g., first-price, second-price, all-pay, or English auctions).

9. Asymptotic results

■ Gabaix et al. (2016) considers an environment where the number of firms is deterministic. The authors show that, in the limit as the number of firms $n \to \infty$, the markup elasticity $\varepsilon_\mu(n)$ converges to the *tail index* γ_G of the distribution of utility shocks. Competition is therefore either *asymptotically price-increasing* (i.e., $\mu'(n) > 0$ as $n \to \infty$) or *asymptotically price-decreasing* (i.e., $\mu'(n) < 0$ as $n \to \infty$) depending on whether the tail index is greater than or less than zero, that is, whether the distribution is fat-tailed or not.

In this section, we generalize these results to our environment. We provide asymptotic results only for our lead example: the Poisson distribution. We summarize the main results here and provide the preliminary lemmas and proofs in Appendix B.

Definition 1. We say that G is *well-behaved* if and only if $\lim_{x\to\bar{x}} \frac{1-G(x)}{g(x)} = a$ where $a \in \mathbb{R}^+ \cup \{+\infty\}$ and G has finite *tail index* $\gamma_G \in \mathbb{R}$ given by $\lim_{x\to\bar{x}} \frac{d}{dx}\left(\frac{1-G(x)}{g(x)}\right) = \gamma_G$.

We make the following assumption on the distribution G for our asymptotic results.

Condition 4. The distribution G is well-behaved with tail index $\gamma_G < 1$.

We adopt standard notation and write $F(y) \sim_{y\to\infty} F_L(y)$, or simply $F(y) \sim F_L(y)$, if and only if $\lim_{y\to\infty} \frac{F(y)}{F_L(y)} = 1$. To derive our results, we make use of Proposition 1, which implies $\mu(\theta) \sim_{\theta\to\infty} M'(\theta)\theta$, and Proposition 3, which says $\varepsilon_\mu(\theta) = \varepsilon_D(\theta) - r_M(\theta)$.

Proposition 5 presents the *asymptotic expected markup*, which is analogous to Proposition 2 of Gabaix et al. (2016). We recover that article's result that the markup elasticity converges to the *tail index*, that is, $\varepsilon_\mu(\theta) \to \gamma_G$. We also provide a necessary and sufficient condition under which the expected markup $\mu(\theta)$ converges to zero as $\theta \to \infty$.

Proposition 5. If P_n is Poisson and G satisfies Condition 4, then

1. The asymptotic expected markup is

$$\mu(\theta) \sim_{\theta\to\infty} \frac{\Gamma(1-\gamma_G)}{\theta g\big(G^{-1}\big(1-\frac{1}{\theta}\big)\big)}, \tag{40}$$

   where $\Gamma(t) \equiv \int_0^\infty y^{t-1}e^{-y}dy$ is the Gamma function.
2. In the limit as $\theta \to \infty$, we have $\mu(\theta) \to 0$ if and only if $\lim_{x\to\bar{x}} \frac{1-G(x)}{g(x)} = 0$.
3. In the limit as $\theta \to \infty$, we have $\varepsilon_\mu(\theta) \to \gamma_G$.

We can also determine the asymptotic value of the consumer surplus share, $\Delta_s(\theta) \equiv \Delta(\theta)/M(\theta)$. Proposition 6 says the consumer surplus share converges to one as $\theta \to \infty$ if the distribution G is bounded, but it converges to $1-\gamma_G \in (0,1]$ if G is unbounded.

Proposition 6. If P_n is Poisson and G satisfies Condition 4, then

1. In the limit as $\theta \to \infty$, we have $\Delta_s(\theta) \to 1-\gamma_G \in (0,1]$ if $\bar{x} = \infty$.
2. In the limit as $\theta \to \infty$, we have $\Delta_s(\theta) \to 1$ if $\bar{x} < \infty$.

10. Examples

■ We now present some examples in order to bring to life our results. For each distribution of utility shocks G, we consider the expected markup when the distribution P_n is Poisson, geometric, or degenerate (that is, the number of firms is deterministic).

Example 1: Exponential. Let $G(x) = 1 - e^{-a(x-x_0)}$ for $x \in [x_0, \infty)$ where $a \in (0, \infty)$. We know from Gabaix et al. (2016) that the asymptotic markup elasticity is $\varepsilon_\mu(\theta) = 0$, the tail index of G, and the asymptotic markup is $\mu(\theta) \sim 1/a$. If P_n is degenerate, the expected markup for $\theta \geq 2$ is

$$\mu(\theta) = \frac{1}{a}. \tag{41}$$

If P_n is Poisson, the expected markup is given by

$$\mu(\theta) = \frac{1}{a} + \frac{\theta e^{-\theta}}{1-e^{-\theta}}(x_0 - z). \tag{42}$$

If P_n is geometric, the expected markup is given by

$$\mu(\theta) = \frac{1}{a} + \frac{1}{1+\theta}(x_0 - z). \tag{43}$$

For both the Poisson and the geometric distribution, the expected markup is constant, $\mu(\theta) = 1/a$, and $\varepsilon_\mu(\theta) = 0$ if $z = x_0$ but it is *decreasing* in the expected number of firms, that is, $\mu'(\theta) < 0$, if $z < x_0$. In the limit as $\theta \to \infty$, we have $\mu(\theta) \sim 1/a$ and $\varepsilon_\mu(\theta) \to 0$.

Figure 1 provides a comparison of the behavior of the expected markup and the markup elasticity for the Poisson and geometric distributions when G is exponential with parameter values $a = 1$, $x_0 = 1$, and $z = 0$. We also show the deterministic markup and its elasticity.

Example 2: Uniform. Let $G(x) = x - x_0$ for $x \in [x_0, x_0 + 1]$. The uniform distribution has a log-concave density and therefore we know that the deterministic markup (44) is decreasing in θ by the standard results. The asymptotic markup elasticity is $\varepsilon_\mu(\theta) = -1$, the tail index of G, and the asymptotic markup is $\mu(\theta) \sim 1/\theta$. If P_n is degenerate, the expected markup for $\theta \geq 2$ is

$$\mu(\theta) = \frac{1}{\theta+1}. \tag{44}$$

FIGURE 1

Expected markup (left panel) and markup elasticity (right panel). The distribution G is exponential with $a = 1$, $x_0 = 1$, and $z = 0$. The asymptotic elasticity is the tail index $\gamma = 0$

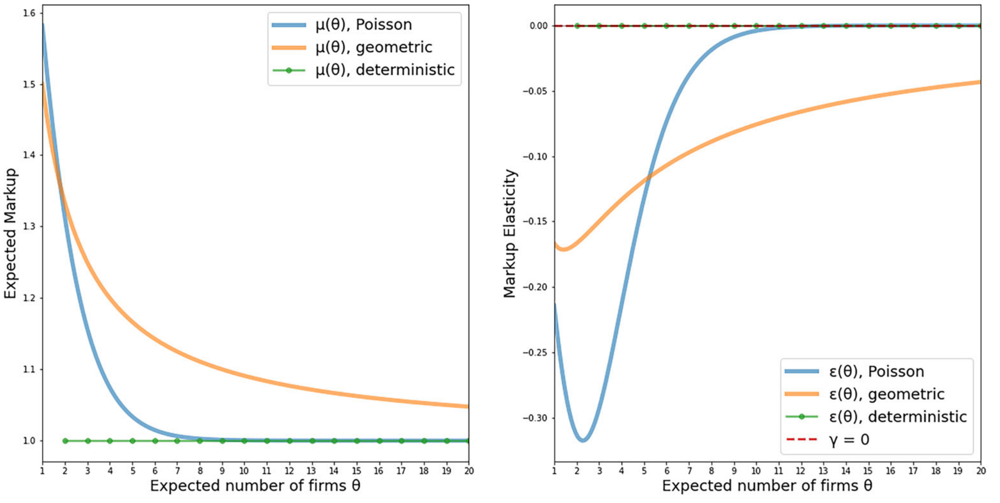

If P_n is Poisson, the expected markup is given by

$$\mu(\theta) = \frac{1}{\theta}\left(\frac{1 - e^{-\theta} - \theta e^{-\theta}}{1 - e^{-\theta}}\right) + \frac{\theta e^{-\theta}}{1 - e^{-\theta}}(x_0 - z). \tag{45}$$

If P_n is geometric, the expected markup is given by

$$\mu(\theta) = \frac{1}{\theta}\left(\frac{(1+\theta)\ln(1+\theta)}{\theta} - 1\right) + \frac{1}{1+\theta}(x_0 - z). \tag{46}$$

For both the Poisson and the geometric distribution, the expected markup is *decreasing* in the expected number of firms, that is, $\mu'(\theta) < 0$, regardless of the value of the outside option z. In the limit as $\theta \to \infty$, we have $\mu(\theta) \sim 1/\theta$ and $\varepsilon_\mu(\theta) \to -1$.

Figure 2 provides a comparison of the behavior of the expected markup and the markup elasticity when G is uniform with parameter values $x_0 = 1$ and $z = 0$. For the Poisson and geometric distributions, as for the deterministic case, the markup elasticity is always negative. However, its behavior is quite different: the markup elasticity is strictly decreasing for both the geometric and deterministic cases, but nonmonotonic for the Poisson.

