The impact of competition on prices with numerous firms [☆]

Xavier Gabaix [a,b,c], David Laibson [d,c], Deyuan Li [e], Hongyi Li [f,*],
Sidney Resnick [g], Casper G. de Vries [h,i,j]

[a] *NYU Stern, United States*
[b] *CEPR, United Kingdom*
[c] *NBER, United States*
[d] *Harvard University, United States*
[e] *Fudan University, China*
[f] *UNSW Australia, Australia*
[g] *Cornell, United States*
[h] *Erasmus University Rotterdam, Netherlands*
[i] *Tinbergen Institute, Netherlands*
[j] *Chapman University, United States*

Abstract

This paper describes a mechanism that sustains high markups, even in markets with homogenous goods and many competing firms. We show that random utility models with idiosyncratic taste shocks driven by standard noise distributions produce, in large markets, robustly high equilibrium markups that are insensitive to the degree of competition. For example, with Gaussian noise and n firms, markups are asymptotically proportional to $1/\sqrt{\ln n}$; consequently, a hundred-fold increase in n, from 10 to 1000 competing firms, only

halves the equilibrium markup. The elasticity of the markup with respect to n asymptotically equals the distribution's tail exponent from extreme value theory. Only noise distributions with very thin tails have negative asymptotic markup elasticities.

1. Introduction

This paper studies the impact of competition on prices in large markets. It focuses on random utility models in a setting of monopolistic competition, where consumer choice is influenced by firm-specific 'noise' shocks (e.g., Luce, 1959; McFadden, 1981; Anderson et al., 1992). We derive a tractable general expression for equilibrium markups in symmetric random utility models with many competing firms.[1] This expression allows us to characterize the impact of different noise distributions on competitive outcomes. We find that high mark-ups are a robust feature of such models. Specifically, random utility models with standard (thin-tailed) noise distributions produce high markups, even with homogenous goods *and* many competing firms; increased competition in large markets only weakly drives down equilibrium mark-ups.

Explicit expressions for equilibrium markups in random-utility settings have previously been derived only for some specific distributions of noise. In these special cases, equilibrium markups turn out to be either completely unresponsive or extremely responsive to competition. Consider the Perloff and Salop (1985) random utility model. If consumer noise has an exponential density or a logit (i.e., Gumbel) density, then markups converge to a strictly positive value as the number of competing firms n goes to infinity: asymptotic markups have zero elasticity with respect to n (Perloff and Salop, 1985; Anderson et al., 1992). In contrast, when noise is uniformly distributed, markups are proportional to $1/n$: markups have unit elasticity and thus decrease strongly with n (Perloff and Salop, 1985).

These special cases — exponential, logit, and uniform — are appealing for their analytic tractability rather than their realism. Relative to the Gaussian distribution, the exponential and logit cases have relatively fat tails while the uniform case has no tails. We seek to understand how prices respond to competition in the general case; in particular, for empirically realistic noise distributions.

Applying tools from Extreme Value Theory (EVT), we show that markups are asymptotically proportional to $1/\left(nf\left[F^{-1}\left(1 - 1/n\right)\right]\right)$, where F is the cumulative distribution function (CDF) of the noise and $f = F'$ is the corresponding density function. This expression is easy to compute. Further, it highlights a simple 'limit pricing' logic for the determination of equilibrium markups. Heuristically, each firm sets prices by conditioning on receiving the best random shock amongst all competing firms, then choosing a markup corresponding to the expected difference

---

[1] The restriction to the symmetric-firm case maintains tractability. This precludes us from addressing instances of asymmetries; see the discussions in Bajari and Benkard (2003) and Armstrong (2016). Consequently, our propositions can be viewed as only suggestive of what happens in the richer structural models that are most frequently used in empirical industrial organization.

between its random shock and that of the next-best firm. In more formal terms, our markup expression is asymptotically proportional (and often equal) to the expected gap between the highest draw and second highest draw in a sample of n random draws of noise. Thus for large n, markups are pinned down by the tail properties of the noise distribution.

The Gaussian case – which has relatively thin tails – is illustrative. In our setting of consumer choice, the Gaussian distribution is a natural benchmark: if a consumer receives many small, idiosyncratic influences on his preferences or beliefs, then (under appropriate assumptions) the sum of these influences produces a Gaussian-distributed 'taste' shock.[2] No closed-form solutions for equilibrium markups associated with Gaussian noise have previously been derived. We show that markups in the Gaussian case are asymptotically proportional to $1/\sqrt{\ln n}$. This implies that an increase in the number of competing firms from 10 to 1000 firms results in only a halving of the equilibrium markup. In contrast, with Cournot competition, where markups are proportional to $1/n$, such an increase would result in the markup becoming 100 times smaller. This example shows that in large markets, competition with plausible noise distributions may only exert weak pressure on prices (even in the extreme case of homogeneous goods).

Further, we argue that insensitive prices are the norm rather than the exception. Specifically, we find that the elasticity of the markup with respect to the number of firms asymptotically equals the EVT tail index of the noise distribution, an easy-to-calculate magnitude that captures the notion of tail "fatness". Using this result, we show that markups have zero asymptotic markup elasticity for a wide range of noise distributions, characterized by intermediate tail fatness. Only distributions with very thin or fat tails have asymptotic markup elasticities different from zero.

Distributions with right-bounded support (i.e., finite right tails), such as the bounded power-law, correspond to settings where consumers' per-unit valuations of goods are bounded; such settings are increasingly popular in models of trade (e.g., Arkolakis et al., 2015).[3] In our setting, with bounded consumer valuations, one might expect intense competition between homogenous firms to generate large and negative markup elasticities (see also Vives, 1985). We show that this intuition does not always hold: although asymptotic markup elasticities are indeed negative for common right-bounded distributions such as the uniform and the bounded power-law, there exist right-bounded distributions where the asymptotic markup elasticity equals zero.

For distributions with fatter-than-exponential tails (e.g., the lognormal and Pareto distributions), mark-ups paradoxically *increase* as the number of competing firms increase. While the possibility of price-increasing competition is not new to the literature (see, e.g., Weyl and Fabinger, 2013), our limit-pricing logic highlights a simple intuition for this phenomenon: with sufficiently fat tails, the expected gap between the highest and second-highest of n random draws is asymptotically increasing in n (i.e., the number of competing firms).

Importantly, our findings exhibit "detail-independence". They hold for all of the random utility models that we consider: Perloff and Salop (1985), Sattinger (1984), Hart (1985b). The Perloff–Salop, Sattinger, and Hart models differ in a host of important ways.[4] Yet, these three models

---

[2] As implied by various versions of the central limit theorem; see, for example, Feller (1971, p. 262).

[3] In these models, bounded consumer valuations are typically imposed either as a consequence of convenient assumptions (e.g., linear demand functions; see Melitz and Ottaviano, 2008) or to produce certain desiderata. For example, in Arkolakis et al. (2015), the assumption of bounded consumer valuations (leading, in their setting, to a finite choke price) ensures that marginal changes in trade costs have continuous effects on prices and welfare, thus enabling tractable analysis.

[4] For instance, in the Perloff–Salop model, consumers need to buy one unit of the good. In the Sattinger model, they allocate a fixed dollar amount to the good. The Hart model does not impose either constraint.

produce (asymptotically) the same equilibrium markup up to a scaling constant, for a wide range of different noise distributions. Such detail-independence permits a more robust analysis than would be possible if results depended on the specific properties of the demand specification.

In addition to an understanding of the economics of random utility models, the tools that we develop allow us to calculate the large-n asymptotic behavior of integrals for a class of functions $h(x)$, of the form

$$\int h(x) f^k(x) F^n(x) dx, \; k \geq 1. \tag{1}$$

This integral can be used to calculate the expected value of a function of the maximum of n random variables, or the gap between the maximum and the second largest value of those random variables. Using EVT, we derive robust approximations of this integral for large n.

These mathematical results have broad applications to various economic settings related to market and auction mechanisms in large economies.[5] In particular, Mangin (2015a) analyzes an elegant model of frictional labor markets where firms compete via auction to hire workers, and points out that the present paper's results may be applied to calculate the asymptotic value of key economic quantities such as the income share of labor. Using a closely related framework, Mangin (2015b) points out that Theorem 3 of the present paper may be applied to derive aggregate production functions from an underlying productivity distribution.[6]

Theoretical results quantifying the effect of competition on markups in large markets have so far been limited. Some papers derive explicit markup expressions under specific functional form assumptions (e.g., Dixit and Stiglitz, 1977; Perloff and Salop, 1985). One exception is Vives (1985), who derives asymptotic bounds on the markup function under Bertrand and Cournot competition given bounded consumer valuations and "sufficiently substitutable" competing products; this corresponds to the case of right-bounded distributions in our model. We briefly discuss the comparison with Vives (1985) in Section 2.3.

A related theoretical literature studies qualitative features of the effect of competition on markups; the key insights are laid out in Vives (2001, Section 6.4). If competing products become infinitely substitutable as the number of competing firms grows large (as in Vives, 1985), then markups asymptotically approach zero. Indeed, given bounded noise distributions in our random-demand setting, each product will (in expectation) have "nearby" competitors with similar taste shock realizations as n grows large, so markups vanish asymptotically (see also Perloff and Salop, 1985). Similarly, Mas-Colell (1975) argues that if consumers have continuous preferences over a compact product space, then the competitive outcome obtains at the large-n limit. The converse point – that markups do not vanish asymptotically if products remain imperfectly substitutable even at the large-n asymptotic limit – has also been made; see, e.g., Hart (1985a). This latter point is analogous to our finding that for sufficiently fat-tailed distributions, markups

---

[5] More generally, EVT techniques are important in many areas of economics, such as industrial organization and discrete choice (e.g., Luce, 1959; McFadden, 1981; Anderson et al., 1992; Dagsvik, 1994; Bulow and Klemperer, 2002, 2012; Dagsvik and Karlstrom, 2005; Ibragimov and Walden, 2010; Weyl and Fabinger, 2013, and Armstrong, 2016), international trade (e.g., Eaton and Kortum, 2002; Bernard et al., 2003, and Chaney, 2008, 2015), macroeconomics and growth (e.g., Gabaix, 1999, 2011; Jones, 2005; Luttmer, 2007, and Acemoglu et al., 2012), systemic risk analysis (e.g., Jansen and de Vries, 1991 and Ibragimov et al., 2009, 2011) and auction theory (e.g., Hong and Shum, 2004).

