The CES distribution circle and its decoupling

Simon P. Anderson and André de Palma*

Revised January 2018

Abstract

We show for CES demands with heterogeneous production costs that profit, revenue, and output distributions lie in the same class as the productivity distribution (e.g., the "Pareto circle"), although the price distribution lies in the inverse class. We relate distribution shapes via the elasticities of their densities. Introducing product quality decouples the CES circle. Then, all distributions can lie in the same class. For instance, it becomes possible to reconcile Pareto price and Pareto profit distributions. We also describe discrete choice underpinnings for the representative consumer, and determine the equilibrium distribution of heterogeneous individual consumers' welfare.

JEL Classification: L13, F12

Keywords: CES; monopolistic competition; quality; distributions; Pareto, power, and log-normal distributions.

*Department of Economics, University of Virginia, USA, sa9w@virginia.edu; CES-University Paris-Saclay, FRANCE. andre.depalma@ens-cachan.fr. The first author gratefully acknowledges research funding from the NSF. We thank Julien Monardo, Maxim Engers, Isabelle Mejean, Farid Toubal, James Harrigan, and Ariell Resheff for valuable comments, and seminar participants at Hebrew University, Melbourne University, Stockholm University, KU Leuven, Laval, Vrij Universiteit Amsterdam, CUHK, Hitotsubashi, University of Tokyo, and Paris Dauphine. This is a revised version of part of CEPR DP 10748. Some of the results in Section 2 are in Mrazova, Neary, and Parenti's CEPR 12044.

1 Introduction

Distributions have been studied for a long time in social sciences by geographers and others, and has recently attracted more interest in economics (Axtell, 2001, Gabaix, 2016). Indeed, interest in revenue distributions in international trade has exploded over the last decade (following Melitz, 2003). However, very little research has been addressed to studying multiple related distributions, such as profit, price, and output distributions. Profit, prices, and outputs are clearly related through profit-maximizing relations. Here we develop equilibrium connections between *distributions* in the paradigm CES demand formulation. For example, can a log-normal distribution of unit costs be consistent with a log-normal distribution of firms' outputs, or a Pareto distribution for profit be consistent with a Pareto distribution for firm prices? The answer is affirmative for the first one: in fact if unit costs are log-normally distributed, then output *must* be log-normally distributed. But (for the second question) a Pareto profit distribution function is only consistent with a power distribution of prices, although when we extend the CES for quality heterogeneity we can indeed render Pareto distributions for both.

The CES representative consumer model is widely used in economics in conjunction with monopolistic competition.[1] A flurry of recent contributions deploy the CES and variants thereof (e.g., Dhingra and Morrow, 2017, Zhelobodko, Kokovin, Parenti, and Thisse, 2012, Bertoletti and Etro, 2017, etc.). The current most intensive use of the model is in International Trade, following Melitz (2003), where it is at the heart of empirical estimation, and it has enjoyed a huge spurt in popularity in the new international trade literature.[2] It is used as a theoreti-

[1] See the original book on monopolistic competition by Chamberlin (1933).
[2] Although note that Fajgelbaum, Grossman, and Helpman (2011) take a nested multinomial logit approach.

cal component in the New Economic Geography and Urban Economics, it is the linchpin of Endogenous Growth Theory, Keynesian underpinnings in Macro, and Industrial Organization. The convenience of the model stems from its analytic manipulability. The CES model delivers equilibrium mark-ups proportional to marginal costs, and so delivers market power (imperfect competition) in a simple way without complex market interaction. The standard models in this vein (following Melitz, 2003) assume that firms' unit production costs are heterogeneous.

However, when we apply this model to distributions, if one distribution (such as profit) is Pareto (1965, 1896), then the distributions of all the economic variables (productivity, revenue, profit, output, and price reciprocal) lie in the Pareto class. This we call the "Pareto circle". More generally, we establish the *CES circle* because the result applies to any distribution class. Previous authors have derived special cases of this relation for productivity and sales revenue for particular distributions: Helpman, Melitz, and Yeaple (2004) cover Pareto, while Head, Mayer, and Thoenig (2014) treat the log-normal distribution. We extend these results to other distribution classes and to the other economic variables.

This analysis enables us to describe the relation between the *shapes* of the linked distributions, such as inherited concavity/convexity properties and how these depend on the demand elasticity of the CES. In particular, we show how the elasticity of the density of costs is a key statistic, along with the demand elasticity, in determining the elasticities of both output and profit densities, and we show the three-way linkage between all three elasticity densities. This analysis generalizes the insights from the constant elasticity case associated to the Pareto distribution.

The CES circle imposes severe restrictions on the linked distributions. The (ubiquitous) CES demand side can be retained while decoupling the CES circle by introducing qualities. These we introduce in the same way as do Baldwin and Harrigan (2011) and Feenstra and Romalis (2014).[3] Doing so delivers two fundamental drivers of equilibrium distributions (instead of just one) – the cost distribution and the quality/cost one. Even if one distribution is Pareto, others can belong to different classes. Most notably, the output distribution depends on the cost distribution (as before) but now also on the quality/cost distribution.

The CES demand system is usually derived from a representative consumer's utility function, although individual consumers typically consume few variants in practice. The demand system though can be shown to be consistent with a population of consumers with heterogeneous preferences, and making discrete choices across products. Then there is another distribution (on the consumer side) associated to CES. This is the distribution of consumer benefits in the population. We show this is a Gumbel (Type 1 Extreme Value) distribution, and its equilibrium shape (more precisely, its location parameter) depends on the distribution of quality/cost.