Example 3: Pareto. Let $G(x) = 1 - \left(\frac{x}{x_0}\right)^{-1/\lambda}$ for $x \in [x_0, \infty)$ where $\lambda \in (0, 1)$. The Pareto distribution has a log-convex density, so we expect that the deterministic markup (47) will be increasing in θ. The asymptotic markup elasticity is $\varepsilon_\mu(\theta) = \lambda > 0$, the tail index of G, and the asymptotic markup is $\mu(\theta) \sim \lambda x_0 \theta^\lambda \Gamma(1-\lambda)$. If P_n is degenerate, the expected markup for $\theta \geq 2$ is

$$\mu(\theta) = \frac{\lambda x_0 \Gamma(\theta+1)\Gamma(1-\lambda)}{\Gamma(\theta+1-\lambda)}. \tag{47}$$

If P_n is Poisson, the expected markup is given by

$$\mu(\theta) = \frac{\lambda x_0 \theta^\lambda \gamma(1-\lambda, \theta)}{1 - e^{-\theta}} + \frac{\theta e^{-\theta}}{1 - e^{-\theta}}(x_0 - z), \tag{48}$$

where $\gamma(s, z) \equiv \int_0^z t^{s-1} e^{-t} dt$, the Lower Incomplete Gamma Function.

FIGURE 2

Expected markup (left panel) and markup elasticity (right panel). The distribution G is uniform with $x_0 = 1$ and $z = 0$. The asymptotic elasticity is the tail index $\gamma = -1$

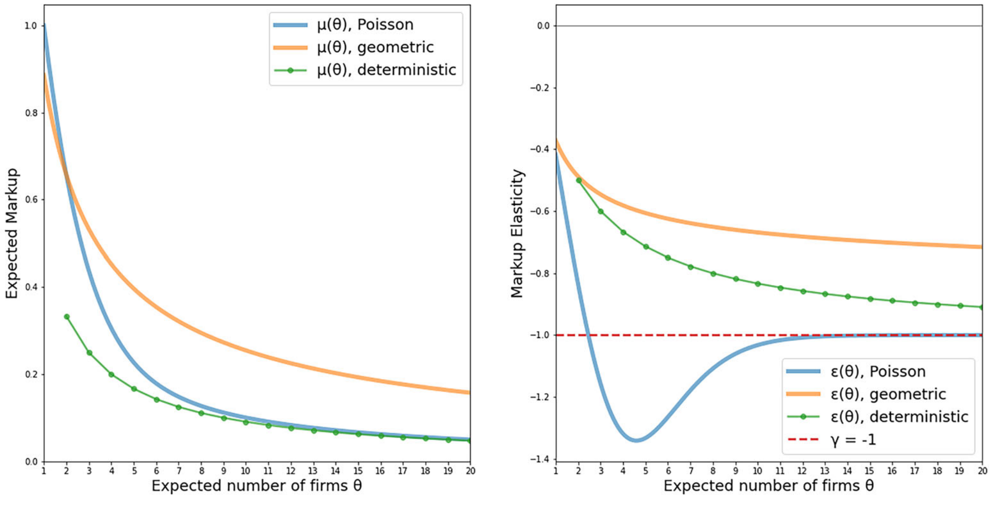

If $z = x_0$, the expected markup is always *increasing* in the expected number of firms, θ. If $z = 0$, however, the expected markup $\mu(\theta)$ varies *nonmonotonically* with the expected number of firms. In the limit as $\theta \to \infty$, we have $\varepsilon_\mu(\theta) \to \lambda$, the tail index of G.

If P_n is geometric, the expression for the expected markup is more complicated.[23]

$$\mu(\theta) = \frac{\lambda x_0 \theta^\lambda}{(1+\lambda)}\left(\frac{1+\theta}{\theta}\right)(L(\theta,\lambda) - L(\theta,\lambda,x_0)) + \frac{1}{1+\theta}(x_0 - z). \tag{49}$$

Figure 3 provides a comparison of the behavior of the expected markup and the markup elasticity when G is Pareto with parameter values $\lambda = 0.25$, $x_0 = 1$, and $z = 0$. For the Poisson distribution, competition is price-decreasing when the expected number of firms θ is less than around five, but price-increasing after that. For the geometric distribution, competition is price-decreasing when θ is less than eight, but price-increasing after that.

Discussion of examples. To understand these examples better, we can decompose the expected markup as follows:

$$\mu(\theta; P_n, G, z) = \rho(\theta; P_n)\mu_1(G, z) + (1 - \rho(\theta; P_n))\mu_2(\theta; P_n, G), \tag{50}$$

where $\rho(\theta; P_n) = \frac{P_1(\theta)}{1-P_0(\theta)}$, the probability that $n = 1$ (i.e., a "local" monopoly), $\mu_1(G, z) = \mathbb{E}_G(x) - z$, the monopoly markup, and $\mu_2(\theta; P_n, G)$ is the expected markup if $n \geq 2$. The deterministic markup is a special case of (50) where $\rho(\theta; P_n) = 0$ and P_n is degenerate.

Effect of outside option. The behavior of the expected markup $\mu(\theta)$ depends crucially on the value of the consumer's outside option, z. This is because there may be only one firm selling to the consumer. The probability of this outcome is $\rho(\theta; P_n)$, which depends on the distribution P_n. When there is only one firm, the expected markup is given by $\mu_1(G, z) = \mathbb{E}_G(x) - z$, which clearly depends on z. In environments where the number of firms is large, this may not be rele-

[23] Here, we define $L(\theta, \lambda, x) \equiv \left(\frac{x}{\theta^\lambda}\right)^{1/\lambda+1} {}_2F_1\left(2, 1+\lambda; 2+\lambda; -\left(\frac{x}{\theta^\lambda}\right)^{1/\lambda}\right)$ and $L(\theta, \lambda) \equiv \lim_{x\to\infty} L(\theta, \lambda, x)$, where ${}_2F_1(a, b; c; d)$ is the hypergeometric function.

FIGURE 3

Expected markup (left panel) and markup elasticity (right panel). The distribution G is Pareto with $\lambda = 0.25$, $x_0 = 1$, and $z = 0$. The asymptotic elasticity is the tail index $\gamma = 0.25$

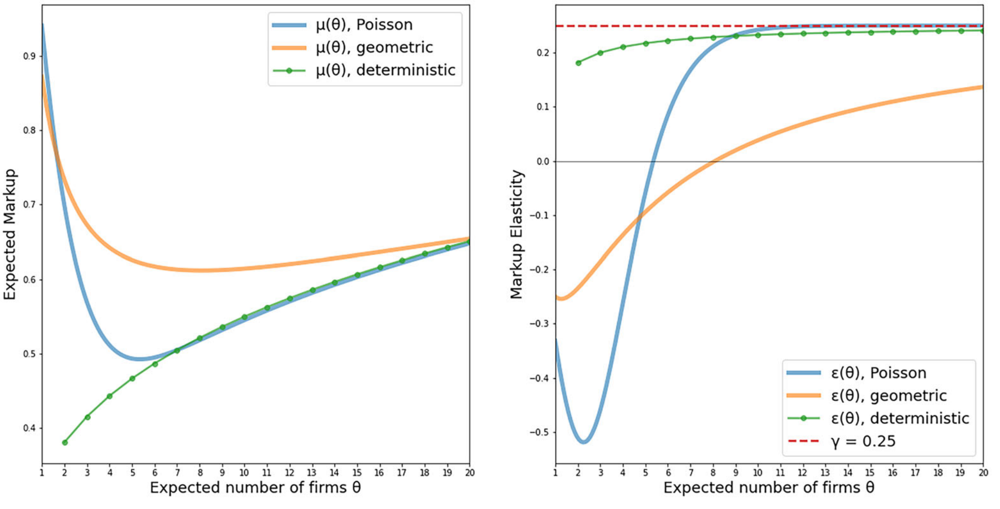

vant. However, in environments with a relatively small expected number of competing firms, the possibility of a local monopoly may be significant.

Nonmonotonicity of expected markup. The value of the consumer's outside option can influence whether the behavior of the expected markup $\mu(\theta)$ is nonmonotonic. For the Pareto distribution, the expected markup is always a nonmonotonic function of the expected number of firms (for both the Poisson and geometric examples) whenever the consumer's outside option is strictly less than the minimum firm-specific utility shock, that is, $z < x_0$. On the other hand, in the special case where $z = x_0$, there is *no* nonmonotonicity at all.

For other distributions, however, nonmonotonicity can still arise even when $z = x_0$. In this case, $\mu_1(G) = \mathbb{E}_G(x) - x_0$ and the effect of the outside option is eliminated. For example, if G is the Fréchet distribution (which has neither a log-concave nor a log-convex density), the expected markup is always a nonmonotonic function of the expected number of firms *regardless* of the outside option (including when $z = x_0$, as shown in Figure 4).