[6] In a separate application, Gabaix and Landier (2008) use some of our results to analyze the upper tail of the distribution of CEO talents.

do not vanish because the expected gap between the largest and second-largest noise shocks remains bounded away from zero.

Perhaps closest to our paper – and in particular to our finding that markups decrease (increase) with competition for thin-tailed (fat-tailed) distributions – are Weyl and Fabinger (2013) and Quint (2014), who show how comparative statics of pricing behavior hinge crucially on log-concavity of the demand function; relating this insight to our results, Weyl and Fabinger (2013) point out that competition increases (decreases) markups if the distribution of consumer valuations is log-convex (log-concave).[7] These papers precisely demarcate the boundary between price-increasing and price-decreasing competition. Complementing these papers, we quantify the impact of competition on prices for general distributions. In fact, we show that for a wide range of distributions above and below the aforementioned boundary, prices are relatively insensitive to competition.

Our results also relate to some stylized facts about competition and markups. A number of papers document high markups in industries with homogenous goods and many competing firms. Hortaçsu and Syverson (2004) document high mark-ups in the mutual fund market, representing more than 1% of assets under management for most asset classes, even in asset classes with hundreds of competing funds. Ausubel (1991) and Stango (2000) show that interest rates on credit cards have been much greater than the cost of funds, despite the presence of hundreds of competing card-issuing banks.[8] Our model provides some relevant conceptual insights for interpreting these facts. A more rigorous theoretical investigation of these industries would require a model with ex-ante heterogenous firms to allow for dispersed markups, something that our symmetric-firm model does not address.

An empirical literature studies the rate at which markups change with entry. Bresnahan and Reiss (1991) show, in a study of US firms, that oligopolistic markups decrease rapidly with n when n is small (between 1 and 3), but level off (above zero) beyond three competitors.[9] While Bresnahan and Reiss' (1991) sample is limited to small markets (typically with no more than ten competitors) and thus does not overlap with the large-n focus of our paper, their results are broadly consistent with our conceptual point that for sufficiently large markets, under a wide range of circumstances, markups are relatively insensitive to increased competition. On the other hand, Campbell and Hopenhayn (2005) show that establishment size is positively correlated with market size in large retail markets, and suggest that this correlation arises because markups decrease with market size (so that firms in large markets have to exploit economies of scale to overcome fixed costs).

The paper proceeds as follows. Section 2 presents the main economic result using the random utility model of Perloff and Salop (1985), demonstrates the equivalence of our results to the limit-pricing and auction settings, and discusses welfare implications. Section 3 considers

---

[7] Relatedly, a number of other papers focus on the point that prices may rise with more intense competition: see, e.g., Chen and Riordan (2008), as well as Rosenthal (1980), Bénabou and Gertner (1993), Bulow and Klemperer (2002), Carlin (2009), and Zhelobodko et al. (2012) for perverse competitive effects generated by different microfoundations.

[8] Further examples of high markups in financial products abound. Bergstresser et al. (2009) find that mutual funds sold by brokers have anomalously high fees and low net-of-fee returns. Henderson and Pearson (2011) find that structured equity products also have robustly high mark-ups, and hypothesize that this is related to investor confusion about product quality. Another complementary explanation is that investors like the psychological comfort given by specific mutual fund brokers (Gennaioli et al., 2015).

[9] Another result in this vein is from Mazzeo (2002), who finds in small, local motel markets that the marginal effect of additional competition on markups decreases quickly with the number of competitors.

alternative random utility models (Sattinger, 1984; and Hart, 1985b), and shows that the details of the demand-side modeling matter little, or not at all, to markups. Section 4 presents the main mathematical result: an asymptotic approximation of a key integral that is needed to characterize economic environments in which extremes matter. We show that the tail of the noise distribution – captured by the tail index – is the crucial determinant of prices. As many common noise distributions have a tail index of zero, our results imply that in a wide range of market contexts additional competition has little effect on prices, once the market goes beyond a small number of firms. Section 5 concludes.

We prove our main results (including Theorems 1 and 3) in Appendix A, and our other results in an online appendix.

2. How much does competition affect prices?

In this section, we describe the random utility model from Perloff and Salop (1985). Postponing some of the mathematical elements of the proof (which are provided in Section 4), we report our key result: an asymptotic expression for price markups under oligopolistic competition. We then discuss implications and applications.

2.1. The Perloff–Salop model

There is a single representative consumer and an exogenously specified number of firms, n. The consumer seeks to purchase exactly one unit of the good from one firm. He perceives that he will receive net utility $U_i = X_i - p_i$ by purchasing the good of firm i, where X_i is a noise term representing a random taste shock, i.i.d. across firms and consumers, and p_i is the price charged by firm i. Thus the consumer chooses to purchase the good that maximizes $X_i - p_i$. Each firm can produce at marginal cost c. The timing is as follows:

1. Firms simultaneously set prices;
2. Random taste shocks are realized;
3. Consumers make purchase decisions;
4. Each firm produces the amount purchased from that firm;
5. Profits are realized.

The key economic object of interest is the price markup in a symmetric equilibrium, which we derive by solving the first-order condition for each firm's profit maximization problem. Firm i's profit function is given by

$$\pi_i = (p_i - c)D(p_1, \ldots, p_n; i) \tag{2}$$

where $D(p_1, \ldots, p_n; i)$ is the demand function for firm i given the price vector $(p_1, \ldots, p_n)$ of the n goods. The first order condition for profit maximization implies the following equilibrium markup in a symmetric equilibrium

$$p - c = -\frac{D(p, p; n)}{D_1(p, p; n)}. \tag{3}$$

Here p is the symmetric equilibrium price, $D(p, p'; n)$ denotes the demand function for a firm that sets price p when there are n goods and all other firms set price p', and $D_1(p, p'; n) \equiv \partial D(p, p'; n)/\partial p$. Denote the markup $p - c$ in a symmetric equilibrium with n firms as μ_n.

In a symmetric-price equilibrium, the demand function of firm i is the probability that the consumer's surplus at firm i, $X_i - p_i$, exceeds the consumer's surplus at all other firms,

$$D(p_1, \ldots, p_n; i) = \mathbb{P}(X_i - p_i \geq \max_{j \neq i} \{X_j - p_j\}) = \mathbb{P}(X_i \geq \max_{j \neq i} \{X_j\}). \tag{4}$$

Let M_n denote $\max\{X_1, \ldots, X_n\}$, which has density $nf(x)F^{n-1}(x)$.[10] Evaluation of (3) gives the following markup expression for the symmetric equilibrium of the Perloff–Salop model:

$$\mu_n = \frac{1}{n\mathbb{E}[f(M_{n-1})]} = \frac{1}{n(n-1)\int f^2(x)F^{n-2}(x)\,dx}. \tag{5}$$

Here F is the distribution function and f is the corresponding density of X_i.

Before proceeding to our analysis of the markup expression (5), let us briefly discuss our modeling approach. We use a stripped-down model of random utility for our analysis. In the model, the consumer's payoff function takes an additive form. We show in Section 3 that our results do not rely on this specification. There, we analyze two other random utility models which feature (as in Perloff and Salop, 1985) a representative consumer who has random i.i.d. taste shocks over producers, but differ in the form of consumer preferences. Our results from the present section are preserved in these alternative models, suggesting that the impact of competition on markups is independent of many of the institutional details of competition in random-utility settings.

A second feature of our model is that firms are completely symmetric ex ante, and thus each firm receives an equal $1/n$ expected market share in equilibrium. This assumption is strong, but enables tractable analysis.

2.2. Extreme value theory: some basics

We very briefly introduce some necessary machinery, and postpone some of the mathematical details to Section 4. As in Section 2.1, define $M_n \equiv \max_{i=1,\ldots,n} X_i$ to be the maximum of n independent random variables X_i with distribution F, and define the counter-cumulative distribution function $\overline{F}(x) \equiv 1 - F(x)$. We are particularly interested in the connection between M_n and $\overline{F}^{-1}(1/n)$[11]; informally (in analogy with the empirical distribution function), one may think of $\overline{F}^{-1}(1/n)$ as the "typical" value of M_n. In fact, the key to our analysis is to formalize this relationship between $\overline{F}^{-1}(1/n)$ and M_n for large n.

Our analysis is restricted to what we call *well-behaved* distributions:

Definition 1. Let F be a distribution function with support on (w_l, w_u), where $w_u \leq \infty$. We say F is well-behaved iff $f = F'$ is differentiable in some neighborhood of w_u, $\lim_{x \to w_u} \overline{F}(x)/f(x) = a$ exists with $a \in [0, \infty]$, and

$$\gamma = \lim_{x \to w_u} \frac{d}{dx}\left(\frac{\overline{F}(x)}{f(x)}\right) \tag{6}$$

exists and is finite. We call γ the tail index of F.

---

[10] Indeed, $\mathbb{P}(M_n \leq x) = \mathbb{P}(X_i \leq x \text{ for } i = 1\ldots n) = \{\mathbb{P}(X_i \leq x)\}^n = F(x)^n$.

[11] Strictly speaking, we abuse notation in cases where F is not strictly increasing by using $\overline{F}^{-1}(t)$ to denote $\overline{F}^{\leftarrow}(t) = F^{\leftarrow}(1-t)$, where $F^{\leftarrow}(t) = \inf\{x \in (w_l, w_u) : F(x) \geq t\}$ is the generalized inverse of F (Embrechts et al. 1997, p. 130). This is for expositional convenience: our results hold with the generalized inverse as well.

Table 1
Asymptotic expressions for markups.