In the sequel, we first develop the analysis for cost heterogeneity alone and show how the productivity distribution delivers the CES circle. We also show that equilibrium densities are related by simple linear relations between their elasticities. We then allow for quality too in order to decouple the CES circle. Section 4 turns to a disaggregate micro-foundation of the CES representative consumer in which the distribution of tastes delivers a Gumbel distribution for individual consumer benefits. This distribution depends only on a simple summary statistic that comes from the quality-cost distribution. A final section gives some concluding remarks.

[3] These authors assume a Pareto distribution for productivities for their empirical work. They do not pursue the implications of the properties of the various distributions and how they are linked, which is our goal here.

2 Standard CES model

2.1 Basic model with cost heterogeneity and the CES circle

We start with the standard CES monopolistic competition model with heterogeneity only in firms' productivities (the reciprocal of unit production costs). This is the basic Melitz (2003) approach. We show how all economic distributions (prices, output, profit, and revenue) are tied down by the productivity distribution.

Several forms of CES representative consumer utility functions are prevalent in the literature. We nest these into one embracing form. The CES representative consumer involves a sub-utility functional for the differentiated product $\chi = \left(\int_{\Omega} q(\omega)^{\rho} d\omega\right)^{1/\rho}$, where the q's are quantities consumed of the differentiated variants, and $\rho \in (0,1)$: variants are perfect substitutes for $\rho = 1$, demands are independent for $\rho \to 0$. The elasticity of substitution is $\sigma = \frac{1}{1-\rho} \in (1,\infty)$: this statistic is often reported in empirical studies.

The individual variants are denoted by ω, and each is produced by a separate firm; the set of variants is denoted by Ω. Common forms of representative consumer formulation are:

(i) Melitz (2003) model (see also Dinghra and Morrow, 2017), where $U = \chi$ so there is only one sector;

(ii) the classic Dixit-Stiglitz (1977) case much used in earlier trade theory (e.g., Helpman and Krugman, 1985), $U = \chi q_0^{\eta}$ with $\eta > 0$, where q_0 is consumption in an outside numeraire sector;

(iii) $U = \ln \chi + q_0$, which constitutes a quasi-linear form (with no income effects) and so constitutes a partial equilibrium approach (see Anderson and de Palma, 2000, and Nocke and

Yeaple, 2008).

The first two formulations have unit income elasticities of demand; hence their popularity in Trade models. Utility is maximized under the budget constraint $\int_{\Omega} q(\omega) p(\omega) d\omega + q_0 \leq I$, where I is income and $p(\omega)$ is the price of variant ω. We need $I > 1$ for case (iii), or else it reverts to case (i) because all income is spent on the differentiated variants.

Let the (constant) unit production costs of variant ω be $c(\omega)$, and let $F_C(c)$ denote the cumulative distribution function of these costs, with $f_C(c)$ the probability density of unit production costs.

The next results are quite standard. For a given set of prices and a set Ω of active firms (with total mass $M = \|\Omega\|$), Firm i's demand is:

$$h(p_i) = \frac{\Xi(I)}{p_i} \frac{p_i^{\frac{\rho}{\rho-1}}}{\int_{\omega \in \Omega} p(\omega)^{\frac{\rho}{\rho-1}} d\omega}, \tag{1}$$

where $\Xi(I)$ is I for case (i), $\frac{I}{1+\eta}$ for case (ii) (which clearly nests case (i) for $\eta = 0$); and 1 for case (iii). In each case, $\Xi(I)$ is the total amount spent on the differentiated commodity. The denominator in (1) represents the aggregate impact of firms' actions on individual demand: under monopolistic competition, each firm's action has no effect on this statistic. Notice that the CES demand system exhibits the IIA (Independence from Irrelevant Alternatives) property that the ratio of the demands for any two products is independent of the price of any other product.[4] The analysis of Section 4 indicates why the CES has this property, which is usually associated to logit discrete choice models.

Firm i's profit maximizing price solves $\max_{p_i} \left\{ \frac{(p_i - c_i)}{p_i} p_i^{\frac{\rho}{\rho-1}} \right\}$, so equilibrium $p_i = \frac{c_i}{\rho}$, and the

[4]That is, from (1), $h(p_i)/h(p_j) = (p_i/p_j)^{1/(\rho-1)}$.

equilibrium Lerner index is

$$\frac{p_i - c_i}{p_i} = (1 - \rho) . \tag{2}$$

Given such pricing, Firm i's equilibrium output is (from (1))

$$y_i = \rho \Xi (I) \frac{c_i^{\frac{1}{\rho - 1}}}{D_C}, \tag{3}$$

where $D_C = M \int c(u)^{\frac{\rho}{\rho-1}} f_C(u) \, du$, and $f_C(.)$ is the density of unit production costs.

Firm i's equilibrium profit, π_i, is proportional to its sales revenue, $r_i = p_i y_i$, so that

$$r_i = \Xi (I) \frac{c_i^{\frac{\rho}{\rho - 1}}}{D_C}, \tag{4}$$

and

$$\pi_i = (1 - \rho) r_i. \tag{5}$$

Hence equilibrium output is elastic with respect to cost and has elasticity $\frac{1}{\rho-1} = -\sigma < -1$, while profit (and also sales revenue) has elasticity $\frac{\rho}{\rho-1} = 1 - \sigma < 0$, which is smaller in absolute terms. Doubling cost causes output to go down by a factor more than half, while profit goes down proportionately less (because of the price increase).