The reason behind the nonmonotonicity (for both the Pareto and Fréchet examples) is the fact that the expected markup μ_n falls from the monopoly markup, $\mathbb{E}_G(x) - z$, to the expected markup for $n = 2$, but is increasing in n after that. This is true if $z = x_0$ for the Fréchet example, but is only true for the Pareto example if $z < x_0$. For the deterministic case where P_n is degenerate and $\theta \geq 2$, the expected markup $\mu(\theta)$ is always strictly increasing. For both the Poisson and geometric distributions, the expected markup $\mu(\theta)$ is a weighted average of the monopoly markup and the expected markups μ_n for $n \geq 2$ with weights that vary depending on both the distribution P_n and the expected number of firms θ.

Effect of distribution P_n. Looking at Figures 1–4, it is clear that the distribution P_n (e.g., Poisson, geometric, or deterministic) can affect outcomes such as (*i*) the level of the expected markup, (*ii*) the level of the markup elasticity, (*iii*) the expected number of firms at which competition switches from being price-decreasing to price-increasing, and (*iv*) the rate of convergence of the markup elasticity to its asymptotic value.

FIGURE 4

Expected markup (left panel) and markup elasticity (right panel). The distribution G is Fréchet with $\lambda = 0.25$, $x_0 = 1$, and $z = x_0$. The asymptotic elasticity is the tail index $\gamma = 0.25$

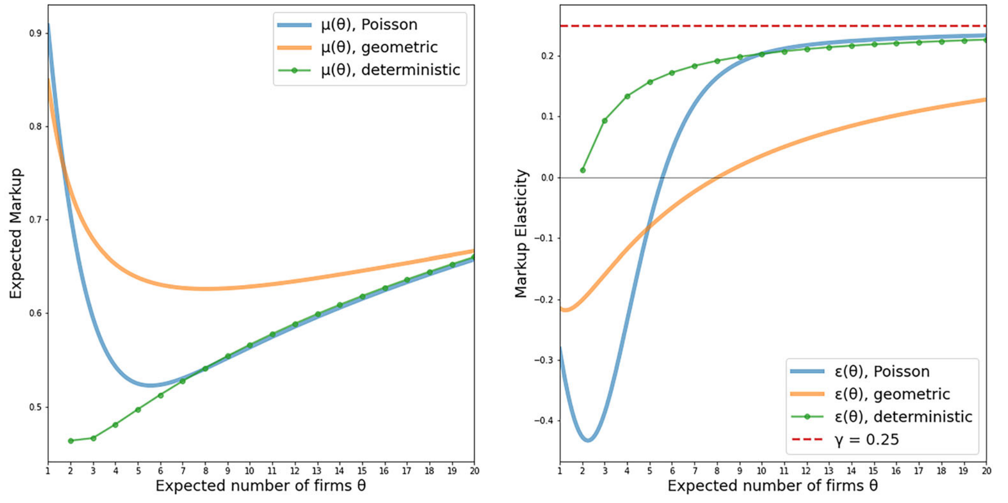

11. Conclusion

■ This article studies the effect of expected competition on markups in a random utility model where the number of competing firms may differ across consumers. There may be either no firms, one firm, or two or more firms competing for a consumer. Prices are determined by "personalized pricing", that is, the equilibrium markup equals the difference between the highest and second-highest utility shock for any given consumer. We show that, under a precise condition on the distribution P_n of the number of competing firms, we can obtain a simple expression for the expected markup in terms of the key object: the consumer's *expected utility* as a function of the expected number of firms. The simplicity of our expression is closely related to another result: firm entry is efficient whenever P_n satisfies the same condition.

Our simple expression for the expected markup reveals that the impact of competition on prices depends crucially on the local curvature of the function $M(.)$. In particular, competition is price-increasing if and only if the *elasticity of marginal utility*, defined as $-M''(\theta)\theta/M'(\theta)$, is strictly less than the elasticity of demand with respect to θ. Whether or not this is true depends not only on properties of the distribution of utility shocks, but also on the expected number of firms and the value of the consumer's outside option. In addition, it depends on the *distribution* of the number of competing firms in a consumer's choice set.

Allowing the number of competing firms to vary across consumers in a random manner is useful for modelling environments that feature various *frictions* (e.g., search frictions). Somewhat surprisingly, however, this approach can still yield a remarkably simple expression for the expected markup, which features significantly greater generality (and arguably greater tractability) than the analogous expression when the number of firms is deterministic. This suggests that a similar approach may be fruitfully applied to many other problems in industrial organization. We leave this as a potential avenue for future research.

Appendix A

Invariance. The probability generating function (PGF) $\mathbb{G}$ of the distribution P_n is defined as follows: $\mathbb{G}(y;\theta) \equiv \sum P_n(\theta)y^n$ for $y \in [0, 1]$ and all $\theta \in \Theta$. A distribution P_n is called *invariant* in Lester, Visschers, and Wolthoff (2015) if and only if the PGF takes the following form.

Definition A1. A distribution P_n is *invariant* if and only if, for $y \in [0, 1]$ and all $\theta \in \Theta$,

$$\mathbb{G}(y;\theta) = P_0(\theta(1-y)). \tag{A1}$$

Invariance is in fact a fairly intuitive assumption. To see this, suppose there is a continuum of red and white balls in an urn and the proportion of red balls is $y \in [0, 1]$. Consider the following two alternative exercises.

First, suppose that we select a random number n of balls from the urn, where $n \sim P_n(\theta)$ and $\mathbb{E}(n) = \theta$. The probability that every draw is red is equal to $\sum P_n(\theta)y^n$.

Second, suppose that we first split the balls into two different urns: the red balls go in a red urn, and the white balls go in a white urn. We select a random number n_r of balls from the red urn, where $n_r \sim P_n(\theta y)$ and $\mathbb{E}(n_r) = \theta y$. We also select a random number n_w of balls from the white urn, where $n_w \sim P_n(\theta(1-y))$ and $\mathbb{E}(n_w) = \theta(1-y)$. The total expected number of balls drawn is again θ. In this case, the probability that every draw is red equals the probability that there are no balls drawn from the white urn, $P_0(\theta(1-y))$.

Invariance of the distribution P_n means that the probability there are no red balls drawn is the *same* for both exercises, that is, this probability is "invariant" to whether we first draw and then split; or first split and then draw. That is, invariance says the following holds:

$$\sum P_n(\theta)y^n = P_0(\theta(1-y)). \tag{A2}$$

The following lemma provides an alternative condition that is equivalent to (A1), as described in Lester, Visschers, and Wolthoff (2015). This will prove useful for proving our equivalence result.

Lemma A1. A distribution P_n is invariant if and only if

$$P_n(\theta) = \frac{(-1)^n\theta^n P_0^{(n)}(\theta)}{n!}, \tag{A3}$$

for all $n \in \mathbb{N}$ such that $P_n(\theta) > 0$ and all $\theta \in \Theta$, where $P_0^{(n)}$ is the n-th derivative of P_0.

Proof. If (A1) then (A3) follows from the general property of probability generating functions that, for all $n \in \mathbb{N}$ such that $P_n(\theta) > 0$ and all $\theta \in \Theta$,

$$P_n(\theta) = \frac{1}{n!}\frac{\partial^n}{\partial y^n}\bigg|_{y=0} \sum P_n(\theta)y^n. \tag{A4}$$

If (A3) then (A1) using the Taylor series expansion of $P_0(z)$ at θ, where $z = \theta(1-y)$. □

We are now in a position to prove our equivalence result.

Lemma A2. A distribution P_n is invariant if and only if

$$-\theta P_n'(\theta) = (n+1)P_{n+1}(\theta) - nP_n(\theta), \tag{A5}$$

for all $n \in \mathbb{N}$ such that $P_n(\theta) > 0$ and all $\theta \in \Theta$.

Proof. If P_n is invariant, then $\mathbb{G}(y;\theta) = P_0(\theta(1-y))$. By Lemma A1, this is true if and only if $P_n(\theta)$ can be written as (A3). Differentiating (A3), we obtain

$$P_n'(\theta) = \frac{(-1)^n\theta^n P_0^{(n+1)}(\theta)}{n!} + \frac{(-1)^n n\theta^{n-1} P_0^{(n)}(\theta)}{n!}, \tag{A6}$$

which is equivalent to

$$-\theta P_n'(\theta) = (n+1)\frac{(-1)^{n+1}\theta^{n+1} P_0^{(n+1)}(\theta)}{(n+1)!} - n\frac{(-1)^n\theta^n P_0^{(n)}(\theta)}{n!}, \tag{A7}$$

and thus (A5) holds. Conversely, if (A5) holds, then (22) implies

$$-\sum P_n'(\theta)y^n = \sum (Q_{n+1}(\theta) - Q_n(\theta))y^n. \tag{A8}$$

Rearranging the right-hand side, using $Q_0(\theta) = 0$, this is equivalent to

$$-\frac{d}{d\theta}\sum P_n(\theta)y^n = \sum\nolimits_1 Q_n(\theta)y^{n-1}(1-y), \tag{A9}$$

which implies that

$$\sum\nolimits_1 Q_n(\theta)y^{n-1}(1-y) + \frac{d}{d\theta}\sum P_n(\theta)y^n = 0. \tag{A10}$$

Rearranging, using the identity $\theta Q_n(\theta) = nP_n(\theta)$, this is equivalent to

$$\frac{1}{\theta}(1-y)\sum\nolimits_1 nP_n(\theta)y^{n-1} + \frac{d}{d\theta}\sum P_n(\theta)y^n = 0. \tag{A11}$$

Letting $\mathbb{G}_1(y;\theta) = \frac{\partial}{\partial y}\mathbb{G}(y;\theta)$ and $\mathbb{G}_2(y;\theta) = \frac{\partial}{\partial \theta}\mathbb{G}(y;\theta)$, this is equivalent to

$$\frac{1}{\theta}(1-y)\mathbb{G}_1(y;\theta) + \mathbb{G}_2(y;\theta) = 0. \tag{A12}$$

Applying Proposition 2 in Lester, Visschers, and Wolthoff (2015), this implies that the meeting fee equals zero in their environment.[24] By Proposition 4 in Lester, Visschers, and Wolthoff (2015), P_n is therefore invariant. □

Negative binomial distribution. The negative binomial family is a two-parameter family of distributions that count the number n of failures before $r \in \mathbb{N}\backslash\{0\}$ successes, where the probability of success is $r/(r+\theta)$. If P_n is negative binomial, the probability there are $n \in \mathbb{N}$ competing firms is

$$P_n(\theta) = \binom{n+r-1}{n}\left(\frac{r}{r+\theta}\right)^r\left(\frac{\theta}{r+\theta}\right)^n. \tag{A13}$$

We have $P_n : \Theta \to [0,1]$ where $\Theta = (0,\infty)$, and $\sum P_n(\theta) = 1$ where $\mathbb{E}(n) = \theta$.