Name of distribution	f	γ	μ_n	$\lim_{n\to\infty}\mu_n$
Uniform	$1,\quad x\in[-1,0]$	-1	$1/n$	0
Bounded Power Law	$\alpha(-x)^{\alpha-1}$ $\alpha\geq 1, x\in[-1,0]$	$-1/\alpha$	$\frac{\Gamma(1-1/\alpha+n)}{\alpha\Gamma(2-1/\alpha)\Gamma(1+n)}\sim\frac{n^{-1/\alpha}}{\alpha\Gamma(2-1/\alpha)}$	0
Weibull	$\alpha(-x)^{\alpha-1}e^{-(-x)^{\alpha}}$ $\alpha\geq 1, x<0$	$-1/\alpha$	$\frac{1}{\alpha\Gamma(2-1/\alpha)}\frac{n^{1-1/\alpha}}{n-1}\sim\frac{n^{-1/\alpha}}{\alpha\Gamma(2-1/\alpha)}$	0
Bounded Exponential-like	$\frac{e^{-x/(1-x)}}{(1-x)^2}$ $x\in[0,1]$	0	$\sim\frac{1}{(\ln n)^2}$	0
Gaussian	$(2\pi)^{-1/2}e^{-x^2/2}$	0	$\sim(2\ln n)^{-1/2}$	0
Rootzen class, $\phi>1$	$\kappa\lambda\phi x^{a+\phi-1}e^{-x^{\phi}}$	0	$\sim\frac{1}{\phi\lambda^{1/\phi}}(\ln n)^{1/\phi-1}$	0
Gumbel	$\exp(-e^{-x}-x)$	0	$\frac{n}{n-1}$	1
Exponential	$e^{-x},\quad x>0$	0	1	1
Rootzen Gamma	$\tau x^{\tau-1}e^{-x^{\tau}}$ $x>0, \tau<1$	0	$\sim\frac{1}{\tau}(\ln n)^{1/\tau-1}$	∞
Lognormal	$\frac{\exp(-2^{-1}\log^2 x)}{x\sqrt{2\pi}}$ $x>0$	0	$\sim\frac{1}{\sqrt{2\ln n}}e^{\sqrt{2\ln n}}$	∞
Pareto	$\alpha x^{-\alpha-1}$ $\alpha>1, x\geq 1$	$1/\alpha$	$\frac{\Gamma(1+1/\alpha+n)}{\alpha\Gamma(2+1/\alpha)\Gamma(1+n)}\sim\frac{n^{1/\alpha}}{\alpha\Gamma(2+1/\alpha)}$	∞
Fréchet	$\alpha x^{-\alpha-1}e^{-x^{-\alpha}}$ $\alpha>1, x\geq 0$	$1/\alpha$	$\frac{1}{\alpha\Gamma(2+1/\alpha)}\frac{n^{1+1/\alpha}}{n-1}\sim\frac{n^{1/\alpha}}{\alpha\Gamma(2+1/\alpha)}$	∞

This table lists asymptotic markups (under symmetric equilibrium) for the Perloff–Salop model for various noise distributions as a function of the number of firms n. f specifies the density function, and γ specifies the distribution's tail index. Distributions are listed in order of increasing tail fatness. Asymptotic approximations are calculated using Theorem 1 except where the markup can be exactly evaluated.

The tail index γ measures the fatness of F's right tail. The case $\gamma<0$ consists of very thin-tailed distributions such as the uniform distribution. The case $\gamma=0$ consists of distributions with tails of intermediate thickness. A wide range of economically interesting distributions fall within this domain, ranging from the relatively thin-tailed Gaussian distribution to the relatively fat-tailed lognormal distribution, as well as other distributions in between, such as the exponential distribution. The case $\gamma>0$ consists of fat-tailed, Pareto-like distributions such as the Pareto and the Fréchet distributions.

Being well-behaved in the sense of Definition 1 is not a particularly strong restriction. It is satisfied by most distributions of interest, and is easy to verify.[12] In Section 2.3, Table 1 lists a number of popular densities and the corresponding tail index γ. Note that distributions with an exponential-like right tail all have $\gamma=0$.

To ensure that the quantities that we are calculating do not diverge, we also impose some restrictions on the rest of F.

Definition 2. Let $j:\mathbb{R}\to\mathbb{R}$ have support on (w_l, w_u). The function $j(x)$ is (w_l, w_u)-integrable iff $\int_{w_l}^{w}|j(x)|\,dx<\infty$ for all $w\in(w_l, w_u)$.

---

[12] Condition (6) is well-known in the EVT literature as a von Mises condition.

For example, in Theorem 1 we require that f^2 be (w_l, w_u)-integrable. Verification of this condition is typically straightforward; for example, $f^2(x)$ is (w_l, w_u)-integrable if $f = F'$ is uniformly bounded. Throughout the paper, we use $w_l = \inf\{x : F(x) > 0\}$ and $w_u = \sup\{x : F(x) < 1\}$ to denote, respectively, the lower and upper bounds of the support of the noise distribution F.

2.3. How do markups change with competition?

The next theorem is our key result: it characterizes, asymptotically, the equilibrium markup as a function of the noise distribution and the number of competing firms.[13] Assume that F is well-behaved, $f^2(x)$ is (w_l, w_u)-integrable, and the tail index satisfies $-1.45 \leq \gamma \leq 0.64$.[14] Adopting standard notation, we write $a_n \sim_{n \to \infty} b_n$ (or simply $a_n \sim b_n$) if and only if $\lim_{n \to \infty} a_n/b_n = 1$.

Theorem 1. *The symmetric equilibrium markup in the Perloff–Salop model is, asymptotically (for $n \to \infty$),*

$$\mu_n \sim \frac{1}{nf\left(\overline{F}^{-1}\left(\frac{1}{n}\right)\right)\Gamma(\gamma + 2)},$$

where $\overline{F}(x) \equiv 1 - F(x)$, and $\Gamma(t) \equiv \int_0^\infty y^{t-1}e^{-y}dy$ is the standard Gamma function.

Theorem 1 allows us to calculate the equilibrium markup for various noise distributions. Markups for some common distributions are listed in Table 1, where distributions are presented in increasing order of tail fatness. At one extreme, the uniform distribution has the thinnest tails ($\gamma = -1$). The uniform distribution produces an asymptotic markup that is proportional to $1/n$. In this case, competition dramatically reduces markups. Going down the table, as tail fatness increases, markups become less sensitive to n. This cursory inspection suggests a tight connection between the tail fatness and the impact of competition on markups. We state this point more precisely in the following proposition. It shows that the tail index γ in (6) has a concrete economic implication: γ is the asymptotic elasticity of the markup with respect to the number of firms. In other words, interpreting n as a continuous variable, the markup behaves locally as $\mu \sim kn^\gamma$. Assume that the conditions in Theorem 1 hold, and further that $f^2(x)\log F(x)$ is (w_l, w_u)-integrable.

Proposition 1. *The asymptotic elasticity of the Perloff–Salop markup with respect to the number of firms n is*

$$\lim_{n \to \infty} \frac{n}{\mu_n}\frac{d\mu_n}{dn} = \gamma.$$

The case $\gamma < 0$ consists of distributions with very thin tails. For these distributions, asymptotic elasticity is strictly negative: markups are sensitive to the number of firms n, even as $n \to \infty$.

---

[13] The proof relies on Theorem 3, proven later; for expositional convenience, we start with the main economic results.

[14] This is the range over which the second order condition holds (see the online appendix for details); the first order condition holds whenever $\gamma > -2$. Note that this assumption on γ is not very restrictive. It permits thin-tailed distributions such as the Weibull, and all (fat-tailed) Pareto distributions with finite variance.

The extreme case is the uniform distribution, with asymptotic elasticity $\gamma = -1$.[15] The bounded power law and Weibull distributions are particularly flexible, and allow us to span the gamut of negative asymptotic elasticities within $(-1, 0)$. One common feature of these distributions is that they have right-bounded support, so that consumers' valuations of each good are bounded above by $w_u < \infty$.

While all distributions with $\gamma < 0$ have right-bounded support (e.g., uniform, bounded power-law, and Weibull), the converse is not true: there exist right-bounded distributions with zero asymptotic elasticity. Table 1 includes a bounded exponential-like distribution (see Gnedenko, 1943; and Resnick, 1987, p. 39) where markups decrease very slowly for large n: specifically, as $1/(\log n)^2$. In other words, prices may be insensitive to additional competition in large markets, even in settings with bounded consumer valuations. These results are somewhat in contrast with those of Vives (1985), who shows that in a differentiated-product setting with bounded consumer valuations and additional assumptions about the substitutability of competing goods, asymptotic markup elasticities are large and negative (≤ -1) in the sense that markups go to zero at a rate of at least $1/n$.[16] In our model, such elasticities are only achieved in the extreme thin-tailed case of the uniform distribution.

At the other end of Table 1, for distributions with fatter tails than the exponential distribution, markups *rise* as the number of competitors *increases*. This set includes the fat-tailed, Pareto-like distributions, which have $\gamma > 0$; for these distributions, the asymptotic elasticity is strictly positive, so markups increase quite dramatically with n. It also includes some relatively fat-tailed distributions with $\gamma = 0$, such as the lognormal. Intuitively, for sufficiently fat-tailed noise, as n increases, the difference between the best draw and the second-best draw, which is proportional to $nf\left(\overline{F}^{-1}(1/n)\right)$, increases with n (see also Section 2.4 below).[17]

Returning to the middle of Table 1, a wide range of distributions – those with intermediate tail fatness, $\gamma = 0$ – have zero asymptotic elasticity. This range encompasses cases of price-decreasing competition (e.g., Bounded Exponential-like, and Gaussian) as well as price-increasing competition (e.g., lognormal). It encompasses cases of bounded and unbounded support. Within this range, Proposition 1 tells us that increased competition has remarkably little effect on markups in large markets.

2.3.1. Markup sensitivity: a numerical example

Consider the (relatively thin-tailed) case of Gaussian noise. In this case, the markup μ_n is proportional to $1/\sqrt{\ln n}$. Accordingly, μ_n converges to zero, but this convergence proceeds at a glacial pace.

To conceptually illustrate this slow convergence, we calculate μ_n when noise is Gaussian for a series of values of n. Table 2 shows that with Gaussian noise, a highly competitive industry with

---

[15] Distributions with tail index $\gamma < -1$ may be considered "thinner" than the uniform and have even larger negative asymptotic elasticities. However, they have the unattractive and unnatural feature that their density diverges at the (finite) upper bound of their support.

[16] Relatedly, Mas-Colell (1975) shows in a setting where the space of product characteristics is compact that markups vanish as $n \to \infty$ (see also Vives, 2001, Section 6.4).

[17] This logic also explains why the exponential distribution, and asymptotic tail equivalents such as the Gumbel distribution, demarcate the threshold between price-decreasing and price-increasing competition. For such distributions, the expected difference between the best draw and the second-best draw is asymptotically constant in n.