2.2 Distribution classes

We can now tie together the various equilibrium distributions with the help of the following straightforward result, which tells us how distributions are modified by multiplicative and power *transformations.*[5] These transformations relate profit, revenue, output, price reciprocal $(1/p)$, and productivity (the cost reciprocal, $1/c$) in the CES model.

Lemma 1 *(Transformation) Let $F_X(x)$ be the CDF of a random variable X. Then:*

a) (Multiplicative) the CDF of $Y = kX$ with $k > 0$ is $F_Y(y) = F_{kX}(y) = F_X\left[\left(\frac{y}{k}\right)\right]$;

b) (Positive power) the CDF of $Y = kX^\theta$ with $k > 0$ and $\theta > 0$ is $F_Y(y) = F_X\left[\left(\frac{y}{k}\right)^{\frac{1}{\theta}}\right]$;

c) (Negative power) the CDF of $Y = kX^\theta$ with $k > 0$ and $\theta < 0$ is $F_Y(y) = 1 - F_X\left[\left(\frac{y}{k}\right)^{\frac{1}{\theta}}\right]$.

These three transformations define three corresponding classes of distributions. We will refer to distributions that are related by positive multiplicative transforms as being in the *same multiplicative class*. Similarly, we will term those distributions related by positive power transformations as being in the *same power class* (so the multiplicative class is a subset of this). Hence, pairs of distributions with the same functional forms but different (same-signed) parameters are in the same power class. When one distribution is related to another by a negative power transform, we will say they are in *inverse classes.*[6] Hence one distribution is in the inverse class of another one if it is the survival function of a negative power transformation of the other distribution.

For example, power distributions beget power distributions under positive power transforms. They beget Pareto distributions under negative power transforms and conversely: Pareto distributions beget power distributions under negative power transforms. Thus Pareto and power

[5] For example, for property (b) below, $F_Y(y) = \Pr(Y < y) = \Pr\left(kX^\theta < y\right) = \Pr\left(X < \frac{y}{k}^{\frac{1}{\theta}}\right) = F_X\left[\left(\frac{y}{k}\right)^{\frac{1}{\theta}}\right]$.

[6] Two successive applications of negative power transforms are a power transformation because the distribution returns to the original power class.

distributions are power classes that are inverse to each other.[7] Another useful power class is the log-normal, which also has the property that it is self-inverse.

Consider positive power transforms first, and let $Y = kX^\theta$ (as above). Then, if $F_X(x) = ax^\alpha$, $\alpha > 0$, we get from (b) that $F_Y(y) = F_X\left[\left(\frac{y}{k}\right)^{\frac{1}{\theta}}\right] = a\left(\frac{y}{k}\right)^{\frac{\alpha}{\theta}}$, so the algorithm for the transformation is to replace the distribution elasticity α by $\frac{\alpha}{\theta}$. Similarly, for negative power transforms, from (c) we get $F_Y(y) = 1 - F_X\left[\left(\frac{y}{k}\right)^{\frac{1}{\theta}}\right] = 1 - a\left(\frac{y}{k}\right)^{\frac{\alpha}{\theta}}$, with the tail power given by $\frac{\alpha}{\theta}$. This latter relation is the key to moving from a power cost or price distribution to a Pareto output or profit distribution below.

Other authors have noted that for particular distributions, *revenue* and *productivity* take the same distribution form in the CES. The Pareto case was shown by Helpman, Melitz, and Yeaple (2004, p.304), while Head, Mayer, and Thoenig (2014) derived that a log-normal distribution for productivity delivers log-normal revenues.[8] We show analogous results for all distributions through specifying distribution classes. Moreover, we show how the distribution equivalence result extends to encompass all the economic distributions of interest for the CES.

Proposition 1 *(CES circle) (a)The distributions of profit and revenue are in the same multiplicative class, and the distributions of price and cost are in the same multiplicative class. (b) The distributions of profit, revenue, output, price reciprocal, and cost reciprocal (productivity) are all in the same power class. (c) The distributions of price and cost are in the inverse class from the distributions of profit, revenue and output (and conversely).*

[7] A distribution which is a sum of power functions also is a power class because it remains a sum of powers under power transformations. Insofar as any distribution can be approximated arbitrarily closely by a sum of integer power functions, this extends the potential of this approach.

[8] These authors look at the empirical evidence for choosing either log-normal or Pareto distributions for firm size: Nigai (2017) splices the two (with the right tail Pareto) to get the best fit.

Proof. From the analysis above, profit is a positive fraction of revenue (see (5)), so that revenue and profit distributions are in the same multiplicative class. Likewise, cost is a positive fraction of equilibrium price (see (2)), so the same property attains. Profit and revenue distributions are related to the output distribution by a positive power transformation, so that their distributions are in the same power class. From (3), equilibrium output, y_i, is related to the cost reciprocal, $1/c_i$, by a positive power and a positive factor, and similarly for the price reciprocal, $1/p_i$. Hence, the second statement of the Proposition follows. The last statement follows because costs and prices are related to output via a negative power transformation, and to revenue and profit by a negative power transformation. ■

In particular, if any one of the distributions of profit, revenue, output, price reciprocal, and cost reciprocal (i.e., productivity) is Pareto (resp. power), then they all are Pareto (resp. power), although they have different parameters. Moreover, if any one distribution is (truncated) log-normal then they all are. We provide other examples below. This result we term the *CES-circle*.