For any value of $r \in \mathbb{N}\backslash\{0\}$, the corresponding distribution P_n satisfies Condition 2 by Lemma A2 because it is invariant. To see this, note that the probability generating function of the negative binomial distribution is

$$\mathbb{G}(y;\theta) = \left(\frac{r}{r+\theta(1-y)}\right)^r = P_0(\theta(1-y)). \tag{A14}$$

In the limit as $r \to \infty$, we obtain the Poisson distribution, described in Section 5.

In the special case where $r = 1$, we obtain the geometric distribution:

$$P_n(\theta) = \left(\frac{1}{1+\theta}\right)\left(\frac{\theta}{1+\theta}\right)^n. \tag{A15}$$

Given that we provide the geometric distribution ($r = 1$) as an example in Section 10, we derive the key expressions for this distribution here.

Expected markup. Substituting $P_n(\theta)$ into (7) from Lemma 1,

$$\mu(\theta) = \frac{1}{\theta}\sum\nolimits_1\left(\frac{\theta}{1+\theta}\right)^n\int nG(x)^{n-1}(1-G(x))dx + \frac{1}{1+\theta}(x_0 - z). \tag{A16}$$

Rearranging and simplifying the above yields

$$\mu(\theta) = \frac{1}{1+\theta}\int\sum\nolimits_1 n\left(\frac{\theta G(x)}{1+\theta}\right)^{n-1}(1-G(x))dx + \frac{1}{1+\theta}(x_0 - z), \tag{A17}$$

which, using the fact that $\sum_1 nr^{n-1} = \frac{1}{(1-r)^2}$ for $r \in (0,1)$, is equivalent to

$$\mu(\theta) = \int\frac{1+\theta}{(1+\theta(1-G(x)))^2}(1-G(x))dx + \frac{1}{1+\theta}(x_0 - z). \tag{A18}$$

Expected demand. Starting with $P_0(\theta) = \frac{1}{1+\theta}$ and using expression (8) from Lemma 2, the *ex ante* expected demand for a single firm's product is given by

$$D(\theta) = \frac{1}{1+\theta}. \tag{A19}$$

Consumer's expected utility. First, by substituting in $P_n(\theta)$, we have

$$\sum P_n(\theta)G(x)^n = \frac{1}{1+\theta(1-G(x))}. \tag{A20}$$

Starting with (10) and using (A20), plus $h(x;\theta) = \frac{d}{dx}\sum P_n(\theta)G(x)^n$ for $x \in [x_0, \bar{x})$,

$$M(\theta) = \int\frac{\theta x g(x)}{(1+\theta(1-G(x)))^2}dx + \frac{1}{1+\theta}z. \tag{A21}$$

[24] Note that our equation (A12) implies equation (6) in Lester, Visschers, and Wolthoff (2015) with $t = 0$, in the special case where $y = G(x)$ and the seller's own valuation, denoted y in Lester, Visschers, and Wolthoff (2015), is equal to x_0 and $\bar{x} < \infty$.

Next, applying Leibniz' integral rule and using integration by parts yields

$$M'(\theta) = \int \frac{1}{(1+\theta(1-G(x)))^2}(1-G(x))dx + \frac{1}{(1+\theta)^2}(x_0 - z). \tag{A22}$$

It is straightforward to verify that the expected markup is given by $\mu(\theta) = M'(\theta)/D(\theta)$.

Technical lemmas.

Lemma L1. If P_n satisfies Condition 2, then

$$\frac{d}{d\theta}\int \left(1 - \sum P_n(\theta)G(x)^n\right)dx = -\int \sum P_n'(\theta)G(x)^n dx, \tag{A23}$$

and

$$\frac{d}{d\theta}\sum P_n(\theta)G(x)^n = \sum P_n'(\theta)G(x)^n. \tag{A24}$$

Proof. Given that Condition 2 is equivalent to (22), we have $\left|P_n'(\theta)\right| = |Q_n(\theta) - Q_{n+1}(\theta)|$ and therefore $\left|P_n'(\theta)\right| \leq 1$ because $\left|P_n'(\theta)\right| \leq \max\{Q_n(\theta), Q_{n+1}(\theta)\}$ and $Q_n(\theta) \leq 1$ for all $n \in \mathbb{N}$. So, $\left|P_n'(\theta)y^n\right| \leq y^n$ and $\sum y^n$ converges for $y \in (0,1)$. So, if Condition 2 holds, then we have (A24).

Next, we show that Leibniz' integral rule applies to $\frac{d}{d\theta}\int (1 - \sum P_n(\theta)G(x)^n)\,dx$. Noting that $\frac{d}{d\theta}(1 - \sum P_n(\theta)G(x)^n) = -\frac{d}{d\theta}\sum P_n(\theta)G(x)^n$, in order to apply Leibniz' integral rule (because we allow $\bar{x} = \infty$) we need to show there exists a function $\phi(x)$ such that $\left|\frac{d}{d\theta}\sum P_n(\theta)G(x)^n\right| \leq \phi(x)$ and $\int_{x_0}^{\bar{x}} \phi(x) < \infty$. Using the result (A24), we have $\left|\frac{d}{d\theta}\sum P_n(\theta)G(x)^n\right| = \left|\sum P_n'(\theta)G(x)^n\right|$. If Condition 2 holds, then (A8) implies

$$\sum P_n'(\theta)G(x)^n = -\sum (Q_{n+1}(\theta) - Q_n(\theta))G(x)^n. \tag{A25}$$

Next, by rearranging, and using the fact that $Q_0(\theta) = 0$, we have

$$\sum P_n'(\theta)G(x)^n = -(1 - G(x))\sum Q_{n+1}(\theta)G(x)^n. \tag{A26}$$

Now, $G(x)^n \leq 1$ and therefore $\sum Q_{n+1}(\theta)G(x)^n \leq 1$. So, $\left|\sum P_n'(\theta)G(x)^n\right| \leq \phi(x) \equiv 1 - G(x)$ where $\int_{x_0}^{\bar{x}}(1 - G(x))dx < \infty$ because G has a finite mean by Condition 1. So, $\frac{d}{d\theta}\int (1 - \sum P_n(\theta)G(x)^n)\,dx = -\int \frac{d}{d\theta}\sum P_n(\theta)G(x)^n dx$, which equals $-\int \sum P_n'(\theta)G(x)^n$ by (A24). □

Lemma L2. If P_n satisfies Condition 2, then

$$-\frac{d}{d\theta}\int \sum P_n'(\theta)G(x)^n dx = -\int \sum P_n''(\theta)G(x)^n dx, \tag{A27}$$

and

$$\frac{d^2}{d\theta^2}\sum P_n(\theta)G(x)^n = \sum P_n''(\theta)G(x)^n. \tag{A28}$$

Proof. Differentiating condition (21) in Condition 2, we have

$$\theta P_n''(\theta) = (n-1)P_n'(\theta) - (n+1)P_{n+1}'(\theta). \tag{A29}$$

Therefore, letting $y = G(x)$, we obtain

$$\left|\theta P_n''(\theta)y^n\right| \leq \left|nP_n'(\theta)y^n\right| + \left|P_n'(\theta)y^n\right| + \left|(n+1)P_{n+1}'(\theta)y^n\right|, \tag{A30}$$

and thus $\left|P_n''(\theta)y^n\right| \leq \frac{1}{\theta}2(n+1)y^n$ because $\left|P_n'(\theta)\right| \leq 1$ for all $n \in \mathbb{N}$, as shown in Lemma L1. Letting $B_\epsilon(\theta) = [\theta - \epsilon, \theta + \epsilon]$ and $K = 1/(\theta - \epsilon)$, we have $1/\hat{\theta} \leq K$ for all $\hat{\theta} \in B_\epsilon(\theta)$ and therefore $\left|P_n''(\theta)y^n\right| \leq 2K(n+1)y^n$. Also, $\sum (n+1)y^n = \sum_1 ny^{n-1}$ converges for $y \in (0,1)$, so we have $\frac{d}{d\theta}\sum P_n'(\theta)G(x)^n = \sum P_n''(\theta)G(x)^n$. Together with Lemma L1, this implies (A28).