Table 2
Markups with Gaussian noise and uniform noise.

n	Markup with Gaussian noise	Markup under Cournot Competition
10	1	1
100	0.61	0.1
1,000	0.47	0.01
10,000	0.40	0.001
100,000	0.35	0.0001
1,000,000	0.32	0.00001

Markups are calculated for (i) the symmetric equilibrium of the Perloff–Salop model for Gaussian noise and (ii) under Cournot Competition, for various values of the number of firms n, where n is the number of firms in the market. Markups are normalized to equal one at $n = 10$.

$n = 1,000,000$ firms (far more than in any realistic setting) will nonetheless retain a third of the markup of a relatively concentrated industry with only $n = 10$ competitors. We also compare markups in our model to those in the Cournot model, which features markups proportional to $1/n$ and a markup elasticity w.r.t. n of -1 (note that this is equal to markups in the Perloff–Salop model with uniformly distributed noise).

More generally, in cases with moderate tail fatness, such as the Gumbel (i.e., logit), exponential, and lognormal densities, the markup again shows little (zero asymptotic) response to changes in n. Nevertheless, the markups become unbounded for the lognormal distribution. Finally, the bounded exponential-like distribution shows that an infinite support is not necessary for our results. In this case the markup is asymptotically proportional to $1/(\log n)^2$ and markup decay remains slow. In concrete terms, in markets with noisy demand and many competitors, one should not assume that increased competition dramatically reduces markups.

2.4. Limit pricing: an alternative interpretation

We now discuss an alternative model of oligopolistic competition, sometimes called "limit pricing", which has proven to be very useful in trade and macroeconomics (e.g., Bernard et al., 2003; see also Auer and Chaney, 2009). The price-setting mechanism in the limit pricing model is remarkably simple, yet it produces (asymptotically) the same markups as the Perloff and Salop (1985) model. This equivalence result implies that a similar logic underlies the equilibrium markups for these models, and thus generates a simple but useful interpretation of our economic results.

In the limit pricing model, each firm i draws an i.i.d. quality shock X_i. Firms simultaneously set prices p_i *after* observing the other firms' quality shocks. (This is in contrast with the Perloff and Salop (1985) model, where prices are set *before* taste shocks are observed.) The representative consumer purchases one unit of the good, and picks the firm which maximizes $X_i - p_i$. As before, call $M_n = \max_{i=1...n} X_i$ the largest quality draw from the n firms, and S_n the second-largest draw. In the competitive equilibrium, the firm with the highest quality, M_n, captures the entire market, and sets a markup of $\mu_n^{LP} = M_n - S_n$. This is just enough to take all the market away from the firm with the second-highest quality.

The next proposition analyzes the equilibrium markup under Limit Pricing. We assume that F is well-behaved with tail index $\gamma < 1$, and that $\mathbb{E}[X_i] < \infty$.

Proposition 2. *Let M_n and S_n be, respectively, the largest and second largest realizations of n i.i.d. random variables with CDF F. Then limit pricing markup is $\mu_n^{LP} = M_n - S_n$, and*

$$\mathbb{E}\left[\mu_n^{LP}\right] \sim_{n\to\infty} \frac{\Gamma(1-\gamma)}{nf\left(\overline{F}^{-1}\left(\frac{1}{n}\right)\right)}. \tag{7}$$

Notice that this markup is asymptotically proportional to the markup from Theorem 1. This suggests the following intuition for Theorem 1: to set its optimum price, a firm conditions on getting the largest draw, then evaluates the likely draw of the second highest firm and engages in limit pricing, where it charges a markup equal to the difference between its draw and the next highest draw: $\mathbb{E}\left[\mu_n^{LP}\right] \approx M_n - S_n$. (This is analogous to the analysis of a first-price sealed-bid auction.) In fact, this reasoning gets us approximately the correct answer: observe that $\mathbb{E}\left[\overline{F}(M_n)\right] \simeq \frac{1}{n+1}$ and $\mathbb{E}\left[\overline{F}(S_n)\right] \simeq \frac{2}{n+1}$, which suggests that M_n (the highest draw) will be close to $\overline{F}^{-1}\left(\frac{1}{n}\right)$ and that S_n (the second-highest draw) will be close to $\overline{F}^{-1}\left(\frac{2}{n}\right)$. So,

$$\mathbb{E}\left[\mu_n^{LP}\right] \approx M_n - S_n \approx \overline{F}^{-1}(1/n) - \overline{F}^{-1}(1/n + 1/n)$$

$$\approx -\frac{1}{n} \cdot \left(\overline{F}^{-1}\right)'(1/n) \quad \text{by Taylor expansion}$$

$$= \frac{1}{nf\left(\overline{F}^{-1}(1/n)\right)}.$$

In fact, revisiting Theorem 1, we see that this heuristic argument generates the right approximation for the Perloff–Salop markups when $\gamma = 0$ (e.g. Gaussian, logit (Gumbel), exponential, and lognormal distributions), and that the approximation remains accurate up to a corrective constant for the other distributions.

2.4.1. An application to auctions

Our mathematical results can also be applied to the analysis of auctions. Consider a second-price auction with a single good and n bidders where each bidder i privately values the good at X_i, which is i.i.d. with CDF F. It is well-known that if F is strictly increasing on (w_l, w_u), then in equilibrium each bidder bids his private valuation; the bidder with the highest valuation M_n wins and pays the second-highest valuation S_n. Proposition 2 then immediately implies that the expected surplus for the winner of the auction is[18]

$$\mathbb{E}[M_n - S_n] \sim_{n\to\infty} \frac{\Gamma(1-\gamma)}{nf\left(\overline{F}^{-1}(1/n)\right)}. \tag{8}$$

Other key quantities are also easily derived. For example, the seller's expected revenue is

$$\mathbb{E}[S_n] \sim_{n\to\infty} \overline{F}^{-1}(1/n)\,\Gamma(2-\gamma) \quad \text{if } w_u = \infty. \tag{9}$$

For some applications of these results to large auction settings, see, e.g., Mangin (2015a).

---

[18] The case $\gamma \neq 0$ of result (8) appeared in Caserta (2002, Prop. 4.1). Her proof relied on a different argument.

2.5. Consumer surplus

The random utility framework is sometimes criticized for generating an unrealistically high value for consumer surplus and social surplus. Indeed, if the noise distribution is unbounded (and the noise is treated as normatively meaningful taste shocks), then total consumer surplus tends to ∞ as the number of firms increases. Our analytical results allow us to examine this criticism.

To perform welfare analysis in the Perloff and Salop (1985) model, in this subsection we interpret taste shocks as capturing preference heterogeneity among consumers. So our measure of consumer surplus is simply $X_i - p_i$ where X_i is the consumer taste shock and p_i is the price for the purchased good. In this setting, expected consumer surplus is $\mathbb{E}[M_n] - p$ and expected social surplus is $\mathbb{E}[M_n] - c$, where p is the equilibrium price and $M_n \equiv \max_{i=1,\dots,n} X_i$ is the largest quality draw from n firms. For brevity, we restrict ourselves to the case with unbounded distributions and $\gamma \geq 0$.

We can show that $\mathbb{E}[M_n] \sim \Gamma(1-\gamma)\overline{F}^{-1}(1/n)$ for $0 \leq \gamma < 1$.[19] For all the distributions that we study except the Pareto-like cases, $\overline{F}^{-1}(1/n)$ rises only slowly with n. Hence, even for unbounded distributions, and large numbers of producers, surplus can be quite small. For example, for the case of Gaussian noise when consumer preferences have a standard deviation of \$1, $\overline{F}^{-1}(1/n) \sim \sqrt{2\ln n}$; with a million toothpaste producers, consumer surplus averages no more than \$5.25 per tube. Hence, in many instances, the framework — even with unbounded distributions — does not generate counterfactual predictions about surplus or counterfactual predictions about the prices that cartels would set.

3. Detail-independence

This section demonstrates the robustness of our main findings from Section 2 to alternative assumptions about consumer preferences.

3.1. Alternative models

We briefly describe two alternative random-utility models and show that these models also obey the asymptotic markup rule of Theorem 1. These models differ from Perloff and Salop (1985) in the specification of consumer preferences, but otherwise share common features: there is a single representative consumer and n firms, indexed as $i = \{1, \dots, n\}$. The timing is also the same: firms set prices simultaneously, before taste shocks are realized. As before, c is the marginal cost of production, p_i is the price of good i, and the random shocks X_i associated with each good i are i.i.d. randomly distributed with distribution function F.

Sattinger (1984) analyzes the case of multiplicative random utility, where consumers demand a fixed dollar amount. There are two types of goods. Besides the monopolistically competitive market, there is a composite good purchased from an industry with homogenous output. Our focus is on markups in the monopolistically competitive market. The consumer has utility function

$$U = Z^{1-\theta}\left[\sum_{i=1}^{n} A_i Q_i\right]^{\theta}, \tag{10}$$

---

[19] See Resnick (1987, p. 77, Proposition 2.1). Alternatively, this result is an immediate application of Theorem 3, which we present in Section 4.

where Z is the quantity of the composite good, $A_i = e^{X_i}$ is the random taste shock, and Q_i is the quantity consumed of good i. The consumer faces the budget constraint $y = qZ + \sum_i p_i Q_i$, where y is the consumer's endowment and q is the price of the composite good. In the online appendix, we show that the equilibrium markup in this model is

$$\frac{\mu_n^{Satt}}{c} = \frac{1}{n(n-1)\int f^2(x)F^{n-2}(x)\,dx} = \frac{1}{n\mathbb{E}\left[f(M_{n-1})\right]}. \tag{11}$$

Hart (1985b) analyzes a richer setup where consumers' demand is flexible in quantity and value. In comparison, in the Perloff–Salop model, the quantity demanded is fixed; whereas in the Sattinger model, expenditure is fixed. The consumer's utility function is:

$$U^{Hart} = \frac{\psi+1}{\psi}\left(\sum_{i=1}^n A_i Q_i\right)^{\psi/(\psi+1)} - \sum_{i=1}^n p_i Q_i, \tag{12}$$

where $A_i = e^{X_i}$ is the associated random taste shock for good i and Q_i is the quantity consumed. The equilibrium markup of the Hart (1985b) model is[20]

$$\frac{\mu_n^{Hart}}{c} = \frac{1}{\psi + (n-1)\frac{\int e^{\psi x}f^2(x)F^{n-2}(x)dx}{\int e^{\psi x}f(x)F^{n-1}(x)dx}} = \frac{1}{\psi + (n-1)\frac{\mathbb{E}\left[e^{\psi M_{n-1}}f(M_{n-1})\right]}{\mathbb{E}\left[e^{\psi M_{n-1}}\right]}}. \tag{13}$$

3.2. Comparing equilibrium markups

We now characterize equilibrium markups for the Sattinger (1984) and Hart (1985b) models. As in Theorem 1, we assume that F is well-behaved, and that $f^2(x)$ is $[w_l, w_u]$-integrable. For the Sattinger model, assume that $-1.45 \le \gamma \le 0.64$.[21] For the Hart model with parameter ψ, assume that $-1 < \gamma \le 0$; if $\gamma = 0$, we further require that $1 - \psi a > 0$.[22] Denote the Perloff–Salop, Sattinger and Hart markups as, respectively, μ_n, μ_n^{Satt}, and μ_n^{Hart}. The following theorem states that (up to the marginal cost factor c) all three markups are asymptotically equal; in fact, the Sattinger markup is exactly equal to the Perloff–Salop markup.