We first illustrate with the power cost distribution:

$$F_C(c) = \left(\frac{c}{\underline{c}}\right)^{\beta} \text{ for } c \in [0, \underline{c}]. \tag{6}$$

Note here that we have taken the lower bound to cost as 0 to correspond to the infinite upper bound to profit (and output and revenue) for a Pareto distribution. A finite upper bound to profit corresponds to a strictly positive lower bound to cost (in which case the distribution class is a bounded Pareto in what follows).[9] Lemma 1*(a)* implies that the price distribution

[9] The $\underline{c}$ value might be endogenously determined by fixed cost, as in Melitz (2003).

must also be Pareto. We can find the profit distribution using Lemma 1*(c)* with $\theta = \frac{\rho}{\rho-1}$ (the elasticity of profit with respect to c) from (4) and (5):

$$F_{\Pi}(\pi) = 1 - \left(\frac{\underline{\pi}}{\pi}\right)^{\alpha_\pi} \quad \text{for } \pi \geq \underline{\pi}, \tag{7}$$

which is a Pareto distribution with tail parameter $\alpha_\pi = \frac{1-\rho}{\rho}\beta$, and $\underline{\pi}$ is given by (4) and (5) with $c = \underline{c}$. We can apply Lemma 1*(a)* with $k = (1-\rho)$ to give the revenue distribution as $F_R(r) = F_{\frac{\Pi}{(1-\rho)}}(r) = F_{\Pi}(r(1-\rho))$, so that revenue follows the same distribution as profit (with the same tail parameter) except with a higher lower bound for the support, $\underline{r} = \frac{\underline{\pi}}{(1-\rho)}$.[10]

From (6), the shape of the (Pareto) profit distribution is related to the cost distribution's power exponent by the relation $\alpha_\pi = \beta\left(\frac{1-\rho}{\rho}\right)$. Thus the profit distribution's form is jointly determined by the taste parameter and the cost distribution. Likewise, since $p = c/\rho$, $F_P(p)$ has a power distribution - and with the same power:

$$F_P(p) = \left(\frac{p}{\underline{p}}\right)^{\beta} \quad \text{for } p \in \left[0, \underline{p}\right], \text{ where } \underline{p} = \underline{c}/\rho.$$

Finally, along similar lines, because from (3) equilibrium output has elasticity $\frac{1}{\rho-1}$, the equilibrium output distribution has a Pareto distribution with tail parameter

$$\alpha_y = (1-\rho)\beta = \alpha_\pi \rho. \tag{8}$$

[10]From here we can find the distribution of the productivity (cost reciprocal), $F_{\hat{C}}(\hat{c})$, where $\hat{c}$ denotes $1/c$. Proceeding as above, but using Lemma 1*(b)* and (from (4)) $r = \Xi(I)\frac{\hat{c}^{\frac{\rho}{1-\rho}}}{D_C}$ so $\hat{c} = kr^{\theta}$ where k takes the value $\left(\frac{D_C}{\Xi(I)}\right)^{\frac{(1-\rho)}{\rho}} > 0$ and θ takes the value $\frac{\rho}{1-\rho} > 0$. Therefore, we write $F_{\hat{C}}(\hat{c}) = F_{kR^\theta}(\hat{c})$ as $F_{\hat{C}}(\hat{c}) = 1 - (\underline{\hat{c}}/\hat{c})^{\alpha_{\hat{c}}}$ for $\hat{c} \geq \underline{\hat{c}}$, where $\alpha_{\hat{c}} = \alpha_\pi\left(\frac{1-\rho}{\rho}\right)$ and the lower bound of the support is $\underline{\hat{c}} = k\underline{r}^{\theta}$.

For the Pareto distribution, densities are always decreasing (the distribution function is concave), so profit, revenue, and output densities must decrease. However, the price or cost ones can be increasing. This can be seen from the relation $\beta = \alpha_\pi \left(\frac{\rho}{1-\rho}\right)$ and noting that the slope of the density of costs has the sign of $\beta - 1$ so that the cost density is decreasing if and only if $\alpha_\pi < \frac{1-\rho}{\rho}$. For example, if the long-tail parameter is $\alpha_\pi = 1$ (as per Axtell, 2001) then the cost density is decreasing if $\rho < 1/2$, i.e. far from close substitutes.[11] And if the "80/20" rule holds, then $\alpha_\pi = 1.16$, with corresponding value $\rho < 0.46$ for decreasing cost density.[12]

The relation in (8) is particular to the Pareto distribution, and does not indicate a more general relation that can be extended to other distributions. However, as seen below and in the next sub-section, the fundamental relationship can be generalized (to a general distribution relation) once we re-express (8) in terms of the elasticities of the corresponding densities. That is, we can determine the elasticities of the output and profit densities in terms of the elasticity of the cost density. Recall that the cost density elasticity density is the same as the price one, so cost and price can be interchanged in the statements below. We prefer to retain costs because they are a primitive to the model.

[11] Axtell (2001) estimates the sales revenue tail parameter as $\alpha_r = .994$ (and recall that $\alpha_\pi = \alpha_r$ for the CES). He estimates the tail parameter for firm size by employee numbers as 1.059, which one might take as a proxy for α_y. Both are estimated for 1997 US Census Bureau data. Intriguingly, (8) then suggests that ρ is close to 1 (indeed, larger!), which stands at odds with empirical estimates of ρ from Broda and Weinstein (2006) and Blonigen and Soderbery (2010). This calls into question the device of proxying output size distribution with employment size distribution, or, indeed, the CES itself.