Next, Leibniz' integral rule applies to $\frac{d}{d\theta}\int \sum P_n'(\theta)G(x)^n dx$ provided there exists a function $\hat{\phi}(x)$ such that $\left|\sum P_n''(\theta)G(x)^n\right| \leq \hat{\phi}(x)$ and $\int_{x_0}^{\bar{x}} \hat{\phi}(x) < \infty$. Using condition (A29) above,

$$\theta \sum P_n''(\theta)G(x)^n = \sum (n-1)P_n'(\theta)G(x)^n - \sum (n+1)P_{n+1}'(\theta)G(x)^n, \tag{A31}$$

which can be shown to be equivalent to

$$\theta \sum P_n''(\theta)G(x)^n = -(1 - G(x))\sum_1 nP_n'(\theta)G(x)^{n-1} - \sum P_n'(\theta)G(x)^n. \tag{A32}$$

Also, applying condition (21) in Condition 2 yields

$$\sum\nolimits_1 nP_n'(\theta)G(x)^{n-1} = \sum\nolimits_1 n(Q_n(\theta) - Q_{n+1}(\theta))G(x)^{n-1}, \tag{A33}$$

which can be rearranged to

$$\sum\nolimits_1 nP_n'(\theta)G(x)^{n-1} = -(1 - G(x))\sum nQ_{n+1}(\theta)G(x)^{n-1} + \sum Q_{n+1}(\theta)G(x)^n. \tag{A34}$$

Therefore, we have

$$\left|\theta\sum P_n''(\theta)G(x)^n\right| = \left|(1 - G(x))\sum\nolimits_1 nP_n'(\theta)G(x)^{n-1} + \sum P_n'(\theta)G(x)^n\right|. \tag{A35}$$

Substituting equation (A34) into the above gives us

$$\theta\left|\sum P_n''(\theta)G(x)^n\right| = \left|\begin{array}{c} -(1 - G(x))^2\sum nQ_{n+1}(\theta)G(x)^{n-1} \\ +(1 - G(x))\sum Q_{n+1}(\theta)G(x)^n + \sum P_n'(\theta)G(x)^n \end{array}\right|, \tag{A36}$$

and thus equation (A26) from the proof of Lemma L1 implies

$$\theta\left|\sum P_n''(\theta)G(x)^n\right| = (1 - G(x))^2\sum nQ_{n+1}(\theta)G(x)^{n-1}. \tag{A37}$$

Now, $\sum nQ_{n+1}(\theta)G(x)^{n-1} \leq \sum_1 nQ_{n+1}(\theta)$ because $G(x)^{n-1} \leq 1$ for all $n \geq 1$ and

$$\sum\nolimits_1 nQ_{n+1}(\theta) = \sum\nolimits_2 (n-1)Q_n(\theta) \leq \sum\nolimits_1 nQ_n(\theta). \tag{A38}$$

Therefore, using the identity $\theta Q_n(\theta) = nP_n(\theta)$, we have

$$\left|\sum P_n''(\theta)G(x)^n\right| \leq (1 - G(x))^2\frac{\sum n^2P_n(\theta)}{\theta^2}. \tag{A39}$$

Letting $\sigma^2(\theta) = \sum n^2P_n(\theta) - \theta^2$, the variance of P_n, we can write $\frac{1}{\theta^2}\sum n^2P_n(\theta) = \frac{\sigma^2(\theta)}{\theta^2} + 1$ where $\sum n^2P_n(\theta) < \infty$ and $\theta > 0$ by assumption. Letting $B_\epsilon(\theta) = [\theta - \epsilon, \theta + \epsilon]$ and $\bar{K} = \max_{\hat{\theta}\in B_\epsilon(\theta)}\left\{\frac{\sigma^2(\hat{\theta})}{\hat{\theta}^2}\right\}$, we have $\frac{\sigma^2(\hat{\theta})}{\hat{\theta}^2} \leq \bar{K}$ for all $\hat{\theta} \in B_\epsilon(\theta)$ and thus $\frac{1}{\theta^2}\sum n^2P_n(\theta) \leq \bar{K} + 1$. Therefore, $\left|\sum P_n''(\theta)G(x)^n\right| \leq \hat{\phi}(x)$ where $\hat{\phi}(x) \equiv (\bar{K} + 1)(1 - G(x))^2$ and $(\bar{K} + 1)\int_{x_0}^{\bar{x}}(1 - G(x))^2 < \infty$ because G has a finite mean by Condition 1. Therefore, (A27) is proven. □

Lemma L3. If P_n satisfies Condition 2, the derivative of $M(.)$ is given by

$$M'(\theta) = -\int\sum P_n'(\theta)G(x)^n dx - P_0'(\theta)(x_0 - z), \tag{A40}$$

and we have

$$\frac{d}{d\theta}\sum P_n(\theta)\mathbb{E}(M_n) = -\int\sum P_n'(\theta)G(x)^n dx. \tag{A41}$$

Proof. We can write

$$M(\theta) = \sum P_n(\theta)\mathbb{E}(M_n) - P_0(\theta)(x_0 - z). \tag{A42}$$

Using integration by parts, we have $\mathbb{E}(M_n) = \int xh_n(x)dx = x_0 + \int(1 - H_n(x))dx$, so

$$M(\theta) = x_0 + \int\sum P_n(\theta)(1 - H_n(x))dx - P_0(\theta)(x_0 - z). \tag{A43}$$

Next, using the fact that $\sum P_n(\theta) = 1$ and $H_n(x) = G(x)^n$, we obtain

$$M(\theta) = x_0 + \int\left(1 - \sum P_n(\theta)G(x)^n\right)dx - P_0(\theta)(x_0 - z). \tag{A44}$$

Finally, differentiating expression (A44) with respect to θ yields

$$M'(\theta) = \frac{d}{d\theta}\int\left(1 - \sum P_n(\theta)G(x)^n\right)dx - P_0'(\theta)(x_0 - z). \tag{A45}$$

If P_n satisfies Condition 2, applying Lemma L1 yields the following:

$$M'(\theta) = -\int\frac{d}{d\theta}\sum P_n(\theta)G(x)^n dx - P_0'(\theta)(x_0 - z). \tag{A46}$$

Applying Lemma L1 again yields (A40). Together with (A42), this implies (A41). □

Additional proofs.

Proof of Lemma 1. Substituting expressions (5) and (2) into (6), we obtain

$$\mu(\theta) = \frac{\sum_2 P_n(\theta) \int nG(x)^{n-1}(1 - G(x))dx + P_1(\theta)(\mathbb{E}_G(x) - z)}{1 - P_0(\theta)}, \tag{A47}$$

which is equivalent to

$$\mu(\theta) = \frac{\sum_1 P_n(\theta) \int nG(x)^{n-1}(1 - G(x))dx + P_1(\theta)\big(\mathbb{E}_G(x) - \int(1 - G(x))dx - z\big)}{1 - P_0(\theta)}. \tag{A48}$$

Given that G has a finite mean, $\lim_{x\to\bar{x}} x(1 - G(x)) = 0$ and therefore $\mathbb{E}_G(x) - \int(1 - G(x))dx = x_0$ using integration by parts. Substituting into (A48), we obtain (7). □

Proof of Lemma 5. Part (1). Consider expression (A46) for $M'(\theta)$. Clearly, Condition 3 implies $M'(\theta) > 0$. Next, differentiating (A40) yields

$$M''(\theta) = -\frac{d}{d\theta}\int \sum P_n'(\theta)G(x)^n dx - P_0''(\theta)(x_0 - z). \tag{A49}$$

Applying Lemma L2 to (A49), we obtain

$$M''(\theta) = -\int \frac{d^2}{d\theta^2}\sum P_n(\theta)G(x)^n dx - P_0''(\theta)(x_0 - z). \tag{A50}$$

Clearly, Condition 3 implies $M''(\theta) < 0$.

Part (2). First, Condition 2 implies $M'(\theta) = \mu(\theta)D(\theta)$ and thus $\lim_{\theta\to 0} M'(\theta) = \lim_{\theta\to 0}\mu(\theta)D(\theta)$. So, using expression (A47), plus the fact that $\theta Q_n(\theta) = nP_n(\theta)$, we have

$$\lim_{\theta\to 0} M'(\theta) = \lim_{\theta\to 0}\sum\nolimits_2 Q_n(\theta)\int G(x)^{n-1}(1 - G(x))dx + \lim_{\theta\to 0}\frac{P_1(\theta)}{\theta}(\mathbb{E}_G(x) - z). \tag{A51}$$

Next, $\lim_{\theta\to 0}\sum_2 Q_n(\theta)\int G(x)^{n-1}(1 - G(x))dx = 0$, using the fact that

$$0 \le \sum\nolimits_2 Q_n(\theta)\int G(x)^{n-1}(1 - G(x))dx \le \int(1 - G(x))dx \sum\nolimits_2 Q_n(\theta), \tag{A52}$$

where $\int(1 - G(x))dx < \infty$ and $\lim_{\theta\to 0}\sum_2 Q_n(\theta) = 0$ because $\lim_{\theta\to 0} Q_1(\theta) = 1$ by Condition 3. Finally, $\lim_{\theta\to 0}\frac{P_1(\theta)}{\theta} = \lim_{\theta\to 0} Q_1(\theta) = 1$, and thus $\lim_{\theta\to 0} M'(\theta) = \mathbb{E}_G(x) - z$.