Theorem 2. *The symmetric equilibrium markups in the Perloff–Salop, Sattinger and Hart models are asymptotically*

$$\mu_n = \mu_n^{Satt}/c \sim \mu_n^{Hart}/c \sim \frac{1}{nf\left(\overline{F}^{-1}\left(\frac{1}{n}\right)\right)\Gamma(\gamma+2)}, \tag{14}$$

with $\overline{F}(x) \equiv 1 - F(x)$.

---

[20] Note that in the special case $\psi = 0$, by comparing (11) with (13), we see that $\mu_n^{Hart} = \mu_n^{Satt}$; that is, the Hart model generates the same demand functions and markups as the Sattinger model.

[21] As with the Perloff and Salop (1985) model, this is the range over which the second order condition holds (see the online appendix for details).

[22] For distributions violating this condition, no symmetric price equilibrium can be calculated in the Hart model because each firm would face infinite demand.

Table 3
Asymptotic expressions for Sattinger and Hart markups.

Distribution	f	$\mu_n = \mu_n^{Satt}/c$	μ_n^{Hart}/c	$\lim_{n\to\infty}\mu_n$
Uniform	$1, \quad x \in [-1, 0]$	$1/n$	$\sim 1/n$	0
Bounded Power Law	$\alpha\,(-x)^{\alpha-1}$ $\alpha \geq 1, x \in [-1, 0]$	$\frac{\Gamma(1-1/\alpha+n)}{\alpha\Gamma(2-1/\alpha)\Gamma(1+n)} \sim \frac{n^{-1/\alpha}}{\alpha\Gamma(2-1/\alpha)}$	$\sim \frac{n^{-1/\alpha}}{\alpha\Gamma(2-1/\alpha)}$	0
Weibull	$\alpha\,(-x)^{\alpha-1}e^{-(-x)^\alpha}$ $\alpha \geq 1, x < 0$	$\frac{1}{\alpha\Gamma(2-1/\alpha)}\frac{n^{1-1/\alpha}}{n-1} \sim \frac{n^{-1/\alpha}}{\alpha\Gamma(2-1/\alpha)}$	$\sim \frac{n^{-1/\alpha}}{\alpha\Gamma(2-1/\alpha)}$	0
Bounded Exponential-like	$\frac{e^{-x/(1-x)}}{(1-x)^2}$ $x \in [0, 1]$	$\sim \frac{1}{(\ln n)^2}$		0
Gaussian	$(2\pi)^{-1/2}e^{-x^2/2}$	$\sim (2\ln n)^{-1/2}$		0
Rootzen class, $\phi > 1$	$\kappa\lambda\phi x^{a+\phi-1}e^{-x^\phi}$	$\sim \frac{1}{\phi\lambda^{1/\phi}}(\ln n)^{1/\phi-1}$		0
Gumbel	$\exp(-e^{-x} - x)$	$\frac{n}{n-1}$	~ 1	1
Exponential	$e^{-x}, \quad x > 0$	1		1
Rootzen Gamma	$\tau x^{\tau-1}e^{-x^\tau}$ $x > 0, \tau < 1$	$\sim \frac{1}{\tau}(\ln n)^{1/\tau-1}$	$-$	∞
Lognormal	$\frac{\exp(-2^{-1}\log^2 x)}{x\sqrt{2\pi}}$ $x > 0$	$\sim \frac{1}{\sqrt{2\ln n}}e^{\sqrt{2\ln n}}$	$-$	∞
Pareto	$\alpha x^{-\alpha-1}$ $\alpha > 1, x \geq 1$	$\frac{\Gamma(1+1/\alpha+n)}{\alpha\Gamma(2+1/\alpha)\Gamma(1+n)} \sim \frac{n^{1/\alpha}}{\alpha\Gamma(2+1/\alpha)}$	$-$	∞
Fréchet	$\alpha x^{-\alpha-1}e^{-x^{-\alpha}}$ $\alpha > 1, x \geq 0$	$\frac{1}{\alpha\Gamma(2+1/\alpha)}\frac{n^{1+1/\alpha}}{n-1} \sim \frac{n^{1/\alpha}}{\alpha\Gamma(2+1/\alpha)}$	$-$	∞

This table reproduces Table 1 and adds asymptotic markups for the Sattinger and Hart models. μ_n, μ_n^{Satt} and μ_n^{Hart} are respectively the asymptotic markup expressions for the Perloff–Salop, Sattinger, and Hart models. Asymptotic approximations are calculated using Theorems 1 and 2 except where the markup can be exactly evaluated. The Hart markup is undefined for distributions fatter than the exponential.

Theorem 2 delivers the perhaps unexpected result that the Perloff and Salop (1985), Sattinger (1984), and Hart (1985b) models generate asymptotically equal (up to a multiplicative constant) markups; see Table 3. Hence, detail-independence holds: equilibrium markups do not depend on the details of the model of competition. The key ingredient in the modeling is the specification of the noise distribution, rather than the details of the particular oligopoly model. In particular, these results suggest that the limit-pricing logic of Section 2.4 has broad applicability to random utility models of monopolistic competition.

4. Methodological results

This section presents our main mathematical results. Solving for the symmetric equilibrium outcome for distribution function F requires the evaluation of integrals of the form

$$\int x^j e^{\psi x} f^k(x) F(x)^{n-l} dx \tag{15}$$

where $k, l \geq 1$ and $j, \psi \geq 0$. For large n, such integrals mainly depend on the tail of the distribution F, which suggests that techniques from Extreme Value Theory (EVT) may be applied. (See Resnick, 1987, and Embrechts et al., 1997 for an introduction to EVT.)

Before evaluating (15), we first introduce the notion of *regular variation*.

Definition 3. A function $h : \mathbb{R}^+ \to \mathbb{R}$ is regularly varying at ∞ with index ρ if h is strictly positive in a neighborhood of ∞, and

$$\forall \lambda > 0, \lim_{x \to \infty} \frac{h(\lambda x)}{h(x)} = \lambda^\rho. \tag{16}$$

We indicate this by writing $h \in RV_\rho^\infty$.

Analogously, we say that $h : \mathbb{R}^+ \to \mathbb{R}$ is regularly varying at zero with index ρ if, $\forall \lambda > 0$, $\lim_{x \to 0} h(\lambda x) / h(x) = \lambda^\rho$, and denote this by $h \in RV_\rho^0$. Intuitively, a regularly varying function $h(x)$ with index ρ behaves like x^ρ as x goes to the appropriate limit, perhaps up to logarithmic corrections. For instance, x^ρ and $x^\rho |\ln x|$ are regularly varying (with index ρ) at both 0 and ∞. Much of our analysis requires the concept of regular variation; specifically, we require that certain transformations of the noise distribution F be regularly varying. In the case $\rho = 0$, we say that h is *slowly varying* (for example, $\ln x$ varies slowly at infinity and zero).

Our core mathematical result documents an asymptotic relationship between M_n and $\overline{F}^{-1}(1/n)$.

Theorem 3. *Let F be a differentiable CDF with support on (w_l, w_u) that is strictly increasing in a left neighborhood of w_u. Let $G : (w_l, w_u) \to \mathbb{R}$ be strictly positive in some left neighborhood of w_u. Suppose that $\widehat{G}(t) \equiv G\left(\overline{F}^{-1}(t)\right) \in RV_\rho^0$ with $\rho > -1$, and that $\left|\widehat{G}(t)\right|$ is integrable on $t \in (\bar{t}, 1)$ for all $\bar{t} \in (0, 1)$ (or, equivalently, $G(x) f(x)$ is (w_l, w_u)-integrable in the sense of Definition 2). Then, for $n \to \infty$,*

$$\mathbb{E}[G(M_n)] = \int_{w_l}^{w_u} n G(x) f(x) F(x)^{n-1} dx \sim G\left(\overline{F}^{-1}\left(\frac{1}{n}\right)\right) \Gamma(\rho + 1) \tag{17}$$

where M_n is the largest realization of n i.i.d. random variables with CDF F.

We provide some intuition for equation (17). Note that M_n is the maximum of n F-distributed random variables and exceeds an independent F-distributed random variable n times out of $n + 1$. In other words, $\mathbb{E}[F(M_n)] = \frac{n}{n+1}$, or equivalently $\mathbb{E}\left[\overline{F}(M_n)\right] = \frac{1}{n+1} \approx \frac{1}{n}$. Consequently, we might conjecture – via heroic commutation of the expectations operator – that $\mathbb{E}[M_n] \approx \overline{F}^{-1}\left(\frac{1}{n}\right)$, and more generally that $\mathbb{E}[G(M_n)] \approx G(\mathbb{E}[M_n]) \approx G\left(\overline{F}^{-1}\left(\frac{1}{n}\right)\right)$. It turns out that this heuristic argument gives us the correct approximation, up to a correction factor $\Gamma(\rho + 1)$.[23]

We next present an intermediate result that is neither novel (see Pickands, 1986) nor technically demanding, but allows us to apply Theorem 3 to expressions of the form (15). The proof is straightforward, and consequently omitted.