[12] Values of σ from Broda and Weinstein (2006, Table V) vary from 1.2 for footwear (very differentiated) through 17.1 for crude oil (very homogenous). The corresponding ρ values are 0.17 and 0.94. Coffee comes in at $\rho = 0.6$: the CES would then suggest that unit costs for coffee should be close to uniform if (8) is to hold.

The cost (or price) density elasticity is $\eta_{f_C} = \beta - 1$, while the output and profit density elasticities are $\eta_{f_Y} = -(\alpha_y + 1)$ and $\eta_{f_\Pi} = -(\alpha_\pi + 1)$ respectively.[13] Then we can write (8) as:

$$\eta_{f_Y} = -(1-\rho)\left(\eta_{f_C} + 1 + \frac{1}{1-\rho}\right) \tag{9}$$

and

$$\eta_{f_\Pi} = -\left(\frac{1-\rho}{\rho}\right)\left(\eta_{f_C} + \frac{1}{1-\rho}\right). \tag{10}$$

Notice first that a greater cost density elasticity feeds through to a smaller output or profit density elasticity. Recalling that the latter are negative, this means that those densities become more elastic (more responsive to output or profit levels): more variation in the fundamental variable causes more variation in the induced economic variables.

The impact on these elasticities of the degree of industry product differentiation is rather interesting. The derivatives of each of these expressions with respect to ρ have the sign of $\left(\eta_{f_C} + 1\right)$, which is positive given that η_{f_C} must exceed -1 for the underlying the cost distribution to be increasing in c. This means that the density elasticities are larger, and (being negative) hence actually *less* elastic when there is less product differentiation (higher ρ). This is despite the fact that the equilibrium profit and output functions become more elastic (as expected) when ρ rises, which is because the underlying cost heterogeneity is parlayed into more profit variability when products in the industry are closer substitutes.[14] Nonetheless, profit and output densities (i.e., as functions of profit and output respectively) get less responsive to their

[13] The output and profit ones are elastic (below -1) because the α's are positive. The cost one must exceed -1 (but can be positive) because $\beta > 0$ for the cost distribution to be increasing.

[14] For example, the profit elasticity is $\frac{\rho}{\rho-1}$, so its derivative is negative. The elasticity is negative, and becomes more so as ρ increases.

arguments. This can be seen most clearly with a uniform cost density, so $\eta_{f_C} = 0$. Then the density elasticities are $\eta_{f_Y} = -(2-\rho)$ and $\eta_{f_\Pi} = -\frac{1}{\rho}$ and so higher ρ increases both and they become less responsive.

Combining the two equations above ((9) and (10)) gives the relation between output and profit density elasticities as

$$\eta_{f_\Pi} = \frac{1}{\rho}\eta_{f_Y} + \frac{1-\rho}{\rho}.$$

This indicates that these elasticities are positively related, with the profit one both more leveraged (i.e., its coefficient, $\frac{1}{\rho} > 1$) and more responsive to the output one the smaller is ρ. Moreover, recalling that both elasticities are below -1,[15] a higher ρ entails a higher (less elastic) profit density elasticity: less product differentiation in the sector means a flatter profit density even conditioning on a given output density.

We next provide the density elasticity analysis for general distributions, and show that the relations just derived hold for non-constant elasticities. We then consider some examples.

2.3 Distribution shape relations

While the Pareto imposes decreasing densities, this property is not true generally. We next show the connections between the densities of the satellite distributions by determining the inheritance properties of distributions under the (positive or negative) power transformation. That is, we determine the properties of the CDF of $V = kU^\theta$ (with $k > 0$) that characterizes the relations between distributions for the CES model. The applications we pursue are between cost, output, and profit distributions.

[15] The derivative with respect to ρ has the sign of $-\left(\eta_{f_Y} + 1\right)$, which is positive because $\eta_{f_Y} < -1$.

Proposition 2 *(Inheritance) Let $F_U(u)$ be the CDF of a random variable U and let $V = kU^\theta$ with $k > 0$. Then*

$$\eta_{f_V} = \frac{1}{\theta}\left\{\eta_{f_U}(v) + (1-\theta)\right\}, \tag{11}$$

where $\eta_{f_U}(v)$ is the elasticity of $f_U(u)$ and $u = \left(\frac{v}{k}\right)^{1/\theta}$.

Proof. Consider first $\theta > 0$. Because $F_V(v) = F_{kU^\theta}(v) = F_U\left[\left(\frac{v}{k}\right)^{1/\theta}\right]$, then we have

$$f_V(v) = \frac{1}{\theta}\left(\frac{v}{k}\right)^{1/\theta}\frac{1}{v}f_U\left[\left(\frac{v}{k}\right)^{1/\theta}\right].$$

Differentiating

$$\begin{aligned} f'_V(v) &= \frac{1}{\theta^2}\left(\frac{v}{k}\right)^{2/\theta}\frac{1}{v^2}f'_U\left[\left(\frac{v}{k}\right)^{1/\theta}\right] + \left(\frac{1}{\theta}-1\right)\frac{1}{\theta}\left(\frac{v}{k}\right)^{1/\theta}\frac{1}{v^2}f_U\left[\left(\frac{v}{k}\right)^{1/\theta}\right], \\ \text{or} \qquad \eta_{f_V} &= \frac{1}{\theta}\left\{\left(\frac{v}{k}\right)^{1/\theta}\frac{f'_U\left[\left(\frac{v}{k}\right)^{1/\theta}\right]}{f_U\left[\left(\frac{v}{k}\right)^{1/\theta}\right]} + (1-\theta)\right\}. \end{aligned}$$

Hence

$$\eta_{f_V}(v) = \frac{1}{\theta}\left\{\eta_{f_U}(v) + (1-\theta)\right\},$$

where $\eta_{f_U}(v)$ is the elasticity of the density of U.