Next, starting with (A46), we obtain

$$\lim_{\theta\to\infty} M'(\theta) = -\int \lim_{\theta\to\infty}\frac{d}{d\theta}\sum P_n(\theta)G(x)^n dx - \lim_{\theta\to\infty} P_0'(\theta)(x_0 - z). \tag{A53}$$

Interchanging integral and limit is justified because $\left|\frac{d}{d\theta}\sum P_n(\theta)G(x)^n\right| \le 1 - G(x)$, as in Lemma L1, and $\int 1 - G(x)dx < \infty$ as G has a finite mean by Condition 1. Also, $\lim_{\theta\to\infty}\frac{d}{d\theta}\sum P_n(\theta)y^n = 0$ for all $y \in [0, 1)$ and $\lim_{\theta\to\infty} P_0'(\theta) = 0$ by Condition 3. Thus $\lim_{\theta\to\infty} M'(\theta) = 0$. □

Appendix B

Asymptotic results. Before presenting our results, we first define the notion of *regular variation*.[25]

Definition B1. We say that a function $h : \mathbb{R}_+ \to \mathbb{R}$ is *regularly varying at zero with index* ρ, and denote this by $h \in RV_\rho^0$, if and only if h is strictly positive in a neighborhood of zero and, for all $\lambda > 0$, we have $\lim_{t\to 0}\frac{h(\lambda t)}{h(t)} = \lambda^\rho$.

Theorem B1 is identical to Theorem 3 of Gabaix et al. (2016), except for the following differences. First, $H(.;\theta)$ is the distribution of the consumer's expected utility when n is stochastic and the distribution P_n of the number of competing firms n is Poisson. Second, we are taking the limit as the *expected number of firms*, $\theta = \mathbb{E}(n)$, goes to infinity, not as $n \to \infty$. Third, the consumer's outside option is z. Finally, we restrict attention to the case where $\zeta(x) \ge 0$, which is all that is required for our results.

The proof of Theorem B1 is somewhat simpler than that found in Gabaix et al. (2016).

[25] See Bingham, Goldie, and Teugels (1987) or Resnick (1987).

Theorem B1. Let $\zeta : [x_0, \bar{x}) \cup \{z\} \to \mathbb{R}^+$ be a function that satisfies $\zeta(x) \geq 0$ for all $x \in [x_0, \bar{x})$ and $\int |\zeta(x)g(x)|\, dx < \infty$. Suppose that $\hat{\zeta}(t) \equiv \zeta(G^{-1}(1-t)) \in RV_\rho^0$ with $\rho > -1$. If P_n is Poisson, then in the limit as $\theta \to \infty$, we have

$$\mathbb{E}_H(\zeta(x)) = \int \zeta(x)g(x)\theta e^{-\theta(1-G(x))}dx + e^{-\theta}\zeta(z) \sim \zeta\left(G^{-1}\left(1-\frac{1}{\theta}\right)\right)\Gamma(\rho+1), \tag{B1}$$

where $\Gamma(t) \equiv \int_0^\infty y^{t-1}e^{-y}dy$ is the Gamma function.

To apply Theorem B1, we need Condition 4, which says that the tail index of G is below one, that is, $\gamma_G < 1$. We will see that this assumption is sufficient to ensure that $\rho > -1$, and therefore Theorem B1 can be applied, in all cases required for our results.

If the distribution P_n satisfies Condition 2 (which is true for the Poisson distribution), Proposition 1 says that the expected markup can be expressed in terms of the derivative of the consumer's expected utility function $M(.)$. As a result, all of the asymptotic results we need, including the asymptotic behavior of the expected markup and the markup elasticity, depend *only* on the asymptotic behavior of the consumer's expected utility function $M(.)$ and its derivatives.

We now present Lemmas B2 and B3, which derive some results regarding the asymptotic behavior of $M(.)$ and its derivatives.

Lemma B2. If P_n is Poisson and G satisfies Condition 4, in the limit as $\theta \to \infty$ we have

$$M(\theta) \sim G^{-1}\left(1-\frac{1}{\theta}\right)\Gamma(1-\gamma_G) \text{ if } \bar{x} = \infty, \tag{B2}$$

and

$$\bar{x} - M(\theta) \sim \left(\bar{x} - G^{-1}\left(1-\frac{1}{\theta}\right)\right)\Gamma(1-\gamma_G) \text{ if } \bar{x} < \infty. \tag{B3}$$

Lemma B3 summarizes the asymptotic behavior of the derivative $M'(\theta)$ and the measure of curvature $r_M(\theta) \equiv \frac{-M''(\theta)\theta}{M'(\theta)}$. In the limit as the expected number of firms becomes large, $r_M(\theta) \to 1 - \gamma_G$. This result is used to prove part of Proposition 5, but it is interesting in its own right because it says that the *tail index* of the distribution of utility shocks – which is a measure of tail fatness – is equal to one minus the asymptotic value of $r_M(\theta)$ – a measure of local curvature of the consumer's expected utility function, $M(.)$.

Lemma B3. If P_n is Poisson and G satisfies Condition 4, the following hold:

1. In the limit as $\theta \to \infty$, we have

$$M'(\theta) \sim \frac{\Gamma(1-\gamma_G)}{\theta^2 g\left(G^{-1}\left(1-\frac{1}{\theta}\right)\right)}. \tag{B4}$$

2. We have $\lim_{\theta\to\infty} M'(\theta) = 0$.
3. In the limit as $\theta \to \infty$, we have $r_M(\theta) \to 1 - \gamma_G$.

Proofs of asymptotic results.

Proof of Theorem B1. Suppose that $\hat{\zeta}(t) \equiv \zeta(G^{-1}(1-t)) \in RV_\rho^0$ with $\rho > -1$. Changing variables by letting $t = 1 - G(x)$ and rewriting yields

$$\mathbb{E}_H(\zeta(x)) = \int_0^1 \theta e^{-\theta t}\zeta(G^{-1}(1-t))dt + e^{-\theta}\zeta(z). \tag{B5}$$

Rewriting, this is equivalent to

$$\mathbb{E}_H(\zeta(x)) = \int_0^1 \theta e^{-(\theta-1)t}\zeta(G^{-1}(1-t))e^{-t}dt + e^{-\theta}\zeta(z). \tag{B6}$$

Now define $\tilde{h}(t) \equiv \zeta(G^{-1}(1-t))e^{-t}$ and $\tilde{H}(t) \equiv \int_0^t \tilde{h}(y)dy$. Letting $\theta - 1 = \theta'$, we have

$$\mathbb{E}_H(\zeta(x)) = \int_0^1 \theta e^{-\theta' t}d\tilde{H}(t) + e^{-\theta}\zeta(z). \tag{B7}$$

Defining $\hat{h}(t) = h(t)$ for all $t \in [0, 1]$ and $\hat{h}(t) = 0$ for all $t \in (1, \infty)$, and $\hat{H}(t) \equiv \int_0^t \hat{h}(y)dy$,

$$\mathbb{E}_H(\zeta(x)) = \int_0^\infty \theta e^{-\theta' t}d\hat{H}(t) + e^{-\theta}\zeta(z). \tag{B8}$$

We can apply Karamata's Tauberian Theorem because $\hat{H}(t)$ is weakly positive and weakly increasing in t. This theorem says that if $\hat{H}(t) \in RV_{\alpha}^{0}$ then as $\theta' \to \infty$ we have

$$\int_0^{\infty} e^{-\theta' t} d\hat{H}(t) \sim \hat{H}(1/\theta')\Gamma(\alpha+1). \tag{B9}$$

Now, because $\hat{\zeta}(t) \equiv \zeta(G^{-1}(1-t)) \in RV_{\rho}^{0}$ with $\rho > -1$ by assumption, we have $h(t) \equiv \zeta(G^{-1}(1-t))e^{-t} \in RV_{\rho}^{0}$ with $\rho > -1$ because $e^{-t} \in RV_{0}^{0}$, and therefore also $\hat{h}(t) \in RV_{\rho}^{0}$ with $\rho > -1$. By Lemma A1.6 of Gabaix et al. (2016), this implies that $\hat{H}(t) \equiv \int_0^t \hat{h}(y)dy \in RV_{\rho+1}^{0}$ and therefore $\alpha = \rho + 1$, so we have

$$\int_0^{\infty} e^{-\theta' t} d\hat{H}(t) \sim \hat{H}(1/\theta')\Gamma(\rho+2), \tag{B10}$$

as $\theta' \to \infty$. By Lemma A1.6 of Gabaix et al. (2016), we also have $\lim_{x\to 0} \frac{x\hat{h}(x)}{\hat{H}(x)} = \rho + 1$, and thus $\hat{H}(x) \sim x\hat{h}(x)/(\rho+1)$ as $x \to 0$. Therefore as $\theta' \to \infty$ we have

$$\int_0^{\infty} e^{-\theta' t} d\hat{H}(t) \sim \frac{1}{\theta'} \frac{\hat{h}(1/\theta')\Gamma(\rho+2)}{\rho+1}. \tag{B11}$$