---

[23] Notice that the correction is downward if $0 < \rho < 1$, and upward otherwise. Informally, this is because $G\left(\overline{F}^{-1}(x)\right)$ behaves like x^ρ close to $x = 0$; so $G\left(\overline{F}^{-1}(x)\right)$ is approximately concave near zero if $0 < \rho < 1$, and approximately convex near zero otherwise. Jensen's inequality then suggests that for large n, $\mathbb{E}[G(M_n)] = \mathbb{E}\left[G\left(\overline{F}^{-1}\left(\overline{F}(M_n)\right)\right)\right]$ is smaller than $G\left(\overline{F}^{-1}\left(\mathbb{E}\left[\overline{F}(M_n)\right]\right)\right) \approx G\left(\overline{F}^{-1}(1/n)\right)$, necessitating a downward correction, if and only if $0 < \rho < 1$.

Lemma 1. *Let F be well-behaved with tail index γ. Then*

1. *$f\left(\overline{F}^{-1}(t)\right) \in RV^0_{\gamma+1}$.*
2. *If $w_u = \infty$, then $\overline{F}^{-1}(t) \in RV^0_{-\gamma}$. If $w_u < \infty$, then $w_u - \overline{F}^{-1}(t) \in RV^0_{-\gamma}$.*
3. *If a is finite, then $e^{\overline{F}^{-1}(t)} \in RV^0_{-a}$.*

Lemma 1 ensures that when F is well-behaved, (15) satisfies the conditions imposed in Theorem 3 for a wide range of parameter values. The following proposition is then an immediate implication of Theorem 3 and Lemma 1.

Proposition 3. *Let F be well behaved with tail index γ. Let $j, \psi \geq 0$, $k \geq 1$ and let $x^j e^{\psi x} f^k(x)$ be (w_l, w_u)-integrable. If $j > 0$, assume that $w_u > 0$. If $\psi = 0$, we can treat $\psi a = 0$ in the following expressions. If $(k - j - 1)\gamma - \psi a + k > 0$, then as $n \to \infty$,*

$$\int_{w_l}^{w_u} x^j e^{\psi x} f^k(x) F(x)^{n-l} dx$$

$$\sim \begin{cases} n^{-1} \left(\overline{F}^{-1}(1/n)\right)^j e^{\psi \overline{F}^{-1}(1/n)} f^{k-1}\left(\overline{F}^{-1}(1/n)\right) \Gamma\left((k-j-1)\gamma - \psi a + k\right) : w_u = \infty \\ n^{-1} w_u^j e^{\psi w_u} f^{k-1}\left(\overline{F}^{-1}(1/n)\right) \Gamma\left((k-1)\gamma + k\right) : w_u < \infty \end{cases}.$$

Proposition 3 allows us to approximate (15) for well-behaved distributions.[24] The parameter restriction $(k - j - 1)\gamma - \psi a + k > 0$ is necessary to ensure that (15) does not diverge. For our purposes, this restriction is rather mild. One notable exception is that when $\psi > 0$, we cannot analyze heavy-tailed distributions (which have fatter-than-exponential tails) such as the lognormal distribution; for these distributions, $a = \infty$.[25]

In fact, Theorem 1 is now an immediate corollary of Proposition 3. More generally, these results are relatively easy to apply. For example, the key mathematical objects in Theorem 1, γ and $nf\left(\overline{F}^{-1}(1/n)\right)$, are easy to calculate for most distributions of interest, and are listed for commonly used distributions in Table 4. The following fact, which is verified using Lemma A1.6, may often be useful to simplify calculations further for the case $\gamma \neq 0$: as $n \to \infty$,

$$\frac{1}{nf\left(\overline{F}^{-1}(1/n)\right)} \sim \begin{cases} \gamma \overline{F}^{-1}(1/n), & \gamma > 0 \\ -\gamma(w_u - \overline{F}^{-1}(1/n)), & \gamma < 0 \end{cases}.$$

5. Conclusion

Random utility models are a convenient and tractable tool for analyzing settings of imperfect competition. The choice of noise distributions in random utility models is often influenced by

---

[24] For a antecedent to this result, see Maller and Resnick (1984).

[25] Here we define a distribution to be heavy-tailed if $e^{\lambda x} \overline{F}(x) \to \infty$ as $x \to \infty$ for all $\lambda > 0$. To see why $a = \infty$ in this case, note that $\lim_{x \to \infty} \overline{F}(x)/f(x) = \infty$ implies $-\frac{d}{dx} \log \overline{F}(x) = o(1)$ as $x \to \infty$, so $-\log \overline{F}(x) = o(x)$ and $e^{-\lambda x} = o\left(\overline{F}(x)\right)$ for all λ.

Table 4
Properties of common densities.

Name of distribution	f	γ	$nf\left(\overline{F}^{-1}(1/n)\right)$	$\overline{F}^{-1}(1/n)$
Uniform	$1, \quad x \in [-1,0]$	-1	n	$-\frac{1}{n}$
Bounded Power Law	$\alpha(-x)^{\alpha-1}, \alpha \geq 1, x \in [-1,0]$	$-1/\alpha$	$\alpha n^{1/\alpha}$	$-n^{-1/\alpha}$
Weibull	$\alpha(-x)^{\alpha-1}e^{-(-x)^{\alpha}}, \alpha \geq 1, x < 0$	$-1/\alpha$	$\alpha n^{1/\alpha}$	$\sim -n^{-1/\alpha}$
Bounded Exponential-like	$\frac{e^{-x/(1-x)}}{(1-x)^2}, x \in [0,1]$	0	$(1+\ln n)^2$	$1 - \frac{1}{1+\log n}$
Gaussian	$(2\pi)^{-1/2}e^{-x^2/2}$	0	$\sim \sqrt{2\ln n}$	$\sim \sqrt{2\ln n}$
Rootzen Class	$\kappa\lambda\phi x^{a+\phi-1}e^{-x^{\phi}}, \quad x > 0, \phi > 1$	0	$\sim \phi\lambda^{1/\phi}(\ln n)^{1-1/\phi}$	$\sim (\ln n)^{1/\phi}$
Gumbel	$\exp(-e^{-x}-x)$	0	~ 1	$\sim \ln n$
Exponential	$e^{-x}, \quad x > 0$	0	1	$\ln n$
Lognormal	$(2\pi)^{-1/2}x^{-1}e^{-\left(\log^2 x\right)/2}, x > 0$	0	$\sim \frac{\sqrt{2\ln n}}{\overline{F}^{-1}(1/n)}$	$\sim e^{\sqrt{2\ln n}}$
Pareto	$\alpha x^{-\alpha-1}, \quad \alpha > 0, x \geq 1$	$1/\alpha$	$\alpha n^{-1/\alpha}$	$n^{1/\alpha}$
Fréchet	$\alpha x^{-\alpha-1}e^{-x^{-\alpha}}, \quad \alpha > 0, x \geq 0$	$1/\alpha$	$\alpha n^{-1/\alpha}$	$\sim n^{1/\alpha}$

Densities are listed in order of increasing tail fatness whenever possible.

tractability concerns. It is important to understand the consequences of these modeling choices and, when possible, to expand the set of tractable models. With this challenge in mind, our paper makes three sets of contributions.

First, we derive equilibrium markups for general noise distributions in various types of random utility models of monopolistic competition in large markets. We show that markups are asymptotically determined by the tail behavior of the distribution of taste shocks.

Second, our results reveal a substantial degree of "detail-independence." Specifically, the behavior of price markups are asymptotically identical (up to a constant factor) for all models that we study. Moreover, for the wide class of distributions with a zero extreme value tail exponent – including the canonical case of Gaussian noise – we show that the elasticity of markups to the number of firms is asymptotically zero. In other words, for many types of large markets, markups are relatively insensitive to the degree of competition.

Third, we show how to approximate an integral that is useful for studying a wide range of economic environments in which extreme outcomes determine the equilibrium allocation. For example, our framework can be used to model imperfect competition in large economies, including applications in macroeconomics and trade.

Our analysis is agnostic about the source of noise in consumer choice. The noise may reflect either heterogenous preferences with normative validity or consumer confusion about product quality. Thus, our results are relevant to both the classical literature on imperfect competition and the emerging literature on behavioral industrial organization. That said, we find the behavioral interpretation that noise arises from consumer mistakes particularly intriguing. Consumer errors in product evaluation may arise from a variety of mechanisms. Let us briefly outline two hypotheses. First, firms may engage in obfuscation to confuse naive (boundedly rational) consumers about product quality.[26] This point is developed in a number of recent papers, including Spiegler (2006), Gabaix and Laibson (2006), Ellison and Ellison (2009), Armstrong and Vickers (2012),

---

[26] Our basic model exogenously specifies the degree of "obfuscation". In the online appendix, we augment our model to consider deliberate shrouding/obfuscation by sellers, and show that our key insights are preserved in this richer setting.

and Heidhues et al. (2014, 2016).[27] Second, consumers may be influenced by a multitude of idiosyncratic behavioral cues in their decision-making. For example, a consumer who is evaluating a mutual fund may rely on otherwise uninformative 'tips' from his friends and family.

Outside the scope of the present paper, but of definite interest for future work, is to allow for firm heterogeneity in the model. The analysis of asymmetric outcomes in large markets introduces additional mathematical challenges, but may produce further insights. Such an extension would allow us to address some stylized facts about competition in large monopolistically-competitive markets, such as variation in markups across firms. It could also potentially lead to a richer set of testable empirical implications.

Appendix A. Proofs

This appendix proves the methodological results from Section 4, then applies them to prove the economic results of Sections 2 and 3. To clarify notation: denote $f_n \sim g_n$ if $f_n/g_n \to 1$, $f_n = o(g_n)$ if $f_n/g_n \to 0$ and $f_n = O(g_n)$ if there exists $M > 0$ and $n' \geq 1$ such that for all $n \geq n'$, $|f_n| \leq M |g_n|$.

A.1. Methodological results

We start by collecting some useful facts about regular variation; for background, see Resnick (1987) or Bingham et al. (1989).

Lemma A1.