Similarly, for the negative power transformation (i.e., when $\theta < 0$) in Lemma 1*(c)*, because $F_V(v) = F_{kU^\theta}(v) = 1 - F_U\left[\left(\frac{v}{k}\right)^{1/\theta}\right]$, we again obtain (11). Both cases ($\theta > 0$ and $\theta < 0$) are covered by the statement in the Lemma. ■

The uniform distribution for U clarifies the role of the other term. Then $\eta_{f_U} = 0$, and so $\eta_{f_V}(v) > 0$: hence, $f'_V(v) > 0$ if and only if $\theta \in (0,1)$. Then an increasing $f_U(u)$ is reinforced, but a decreasing one is offset (and may be overturned). That is, $f'_V(v) \geq 0$ is guaranteed for

$f_U'(u) \geq 0$ and $\theta \in (0,1)$; and $\theta > 1$ is necessary for $f_V'(v) \leq 0$ if $f_U'(u) \geq 0$. For $\theta < 0$, $f_V'(v) < 0$ if $f_U'(u) \geq 0$. For example, uniform costs in CES imply a decreasing profit density (due to convexity of the profit function in c).

More specifically, taking U as cost, then we have for elasticities of output (where $k = \frac{\rho \Xi(I)}{D_C}$ and $\theta = -\frac{1}{1-\rho} < 0$) and profit (where $k = \frac{(1-\rho)\Xi(I)}{D_C}$ and $\theta = -\frac{\rho}{1-\rho} < 0$):

$$\eta_{f_Y} = -(1-\rho)\left(\eta_{f_C} + 1 + \frac{1}{1-\rho}\right) \text{ and } \eta_{f_\Pi} = -\left(\frac{1-\rho}{\rho}\right)\left(\eta_{f_C} + \frac{1}{1-\rho}\right).$$ [16]

These expressions (which imply also a relation between η_{f_Y} and η_{f_Π}) show that the relations given above for the constant elasticity case remain true when elasticities vary. They imply the following slope conditions:

$$f_Y'(y) < 0 \text{ iff } \eta_{f_C} + 1 + \frac{1}{1-\rho} > 0 \text{ and } f_\Pi'(\pi) < 0 \text{ iff } \eta_{f_C} + \frac{1}{1-\rho} > 0.$$ [17]

Notice that these relations (see also (11)) combine technological and taste distribution properties. A decreasing profit density implies a decreasing output density, while an increasing output density is necessary for an increasing profit density. The relations also have implications for properties of distribution modes. For example, if both output and profit densities were unimodal, then the modal profit level is higher than the profit level associated to the modal output. Loosely, the most common profit level is higher than the profit of the most common output level: if this were not true in the data, the data could not have been given from a CES model. Also, note from the relations above that an increasing cost density implies both output and

[16] Or, in terms of the elasticity of substitution, $\sigma\eta_{f_Y} = -\eta_{f_C} - 1 - \sigma$ and $\eta_{f_\Pi} = -\frac{1}{\sigma-1}\left(\eta_{f_C} + \sigma\right)$.

[17] Here the η_{f_C}'s depend on c through y and π respectively. When we compare density shapes below, we are comparing at the y and π values that are compatible through the same c that generates them from (3) and (5).

profit densities are decreasing. Recall though that high costs are associated to low output and profit, so if the cost density is unimodal then the others can also be unimodal.[18]

Proposition 1 above implies that the standard CES model with cost heterogeneity alone cannot deliver (say) Pareto distributions for *both* profit and prices. Indeed, if profit is Pareto distributed, then price must follow a power distribution. We introduce quality heterogeneity in Section 3 to decouple the CES-circle into two or more satellite orbits, and so enrich the associated distribution circles. First, though, we consider other distribution classes for the basic CES relations.

2.4 Other distribution classes

Recent empirical work has not supported power distributions for prices (and hence costs). Two eminent studies with big data on prices have recently appeared: Kaplan and Menzio (2015) and Hitsch, Hortacsu, Lin (2017). While the latter do not directly estimate the shape of price distributions, Kaplan and Menzio (2015) tend to support symmetric bell-shaped densities. The following Figure reproduces their Figure 2a (Kaplan and Menzio, 2015, p.1174) with the empirical distribution approximated by a Normal (green curve):

Notice that the empirical distribution looks closer to a Laplace than a Normal: in what follows we consider a class of densities that includes both as special cases.

On the other hand, studies suggest (e.g., Head, Mayer, and Thoenig, 2014) a Pareto or a log-normal for profit. As we have shown above, a Pareto profit distribution implies a power

[18]For the power cost case, the associated profit and output densities are always decreasing (as should be the case since they are Pareto!) To see this, note that for $F_C(c) = (c/\underline{c})^{\beta}$ then $\eta_{f_c} = \beta - 1$ and so the profit slope condition above becomes $f'_{\Pi}(\pi) < 0$ iff $\beta + \frac{\rho}{1-\rho} > 0$, which must hold since both terms are positive.