Given that $\Gamma(\rho+2)/(\rho+1) = \Gamma(\rho+1)$ and $\hat{h}(t) = \zeta(G^{-1}(1-t))e^{-t}$ for all $t \in [0, 1]$,

$$\int_0^{\infty} e^{-\theta' t} d\hat{H}(t) \sim \frac{1}{\theta'}\zeta\left(G^{-1}\left(1-\frac{1}{\theta'}\right)\right)e^{-1/\theta'}\Gamma(\rho+1). \tag{B12}$$

Finally, using the fact that $\theta' = \theta - 1$, in the limit as $\theta \to \infty$ we have

$$\mathbb{E}_H(\zeta(x)) = \int_0^{\infty} \theta e^{-\theta' t} d\hat{H}(t) + e^{-\theta}\zeta(z) \sim \zeta\left(G^{-1}\left(1-\frac{1}{\theta}\right)\right)\Gamma(\rho+1). \tag{B13}$$

□

Proof of Lemma B2. Case 1. Suppose that $\bar{x} = \infty$. Let $f_{\zeta} : \mathbb{N} \to \mathbb{R}$ be defined by $f_{\zeta}(n) = \int \zeta(x)h_n(x)dx$ and $\zeta(x) = x$. If P_n is Poisson, applying Theorem B1 and using $\lim_{\theta\to\infty} e^{-\theta} = 0$ yields

$$\sum P_n(\theta)f_{\zeta}(n) \sim_{\theta\to\infty} G^{-1}\left(1-\frac{1}{\theta}\right)\Gamma(1-\gamma_G). \tag{B14}$$

Theorem B1 applies because $\zeta(x) \geq 0$ and $\int xg(x)dx < \infty$ because G has a finite mean by Condition 1, plus $\hat{\zeta}(t) \equiv \zeta(G^{-1}(1-t)) \in RV_{\rho}^{0}$ where $\rho = -\gamma_G$ by Lemma 1 of Gabaix et al. (2016) because $\bar{x} = \infty$, and $\gamma_G < 1$ by Condition 4. Using (A42) for $M(\theta)$ and $\lim_{\theta\to\infty} e^{-\theta} = 0$ yields (B2).

Case 2. Suppose that $\bar{x} < \infty$. Let $f_{\zeta} : \mathbb{N} \to \mathbb{R}$ be defined by $f_{\zeta}(n) = \int \zeta(x)h_n(x)dx$ and $\zeta(x) = \bar{x} - x$. If P_n is Poisson, applying Theorem B1 and using $\lim_{\theta\to\infty} e^{-\theta} = 0$ yields

$$\sum P_n(\theta)f_{\zeta}(n) \sim_{\theta\to\infty} \left(\bar{x} - G^{-1}\left(1-\frac{1}{\theta}\right)\right)\Gamma(1-\gamma_G). \tag{B15}$$

Theorem B1 applies because $\zeta(x) \geq 0$ and $\int |\zeta(x)g(x)|\, dx < \infty$ as G has a finite mean by Condition 1, plus $\hat{\zeta}(t) \equiv \zeta(G^{-1}(1-t)) \in RV_{\rho}^{0}$ where $\rho = -\gamma_G$ by Lemma 1 of Gabaix et al. (2016) as $\bar{x} < \infty$, and $\gamma_G < 1$ by Condition 4. Using (A42) for $M(\theta)$ and $\lim_{\theta\to\infty} e^{-\theta} = 0$ yields (B3). □

Proof of Lemma B3. Part (1). First, suppose that $F(\theta) \sim_{\theta\to\infty} F_L(\theta)$, and either $\lim_{\theta\to\infty} F(\theta) = \infty$ or $\lim_{\theta\to\infty} F(\theta) = 0$. Then $\lim_{\theta\to\infty} \frac{F(\theta)}{F_L(\theta)} = 1$, which implies $\lim_{\theta\to\infty} \frac{F'(\theta)}{F_L'(\theta)} = 1$ by L'Hôpital's rule, so $F'(\theta) \sim_{\theta\to\infty} F_L'(\theta)$. If P_n is Poisson, we can apply this reasoning to (B2) and (B3) from Lemma B2. For both cases, we obtain

$$M'(\theta) \sim_{\theta\to\infty} \Gamma(1-\gamma_G)\frac{d}{d\theta}G^{-1}\left(1-\frac{1}{\theta}\right), \tag{B16}$$

and differentiating $G^{-1}\left(1-\frac{1}{\theta}\right)$ yields (B4).

Part (2). Letting $x = G^{-1}\left(1-\frac{1}{\theta}\right)$ in (B4), we obtain

$$\lim_{\theta\to\infty} M'(\theta) = \Gamma(1-\gamma_G)\lim_{x\to\bar{x}} \frac{(1-G(x))^2}{g(x)}. \tag{B17}$$

Therefore, $\lim_{\theta\to\infty} M'(\theta) = 0$ if and only if $\lim_{x\to\bar{x}} \frac{(1-G(x))^2}{g(x)} = 0$. Rewriting,

$$\lim_{x\to\bar{x}} \frac{(1-G(x))^2}{g(x)} = \lim_{x\to\bar{x}} x(1-G(x))\left(\frac{1-G(x)}{xg(x)}\right). \tag{B18}$$

We have $\lim_{x\to\bar{x}} x(1-G(x)) = 0$ because G has a finite mean by Condition 1. Given we assume Condition 4, we have $\lim_{x\to\bar{x}} \frac{1-G(x)}{xg(x)} = \lim_{x\to\bar{x}} \frac{\frac{1-G(x)}{g(x)}}{x}$ where $\lim_{x\to\bar{x}} \frac{1-G(x)}{g(x)} = a \in \mathbb{R}^+ \cup \{+\infty\}$. If $a \in \mathbb{R}^+$ then $\lim_{x\to\bar{x}} \frac{1-G(x)}{xg(x)} = 0$, and if

$a = \infty$ then L'Hôpital's rule yields $\lim_{x\to\bar{x}} \frac{1-G(x)}{xg(x)} = \lim_{x\to\bar{x}} \frac{d}{dx}\frac{1-G(x)}{g(x)} = \gamma_G \in \mathbb{R}$. Either way, $\lim_{x\to\bar{x}} \frac{(1-G(x))^2}{g(x)} = 0$ and thus $\lim_{\theta\to\infty} M'(\theta) = 0$.

Part (3). Next, let $M_L'(\theta) = \frac{\Gamma(1-\gamma_G)}{\theta^2 g(G^{-1}(1-\frac{1}{\theta}))}$. Letting $t = 1/\theta$, we can write $M_L'(\theta)$ as $H(t) \equiv \frac{\Gamma(1-\gamma_G)t^2}{g(G^{-1}(1-t))}$ for $t \in (0, \infty)$. Next, we can show that $h(t) \sim_{t\to 0} (1-\gamma_G)H(t)/t$.

Letting $x = G^{-1}(1-t)$, we have

$$H(t) = \frac{\Gamma(1-\gamma_G)(1-G(x))^2}{g(x)}. \tag{B19}$$

Differentiating the above, we obtain

$$h(t) = \Gamma(1-\gamma_G)\frac{2(1-G(x))g(x) + \frac{(1-G(x))^2 g'(x)}{g(x)}}{g(x)^2}. \tag{B20}$$

Therefore, we have

$$\lim_{t\to 0}\frac{h(t)t}{H(t)} = 2 + \lim_{x\to\infty}\frac{(1-G(x))g'(x)}{g(x)^2} = 1 - \lim_{x\to\infty}\frac{d}{dx}\left(\frac{1-G(x)}{g(x)}\right). \tag{B21}$$

So, $\lim_{t\to 0}\frac{h(t)t}{H(t)} = 1-\gamma_G$ by definition of the tail index γ_G. Given that $M_L''(\theta) = \frac{dH}{dt}\frac{dt}{d\theta}$, we have $M_L''(\theta) = -h(t)t^2$ where $-h(t)t^2 \sim_{t\to 0} -(1-\gamma_G)H(t)t$. Therefore, we obtain

$$M''(\theta) \sim_{\theta\to\infty} -(1-\gamma_G)\frac{M_L'(\theta)}{\theta}. \tag{B22}$$

Clearly, $\lim_{\theta\to\infty} r_M(\theta) = \lim_{\theta\to\infty}\frac{-M''(\theta)\theta}{M'(\theta)} = 1-\gamma_G$. □

Proof of Proposition 5. Part (1). If P_n is Poisson, then it satisfies Condition 2, so Proposition 1 says $\mu(\theta) = M'(\theta)/D(\theta)$. Given that $\lim_{\theta\to\infty} P_0(\theta) = 0$, we have $D(\theta) \sim_{\theta\to\infty} 1/\theta$. Therefore, $\mu(\theta) \sim_{\theta\to\infty} M'(\theta)\theta$ and (B4) from Lemma B3 yields (40).