1. *If $g(t) \in RV_a^0$, then the limit $\lim_{t\to 0} g(xt)/g(t) = x^a$ holds locally uniformly (with respect to x) on $(0,\infty)$.*
2. *If $\lim_{x\to 0} h(x)/s(x) = 1$, $\lim_{x\to 0} s(x) = 0$ and $g(x) \in RV_\rho^0$, then $g(h(x)) \sim g(s(x))$.*
3. *If $g(t) \in RV_a^0$ and $h(t) \in RV_b^0$, then $g(t)h(t) \in RV_{a+b}^0$.*
4. *If $g(t) \in RV_a^0$, $h(t) \in RV_b^0$ and $\lim_{t\to 0} h(t) = 0$, then $g \circ h(t) \in RV_{ab}^0$.*
5. *If $g(t) \in RV_a^0$ and non-decreasing, then $g^{-1}(t) \in RV_{a^{-1}}^0$ if $\lim_{t\to 0} g(t) = 0$.*
6. *Let $U \in RV_\rho^0$. If $\rho > -1$ (or $\rho = -1$ and $\int_0^x U(t)\,dt < \infty$), then $\int_0^x U(t)\,dt \in RV_{\rho+1}^0$ and*

$$\lim_{x\to 0} \frac{x U(x)}{\int_0^x U(t)\,dt} = \rho + 1.$$

   If $\rho \leq -1$, then for $\overline{x} > 0$, $\int_x^{\overline{x}} U(t)\,dt \in RV_{\rho+1}^0$ and

$$\lim_{x\to 0} \frac{x U(x)}{\int_x^{\overline{x}} U(t)\,dt} = -\rho - 1.$$

7. *If $\lim_{t\to\infty} t j'(t)/j(t) = \rho$, then $j \in RV_\rho^\infty$. Similarly, if $\lim_{t\to 0} t j'(t)/j(t) = \rho$, then $j \in RV_\rho^0$.*
8. *If $g \in RV_\rho^\infty$ and $\varepsilon > 0$, then $g(t) = o\left(t^{\rho+\varepsilon}\right)$ and $t^{\rho-\varepsilon} = o(g(t))$ as $t \to \infty$; and if $g \in RV_\rho^0$ and $\varepsilon > 0$, then $g(t) = o\left(t^{\rho-\varepsilon}\right)$ and $t^{\rho+\varepsilon} = o(g(t))$ as $t \to 0$.*

---

[27] Relatedly, other papers (e.g. Bordalo et al., 2016) emphasize the impact of endogenous salience on market equilibrium.

Proof. See the online appendix. □

Our proof of Theorem 3 depends critically on the following seminal result.

Theorem A1 (*Karamata's Tauberian Theorem*). *Assume* $U : (0, \infty) \to [0, \infty)$ *is weakly increasing. Let* $\tilde{U}(s) = \int_0^\infty e^{-sx} dU(x)$ *be the Laplace–Stieltjes transform of* $U(x)$*, and assume* $\tilde{U}(s) < \infty$ *for all sufficiently large s. Then with* $\alpha \geq 0$*,* $U(x) \in RV_\alpha^0$ *if and only if* $\tilde{U}(s) \in RV_{-\alpha}^\infty$*. Further, if either condition holds, then*

$$\int_0^\infty e^{-sx} dU(x) \sim_{s \to \infty} U(1/s) \Gamma(\alpha + 1).$$

For a proof, see Bingham et al. (1989, pp. 38, Th. 1.7.1') or Feller (1971, XIII.5, Th. 1).

Proof of Theorem 3. Assume for now that $G(x) \geq 0$ for all $x \in (w_l, w_u)$; we relax this assumption later. Differentiation of $\mathbb{P}(M_n \leq x) = F^n(x)$ gives the density of M_n: $f_n(x) = nf(x)F^{n-1}(x)$. Using the change of variable $x = \overline{F}^{-1}(t)$ and observing that $d\overline{F}^{-1}(t)/dt = -1/f\left(\overline{F}^{-1}(t)\right)$

$$\mathbb{E}[G(M_n)] = \int_{w_l}^{w_u} G(x)nf(x)F^{n-1}(x)dx$$

$$= n \int_{w_l}^{w_u} G(x)F^{n-1}(x)(f(x)dx)$$

$$= n \int_0^1 G(\overline{F}^{-1}(t))[F(\overline{F}^{-1}(t))]^{n-1}dt$$

$$= n \int_0^1 \widehat{G}(t)(1-t)^{n-1}dt.$$

We next use the change in variables $x = -\ln(1-t)$, so $t = 1 - e^{-x}$, $dt = e^{-x}dx$, and so

$$\mathbb{E}[G(M_n)] = n \int_0^\infty \widehat{G}\left(1 - e^{-x}\right) e^{-x} e^{-n'x} dx$$

where $n' = n - 1$. Define $h(x) = \widehat{G}\left(1 - e^{-x}\right) e^{-x}$, and $\mu(x) = \int_0^x h(y) dy$. Since $\widehat{G}$ is regularly varying at zero with index $\rho > -1$, Lemma A1.8 implies that $\int_0^s |\widehat{G}(t)| dt < \infty$ for sufficiently small s. This, with the assumptions $G(t) \geq 0$ and $\int_s^1 |\widehat{G}(t)| dt < \infty$ for all $s \in (0, 1)$, ensure that $\mu(x) = \int_0^{1-e^{-x}} \widehat{G}(t) dt$ is finite and non-decreasing on $[0, \infty)$. By Lemma A1.2, $h(x) \sim_{x \to 0} \widehat{G}(x)$. So $h \in RV_\rho^0$, and by Lemma A1.6

$$\mu(x) = \int_0^x h(y)\,dy \sim_{x\to 0} \frac{1}{1+\rho} h(x) x \sim_{x\to 0} \frac{1}{1+\rho} \widehat{G}(x) x.$$

Therefore, $\mu(x) \in RV_{\rho+1}^0$. Noting our assumption that $\rho + 1 > 0$, we can now apply Karamata's Theorem A1 in combination with the last expression to obtain

$$\int_0^\infty e^{-n'x} d\mu(x) \sim_{n'\to\infty} \mu(1/n') \Gamma(2+\rho)$$

$$\sim_{n'\to\infty} \frac{1}{1+\rho} \widehat{G}(1/n') (n')^{-1} \Gamma(2+\rho)$$

$$\sim_{n\to\infty} \widehat{G}(1/n) n^{-1} \Gamma(1+\rho).$$

Thus

$$\mathbb{E}[G(M_n)] = n \int_0^\infty e^{-n'x} d\mu(x)$$

$$\sim n\widehat{G}(1/n) n^{-1} \Gamma(1+\rho) = G(\overline{F}^{-1}(1/n))\Gamma(1+\rho)$$

holds when $G(x) \geq 0$ for all $x \in (w_l, w_u)$. Now relax the assumption that $G(x) \geq 0$ for all $x \in (w_l, w_u)$. Choose $\bar{t} \in (0, 1)$ such that $G(t) > 0$ for $t \in [0, \bar{t}]$. The assumption that $G(\cdot)$ is strictly positive in a left neighborhood of w_u ensures that such $\bar{t}$ exists. Thus we can write

$$\mathbb{E}[G(M_n)] = n \int_0^{\bar{t}} \widehat{G}(t)(1-t)^{n-1}\,dt + n \int_{\bar{t}}^1 \widehat{G}(t)(1-t)^{n-1}\,dt$$

Consider $\widetilde{G} : (0, 1) \to \mathbb{R}$ defined by

$$\widetilde{G}(t) \equiv \begin{cases} \widehat{G}(t) : t \leq \bar{t} \\ 0 : t > \bar{t} \end{cases}.$$

It is easy to check that $\widetilde{G}$ satisfies the conditions of the theorem and additionally is weakly positive everywhere on (w_l, w_u). The argument above shows that as $1/n \to 0$

$$n \int_0^{\bar{t}} \widehat{G}(t)(1-t)^{n-1}\,dt = n \int_0^1 \widetilde{G}(t)(1-t)^{n-1}\,dt$$

$$\sim \widetilde{G}(1/n)\Gamma(1+\rho) \sim \widehat{G}(1/n)\Gamma(1+\rho). \tag{18}$$

To complete the proof we demonstrate that as $n \to \infty$

$$\left| \int_{\bar{t}}^1 \widehat{G}(t)(1-t)^{n-1}\,dt \right| = o\left(\int_0^{\bar{t}} \widehat{G}(t)(1-t)^{n-1}\,dt \right).$$

First, by (18): for $n \to \infty$,

$$\int\limits_0^{\bar{t}} \widehat{G}(t)(1-t)^{n-1} dt \sim n^{-1} \widehat{G}(1/n)\,\Gamma(1+\rho) \in RV_{-\rho-1}^{\infty}.$$

Lemma A1.8 implies that $\int_0^{\bar{t}} \widehat{G}(t)(1-t)^{n-1} dt > n^{-\rho-1-\varepsilon}$ for sufficiently large n and given some $\varepsilon > 0$. Also,

$$\left| \int\limits_{\bar{t}}^1 \widehat{G}(t)(1-t)^{n-1} dt \right| \le \int\limits_{\bar{t}}^1 \left| \widehat{G}(t) \right| (1-t)^{n-1} dt$$

$$\le \left(1-\bar{t}\right)^{n-1} \int\limits_{\bar{t}}^1 \left| \widehat{G}(t) \right| dt$$

$$\le \left(1-\bar{t}\right)^{n-1} \int\limits_0^1 \left| \widehat{G}(t) \right| dt.$$

By assumption $\int_s^1 \left| \widehat{G}(t) \right| dt < \infty$ for all $s \in (0,1)$, therefore

$$\frac{\left| \int_{\bar{t}}^1 \widehat{G}(t)(1-t)^{n-1} dt \right|}{\int_0^{\bar{t}} \widehat{G}(t)(1-t)^{n-1} dt} \le \frac{\left(1-\bar{t}\right)^{n-1} \int_0^1 \left| \widehat{G}(t) \right| dt}{n^{-\rho-1-\varepsilon}} = o(1) \text{ as } n \to \infty.$$

This completes the proof. □

Proof of Proposition 3. Follows immediately from Theorem 3 and Lemma 1. □

A.2. Economic results: markups, elasticities, and auctions

Proof of Theorem 1. Follows immediately from Proposition 3. □

Proof of Proposition 1. Treating n as continuous, we have

$$\frac{n}{\mu_n^{PS}} \frac{d\mu_n^{PS}}{dn} = -\left(\frac{2n-1}{n-1} + \frac{n \int f^2(x) F^{n-2}(x) \log F(x)\, dx}{\int f^2(x) F^{n-2}(x)\, dx} \right).$$

Noting that $-\log(1-x) \sim x \in RV_1^0$, applying Theorem 3 to $G(x) \equiv \frac{f(x)}{F(x)} \log F(x)$, using Lemma A1.3, we obtain

$$\int f^2(x) F^{n-2}(x) \log F(x)\, dx \sim -n^{-2} f\left(\overline{F}^{-1}(1/n) \right) \Gamma(3+\gamma).$$

Together with Theorem 1, it follows that

$$\frac{n}{\mu_n} \frac{d\mu_n}{dn} = -\left(2 - \frac{n^{-2}nf\left(\overline{F}^{-1}(1/n) \right) \Gamma(3+\gamma)}{n^{-2}nf\left(\overline{F}^{-1}(1/n) \right) \Gamma(2+\gamma)} + o(1) \right) = \gamma + o(1). \quad \square$$

Proof of Proposition 2. See the online appendix. □

Proof of Theorem 2. The Sattinger case follows directly from Proposition 3. For the Hart case, see the online appendix.   □

Appendix. Supplementary material

Supplementary material related to this article can be found online at http://dx.doi.org/10.1016/j.jet.2016.04.001.