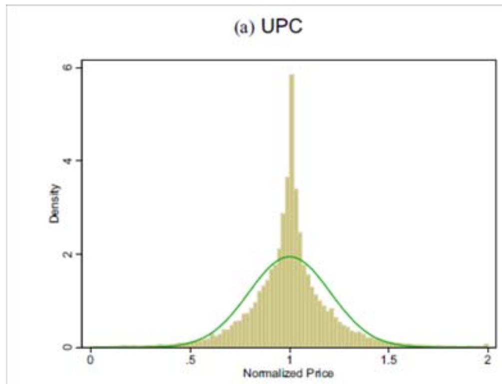

Figure 1:

price distribution under CES. We now go the other direction, and ask what types of distribution are delivered by normal or log-normal price densities (and we pick up log-normal profits along the way). In what follows, it is easier to work directly with the density, with parameters (the K below) determined so that the densities do generate distributions over the relevant support (so K is a normalization factor). In all cases, because the relevant economic variables are non-negative, we have $x \geq 0$.[19]

Consider first the density

$$f_X(x) = Kx^M \exp\left(-\frac{\left|x^B - \xi\right|^A}{\sigma}\right)$$

with parameters $\{M, A, B, \sigma, \xi\}$ and recall we set $x > 0$.[20] This density is continuous, and is

[19] Another nice feature about working directly with densities is that we do not need to figure out how the supports change as we go round the circle. We can simply adjust the constant to make the implied integral a distribution.

[20] This is similar to the double Weibull, which has the form

$$f_X(x) = \frac{c}{2\sigma}\left|\frac{x-\zeta}{\sigma}\right|^{A-1} \exp\left(-\left|\frac{x-\zeta}{\sigma}\right|^A\right).$$

For $A = 1$, we get the Laplace (or the double exponential) density. The double Weibull generalizes the Rayleigh differentiable except at $x = \xi^{1/B} \geq 0$ (although it is differentiable there if $A = 2$). This density has the advantage of nesting the power and Pareto distributions when $\sigma = \infty$ (or $B = 0$). Instead, when $M = 0$ and $B = 1$, it delivers the Normal for $A = 2$ and the Laplace for $A = 1$, and a generalized Normal-Laplace for general A.

Under the CES transformation $Y = kx^{\theta}$, we have the Y density as

$$f_Y(x) = K\left(\frac{y}{k}\right)^{M/\theta} \exp\left(-\frac{\left|\left(\frac{y}{k}\right)^{B/\theta} - \xi\right|^A}{\sigma}\right)$$

which is therefore in the same density class (and so it is in the same distribution class for $\theta > 0$, and in the inverse distribution class for $\theta < 0$). For example, a truncated Normal distribution ($M = 0,\ B = 1,\ A = 2$) for say price, means a truncated exponential power distribution for output.[21] (Truncating the price distribution below entails truncating the output distribution above.)

The log-normal distribution constitutes a power class, and it is self-inverse. To see this property, consider the basic log-normal density

$$f_X(x) = K\exp\left(-\left(\ln x - \xi\right)^2/2\sigma^2\right).$$

Replacing x by $y = kx^{1/\theta}$, as befits the basic power transform fundamental to the CES circle, then gives another log-normal by simply adjusting ξ and σ appropriately, and this is true for θ positive or negative. The immediate implication is that the revenue and productivity

distribution, whose density is $x\exp\left(-x^2\beta\right)$. However, unless $\xi = 0$, the double Weibull does not constitute a power class. However, introducing a power parameter on the x on the RHS rectifies this.

[21] Hence the profit and output densities are proportional to $\exp\left(-\left|x^{\frac{\rho}{\rho-1}} - \xi\right|^2\right)$ and $\exp\left(-\left|x^{\frac{1}{\rho-1}} - \xi\right|^2\right)$ respectively if costs are Normal; if costs are Laplace, the power 2 is removed.

distributions are log-normal if one is (as shown by Head et al., 2014). Furthermore, so are profit and output, and so are price and cost densities (which involve the survivor function distributions because of the flip from high prices associated with low profits, e.g.). Empirically, log-normal has shown up for revenue distributions in Head, Mayer, and Thoenig (2014) among others. A potentially useful generalization that is a power class is to write the generalized log-Laplace-log-Normal, which is

$$f_X(x) = K \exp\left(- \left|\theta \ln kx - \xi\right|^B / 2\sigma^2\right).$$

Here $B = 2$ is log-normal, and $B = 1$ is log-Laplace.[22]

3 CES quality-enhanced model

3.1 General specification

We now extend the model to allow for quality differences across products. Following Baldwin and Harrigan (2011) and Feenstra and Romalis (2014), we rewrite the Representative Consumer's sub-utility functional as $\chi = \left(\int_\Omega z(\omega)^\rho \, d\omega\right)^{1/\rho}$ with $\rho \in (0,1)$ and interpret $z(\omega) = v(\omega) q(\omega)$ as the quality-adjusted consumption of variant ω, where $v(\omega)$ is its quality and $q(\omega)$ is the quantity consumed (as before). We clarify the quality interpretation below in terms of a disaggregate model of individual discrete choice in Section 4. The corresponding demands are:

$$h(p_i, \hat{p}_i) = \frac{\Xi(I)}{p_i} \frac{\hat{p}_i^{\frac{\rho}{\rho-1}}}{\int_{\omega\in\Omega} \hat{p}(\omega)^{\frac{\rho}{\rho-1}} \, d\omega}, \tag{12}$$

[22] Along earlier lines, we could also write a power class that nests power/Pareto and log-normal: $f_X(x) = K x^M \exp\left(- \left|\theta \ln kx - \xi\right|^B / 2\sigma^2\right) x^M \exp\left(- \left|\theta \ln kx - \xi\right|^B / 2\sigma^2\right)$.

where we have defined $\hat{p}_i = p_i/v_i$, which is interpreted as the price per unit of "quality," and $\Xi(I)$ is as above for the three different cases (the amount spent on the differentiated commodity).