Part (2). Starting with (40), and letting $x = G^{-1}\left(1-\frac{1}{\theta}\right)$ or $1/\theta = 1-G(x)$, we have

$$\lim_{\theta\to\infty}\mu(\theta) = \Gamma(1-\gamma_G)\lim_{x\to\bar{x}}\frac{1-G(x)}{g(x)}, \tag{B23}$$

and thus $\lim_{\theta\to\infty}\mu(\theta) = 0$ if and only if $\lim_{x\to\bar{x}}\frac{1-G(x)}{g(x)} = 0$.

Part (3). If P_n is Poisson, then it satisfies Condition 2, so Proposition 3 says that $\varepsilon_\mu(\theta) = \varepsilon_D(\theta) - r_M(\theta)$. Also, $\lim_{\theta\to\infty}\varepsilon_D(\theta) = 1$ because $D(\theta) \sim_{\theta\to\infty} 1/\theta$. Therefore, we can apply Lemma B3, which says that $\lim_{\theta\to\infty} r_M(\theta) = 1-\gamma_G$, to obtain $\lim_{\theta\to\infty}\varepsilon_\mu(\theta) = \gamma_G$. □

Proof of Proposition 6. Part (1). Suppose that $\bar{x} = \infty$. If P_n is Poisson, we have $\Delta_s(\theta) = 1-\eta_M(\theta)$ by Proposition 4. Using (B2) and (B4) yields

$$\eta_M(\theta) = \frac{-M'(\theta)\theta}{M(\theta)} \sim \frac{1}{\theta G^{-1}\left(1-\frac{1}{\theta}\right)g\left(G^{-1}\left(1-\frac{1}{\theta}\right)\right)}, \tag{B24}$$

if P_n is Poisson. Letting $x = G^{-1}\left(1-\frac{1}{\theta}\right)$ or $1/\theta = 1-G(x)$, we obtain

$$\lim_{\theta\to\infty}\eta_M(\theta) = \lim_{x\to\bar{x}}\frac{1-G(x)}{xg(x)} = \gamma_G \in [0,1). \tag{B25}$$

Proposition 4 implies $\Delta_s(\theta) = 1-\eta_M(\theta)$ if P_n is Poisson, so $\lim_{\theta\to\infty}\Delta_s(\theta) = 1-\gamma_G \in (0,1]$.

Part (2). Suppose $\bar{x} < \infty$. Given $\eta_M(\theta) = \frac{-M'(\theta)\theta}{M(\theta)}$, (B4) and $\lim_{\theta\to\infty} M(\theta) = \bar{x}$ yields

$$\lim_{\theta\to\infty}\eta_M(\theta) = \lim_{\theta\to\infty}\frac{\Gamma(1-\gamma_G)}{\theta g\left(G^{-1}\left(1-\frac{1}{\theta}\right)\right)\bar{x}}, \tag{B26}$$

if P_n is Poisson. Letting $x = G^{-1}\left(1-\frac{1}{\theta}\right)$ or $1/\theta = 1-G(x)$, we obtain

$$\lim_{\theta\to\infty}\eta_M(\theta) = \frac{\Gamma(1-\gamma_G)}{\bar{x}}\lim_{x\to\bar{x}}\frac{1-G(x)}{g(x)} = 0, \tag{B27}$$

because $\bar{x} < \infty$. So, $\lim_{\theta\to\infty}\eta_M(\theta) = 0$ and thus $\lim_{\theta\to\infty}\Delta_s(\theta) = \lim_{\theta\to\infty} 1-\eta_M(\theta) = 1$. □

References

Albrecht, J., Gautier, P., and Vroman, S. "A Note on Peters and Severinov, "Competition Among Sellers Who Offer Auctions Instead of Prices"." *Journal of Economic Theory*, Vol. 147 (2012), pp. 389–392.

Albrecht, J., Gautier, P., and Vroman, S. "Efficient Entry in Competing Auctions." *American Economic Review*, Vol. 104 (2014), pp. 3288–3296.

Allen, J., Clark, R., and Houde, J.-F. "Search Frictions and Market Power in Negotiated-Price Markets." *Journal of Political Economy*, Vol. 127 (2019), pp. 1550–1598.

Anderson, S. P., De Palma, A., and Nesterov, Y. "Oligopolistic Competition and the Optimal Provision of Products." *Econometrica*, Vol. 63 (1995), pp. 1281–1301.

Armstrong, M. and Vickers, J. "Consumer Protection and Contingent Charges." *Journal of Economic Literature*, Vol. 50 (2012), pp. 477–493.

Armstrong, M. and Vickers, J. "Patterns of Competitive Interaction." *Econometrica*, Vol. 90 (2022), pp. 153–191.

Bergemann, D., Brooks, B., and Morris, S. "Search, Information, and Prices." *Journal of Political Economy*, Vol. 129 (2021), pp. 2275–2319.

Bernard, A. B., Eaton, J., Jensen, J. B., and Kortum, S. "Plants and Productivity in International Trade." *American Economic Review*, Vol. 93 (2003), pp. 1268–1290.

Bingham, N. H., Goldie, C. M., and Teugels, J. L. *Regular Variation*. Cambridge: Cambridge University Press, 1987.

Chen, Y. and Riordan, M. H. "Price and Variety in the Spokes Model." *The Economic Journal*, Vol. 117 (2007), pp. 897–921.

Chen, Y. and Riordan, M. H. "Price-Increasing Competition." *The RAND Journal of Economics*, Vol. 39 (2008), pp. 1042–1058.

Chen, Y. and Savage, S. J. "The effects of competition on the price for cable modem internet access." *The Review of Economics and Statistics*, Vol. 93 (2011), pp. 201–217.

Coey, D., Larsen, B. J., and Platt, B. C. "Discounts and Deadlines in Consumer Search." *American Economic Review*, Vol. 110 (2020), pp. 3748–3785.

Eeckhout, J. and Kircher, P. "Sorting Versus Screening: Search Frictions and Competing Mechanisms." *Journal of Economic Theory*, Vol. 145 (2010), pp. 1354–1385.

Ellison, G. and Ellison, S. F. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica*, Vol. 77 (2009), pp. 427–452.

Gabaix, X. and Laibson, D. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *The Quarterly Journal of Economics*, Vol. 121 (2006), pp. 505–540.

Gabaix, X., Laibson, D., Li, D., Li, H., Resnick, S., and de Vries, C. G. "The Impact of Competition on Prices with Numerous Firms." *Journal of Economic Theory*, Vol. 165 (2016), pp. 1–24.

Hart, O. D. "Monopolistic Competition in the Spirit of Chamberlin: Special Results." *The Economic Journal*, Vol. 95 (1985), pp. 889–908.

Kim, K. and Kircher, P. "Efficient Competition Through Cheap Talk: The Case of Competing Auctions." *Econometrica*, Vol. 83 (2015), pp. 1849–1875.

Lester, B., Visschers, L., and Wolthoff, R. "Meeting Technologies and Optimal Trading Mechanisms in Competitive Search Markets." *Journal of Economic Theory*, Vol. 155 (2015), pp. 1–15.

Mangin, S. "A Theory of Production, Matching, and Distribution." *Journal of Economic Theory*, Vol. 172 (2017), pp. 376–409.

Perloff, J. M. and Salop, S. C. "Equilibrium with Product Differentiation." *The Review of Economic Studies*, Vol. 52 (1985), pp. 107–120.

Peters, M. and Severinov, S. "Competition among Sellers Who Offer Auctions Instead of Prices." *Journal of Economic Theory*, Vol. 75 (1997), pp. 141–179.

Platt, B. C. "Inferring Ascending Auction Participation from Observed Bidders." *International Journal of Industrial Organization*, Vol. 54 (2017), pp. 65–88.

Quint, D. "Imperfect Competition with Complements and Substitutes." *Journal of Economic Theory*, Vol. 152 (2014), pp. 266–290.

Resnick, S. *Extreme Values, Regular Variation, and Point Processes*. New York: Springer, 1987.

Rhodes, A. and Zhou, J. "Personalized Pricing and Competition." *American Economic Review*, Vol. 114 (2024), pp. 2141–2170.

Rosenthal, R. W. "A Model in Which an Increase in the Number of Sellers Leads to a Higher Price." *Econometrica*, Vol. 48 (1980), pp. 1575–1579.

Salz, T. "Intermediation and Competition in Search Markets: An Empirical Case Study." *Journal of Political Economy*, Vol. 130 (2022), pp. 1–79.

Satterthwaite, M. and Shneyerov, A. "Dynamic Matching, Two-Sided Incomplete Information, and Participation Costs: Existence and Convergence to Perfect Competition." *Econometrica*, Vol. 75 (2007), pp. 155–200.

Satterthwaite, M. A. "Consumer Information, Equilibrium Industry Price, and the Number of Sellers." *The Bell Journal of Economics*, Vol. 10 (1997), pp. 483–502.

Sattinger, M. "Value of an Additional Firm in Monopolistic Competition." *The Review of Economic Studies*, Vol. 51 (1984), pp. 321–332.

Spiegler, R. "Competition over Agents with Boundedly Rational Expectations." *Theoretical Economics*, Vol. 1 (2006), pp. 207–231.

Weyl, E. G. and Fabinger, M. "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy*, Vol. 121 (2013), pp. 528–583.