References

Acemoglu, D., Carvalho, V.M., Ozdaglar, A., Tahbaz Salehi, A., 2012. The network origins of aggregate fluctuations. Econometrica 80 (5), 1977–2016.

Anderson, S.P., De Palma, A., Thisse, J.-F., 1992. Discrete Choice Theory of Product Differentiation. MIT Press, Cambridge, MA.

Arkolakis, C., Costinot, A., Donaldson, D., Rodríguez-Clare, A., 2015. The elusive pro-competitive effects of trade. Working paper.

Armstrong, M., Vickers, J., 2012. Consumer protection and contingent charges. J. Econ. Lit. 50 (2), 477–493.

Armstrong, T., 2016. Large market asymptotics for differentiated product demand estimators with economic models of supply. Working paper.

Auer, R., Chaney, T., 2009. Exchange rate pass-through in a competitive model of pricing-to-market. J. Money Credit Bank. 41 (s1), 151–175.

Ausubel, L.M., 1991. The failure of competition in the credit card market. Am. Econ. Rev. 81 (1), 50–81.

Bajari, P., Benkard, C.L., 2003. Discrete choice models as structural models of demand: some economic implications of common approaches. Working paper.

Bénabou, R., Gertner, R., 1993. Search with learning from prices: does increased inflationary uncertainty lead to higher markups. Rev. Econ. Stud. 60 (1), 69–93.

Bergstresser, D., Chalmers, J.M.R., Tufano, P., 2009. Assessing the costs and benefits of brokers in the mutual fund industry. Rev. Financ. Stud. 22 (10), 4129–4156.

Bernard, A.B., Eaton, J., Jensen, J.B., Kortum, S., 2003. Plants and productivity in international trade. Am. Econ. Rev. 93 (4), 1268–1290.

Bingham, N.H., Goldie, C.M., Teugels, J.L., 1989. Regular Variation. Cambridge University Press, Cambridge, UK.

Bordalo, P., Gennaioli, N., Shleifer, A., 2016. Competition for attention. Rev. Econ. Stud. 83 (2), 481–513.

Bresnahan, T.F., Reiss, P.C., 1991. Entry and competition in concentrated markets. J. Polit. Econ. 99 (5), 977–1009.

Bulow, J., Klemperer, P., 2002. Prices and the winner's curse. Rand J. Econ. 33 (1), 1–21.

Bulow, J., Klemperer, P., 2012. Regulated prices, rent seeking, and consumer surplus. J. Polit. Econ. 120 (1), 160–186.

Campbell, J.R., Hopenhayn, H.A., 2005. Market size matters. J. Ind. Econ. 53 (1), 1–25.

Carlin, B.I., 2009. Strategic price complexity in retail financial markets. J. Financ. Econ. 91 (3), 278–287.

Caserta, S., 2002. Extreme values in auctions and risk analysis. Ph.D. thesis #287. Tinbergen Institute, Rotterdam, Netherlands.

Chaney, T., 2008. Distorted gravity: the intensive and extensive margins of international trade. Am. Econ. Rev. 98 (4), 1707–1721.

Chaney, T., 2015. The gravity equation in international trade: an explanation. Working paper.

Chen, Y., Riordan, M.H., 2008. Price-increasing competition. Rand J. Econ. 39 (4), 1042–1058.

Dagsvik, J.K., 1994. Discrete and continuous choice, max-stable processes, and independence from irrelevant attributes. Econometrica 62 (5), 1179–1205.

Dagsvik, J.K., Karlstrom, A., 2005. Compensating variation and Hicksian choice probabilities in random utility models that are nonlinear in income. Rev. Econ. Stud. 72 (1), 57–76.

Dixit, A.K., Stiglitz, J.E., 1977. Monopolistic competition and optimum product diversity. Am. Econ. Rev. 67 (3), 297–308.

Eaton, J., Kortum, S., 2002. Technology, geography, and trade. Econometrica 70 (5), 1741–1779.

Ellison, G., Ellison, S.F., 2009. Search, obfuscation, and price elasticities on the Internet. Econometrica 77 (2), 427–452.

Embrechts, P., Klüppelberg, C., Mikosch, T., 1997. Modelling Extremal Events. Springer, New York.

Feller, W., 1971. An Introduction to Probability Theory and Its Applications, Volume II. John Wiley & Sons, Hoboken, NJ.

Gabaix, X., 1999. Zipf's law for cities: an explanation. Q. J. Econ. 114 (3), 739–767.

Gabaix, X., 2011. The granular origins of aggregate fluctuations. Econometrica 79 (3), 733–772.

Gabaix, X., Laibson, D., 2006. Shrouded attributes, consumer myopia, and information suppression in competitive markets. Q. J. Econ. 121 (2), 505–540.

Gabaix, X., Landier, A., 2008. Why has CEO pay increased so much? Q. J. Econ. 123 (1), 49–100.

Gennaioli, N., Shleifer, A., Vishny, R., 2015. Money doctors. J. Finance 70 (1), 91–114.

Gnedenko, B., 1943. Sur La Distribution Limite Du Terme Maximum D'Une Série Aléatoire. Ann. Math. 44 (3), 423–453.

Hart, O.D., 1985a. Monopolistic competition in the spirit of Chamberlin: a general model. Rev. Econ. Stud. 52 (4), 529–546.

Hart, O.D., 1985b. Monopolistic competition in the spirit of Chamberlin: special results. Econ. J. 95 (380), 889–908.

Heidhues, P., Koszegi, B., Murooka, T., 2014. Inferior products and profitable deception. Working paper.

Heidhues, P., Koszegi, B., Murooka, T., 2016. Exploitative innovation. Am. Econ. J. Microecon. 8 (1), 1–23.

Henderson, B.J., Pearson, N.D., 2011. The dark side of financial innovation: a case study of the pricing of a retail financial product. J. Financ. Econ. 100 (2), 227–247.

Hong, H., Shum, M., 2004. Rates of information aggregation in common value auctions. J. Econ. Theory 116 (1), 1–40.

Hortaçsu, A., Syverson, C., 2004. Product differentiation, search costs, and competition in the mutual fund industry: a case study of S&P 500 index funds. Q. J. Econ. 119 (2), 403–456.

Ibragimov, R., Jaffee, D., Walden, J., 2009. Nondiversification traps in catastrophe insurance markets. Rev. Financ. Stud. 22 (3), 959–993.

Ibragimov, R., Jaffee, D., Walden, J., 2011. Diversification disasters. J. Financ. Econ. 99 (2), 333–348.

Ibragimov, R., Walden, J., 2010. Optimal bundling strategies under heavy-tailed valuations. Manag. Sci. 56 (11), 1963–1976.

Jansen, D.W., de Vries, C.G., 1991. On the frequency of large stock returns: putting booms and busts into perspective. Rev. Econ. Stat. 73 (1), 18–24.

Jones, C.I., 2005. The shape of production functions and the direction of technical change. Q. J. Econ. 120 (2), 517–549.

Luce, R.D., 1959. Individual Choice Behavior, A Theoretical Analysis. John Wiley, New York.

Luttmer, E.G.J., 2007. Selection, growth, and the size distribution of firms. Q. J. Econ. 122 (3), 1103–1144.

Maller, R.A., Resnick, S.I., 1984. Limiting behaviour of sums and the term of maximum modulus. Proc. Lond. Math. Soc. s3-49 (3), 385–422.

Mangin, S., 2015a. A theory of production, matching, and distribution. Working paper.

Mangin, S., 2015b. How to build a production function. Working paper.

Mas-Colell, A., 1975. A model of equilibrium with differentiated commodities. J. Math. Econ. 2 (2), 263–295.

Mazzeo, M.J., 2002. Product choice and oligopoly market structure. Rand J. Econ. 33 (2), 221–242.

McFadden, D., 1981. Econometric models of probabilistic choice. In: Manski, C.F., McFadden, D. (Eds.), Structural Analysis of Discrete Data with Econometric Applications. MIT Press, Cambridge, MA, pp. 198–272.

Melitz, M.J., Ottaviano, G.I.P., 2008. Market size, trade, and productivity. Rev. Econ. Stud. 75 (1), 295–316.

Perloff, J.M., Salop, S.C., 1985. Equilibrium with product differentiation. Rev. Econ. Stud. 52 (1), 107–120.

Pickands, J., 1986. The continuous and differentiable domains of attraction of the extreme value distributions. Ann. Probab. 14 (3), 996–1004.

Quint, D., 2014. Imperfect competition with complements and substitutes. J. Econ. Theory 152, 266–290.

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

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

Sattinger, M., 1984. Value of an additional firm in monopolistic competition. Rev. Econ. Stud. 51 (2), 321–332.

Spiegler, R., 2006. Competition over agents with boundedly rational expectations. Theor. Econ. 1, 207–231.

Stango, V., 2000. Competition and pricing in the credit card market. Rev. Econ. Stat. 82 (3), 499–508.

Vives, X., 1985. On the efficiency of Bertrand and Cournot equilibria with product differentiation. J. Econ. Theory 36 (1), 166–175.

Vives, X., 2001. Oligopoly Pricing, Old Ideas and New Tools. MIT Press, Cambridge, MA.

Weyl, E.G., Fabinger, M., 2013. Pass-through as an economic tool: principles of incidence under imperfect competition. J. Polit. Econ. 121 (3), 528–583.

Zhelobodko, E., Kokovin, S., Parenti, M., Thisse, J.-F., 2012. Monopolistic competition: beyond the constant elasticity of substitution. Econometrica 80 (6), 2765–2784.