The key feature of (12) is that p_i enters both with and without quality. The standard model (1) ensues when all the v's are the same.

With a continuum of firms (as per the usual monopolistic competition set-up), Firm i's equilibrium price solves $\max_{p_i}\left\{\frac{(p_i - c_i)}{p_i}\hat{p}_i^{\frac{\rho}{\rho-1}}\right\}$, so that the pricing solution $p_i = \frac{c_i}{\rho}$ and the Lerner condition (2) still hold. Hence, using $x_i = v_i/c_i$, which we refer to as quality/cost, all firms set the same proportional mark-up, and the equilibrium profit now depends on quality/cost:

$$\pi_i = (1-\rho)\,\Xi(I)\,\frac{x_i^{\frac{\rho}{1-\rho}}}{\int_{\omega\in\Omega} x(\omega)^{\frac{\rho}{1-\rho}}\,d\omega} = (1-\rho)\,r_i. \tag{13}$$

Equilibrium profit is still a fraction $(1-\rho)$ of sales revenue, so their two distributions are in the same multiplicative class. Likewise, (13) implies that profit, sales revenue, and quality/cost distributions are in the same power class (recall that we define two distributions as in the *same power class* if they have the same functional form.)[23]

Price and cost distributions are still in the same multiplicative class as each other, but reciprocal costs and profits are no longer necessarily in the same power class (because one depends on the distribution of c and the other on the distribution of x). How the cost and profit distributions are linked is determined by the link between cost and quality. A functional relation between cost and quality/cost ties this down (and is illustrated in the next sub-section), along with the other distributions on the profit side. However, the output distribution may be

[23] Profits are increasing in x so that firms would like this as large as possible. We can link cost and quality through a type of production function and have (heterogeneous) firms choose their x. More anon.

in a *hybrid* class of its own, because it draws from both the cost and quality/cost distributions (none of these links are explored in Baldwin and Harrigan, 2011, or Feenstra and Romalis, 2014.) Therefore, there are at most three (linked) distribution classes. To summarize:

Proposition 3 *(Decoupling the CES circle) Consider the quality-enhanced CES model of monopolistic competition. Then:*

i) equilibrium price and cost distributions are in the same multiplicative class;

ii) equilibrium profit and sales revenue distributions are in the same multiplicative class;

iii) equilibrium profit, sales revenue, and quality/cost distributions are in the same power class;

iv) equilibrium output distribution is generally in a different power class from (i) and (ii).

We can derive analogous elasticity relations to those in Section 2.2 to link the elasticities of the densities on the third leg.[24] We next illustrate the Proposition with an example which specifies a power relation between quality and cost.

3.2 Constant elasticity quality/cost relation

Suppose that $x = c^{\gamma}$ so that quality/cost is increasing with cost if $\gamma > 0$ (i.e., quality rises faster than cost), and it is decreasing if $\gamma < 0$. The latter case is embodied in the standard CES model above where $\gamma = -1$ and so "better" firms are those with lower costs. The former case effectively corresponds to Feenstra and Romalis (2014).[25] The advantage of the constant

[24] Namely, $\eta_{f_Y} = (1-\rho)\left(\eta_{f_X} - 1 - \frac{1}{1-\rho}\right)$ and $\eta_{f_\Pi} = \left(\frac{1-\rho}{\rho}\right)\left(\eta_{f_X} - \frac{1}{1-\rho}\right)$.

[25] Along the same lines as Feenstra and Romalis (2014), we can let $v = l^{\alpha}$ be the quality produced at cost $wl + \phi$ with ϕ a firm-specific productivity shock, where l is labor input, w is the wage, and $\alpha \in (0,1)$. Maximizing $x = l^{\alpha}/(wl + \phi)$ gives the optimized value relation between cost and quality as $x = \left(\frac{\alpha}{w}\right)^{\alpha} c^{\alpha - 1}$ and so the quality realtion takes a power form. Here it is decreasing (and depends on the fundamental via $x = \phi^{\alpha-1}$).

elasticity relation is that it allows us to deploy results (Lemma 1) on applying power transforms to random variables.

Because profits are proportional to $x_i^{\frac{\rho}{1-\rho}}$ (see (13)), they are proportional to $c_i^{\gamma\frac{\rho}{1-\rho}}$. Hence if $\gamma > 0$, the profit distribution is in the same power class as the cost distribution. So then too are the sales revenue and quality-cost distributions (see (3)). But if $\gamma < 0$, profits, revenues and quality-costs are in the inverse (or "opposite") class.[26] This is the generalization of the earlier standard CES result. Prices, of course, are in the same multiplicative class as costs.

Output is more intricate because it draws its influences from both sides. Indeed, from (12), output is proportional to $x_i^{\frac{\rho}{1-\rho}}/c_i$, which equals $c_i^{\gamma\frac{\rho}{1-\rho}-1}$ under the constant elasticity quality/cost formulation. This implies that for $\gamma > \left(\frac{1-\rho}{\rho}\right)$ the output and cost distributions are in the same power class, while otherwise they are in inverse classes. A summarizing statement:

Proposition 4 *(constant elasticity quality/cost relation) Consider the quality-enhanced CES model of monopolistic competition with $x = c^\gamma$. Then:*

i) equilibrium price and cost distributions are in the same multiplicative class;

ii) equilibrium profits, sales revenue, and quality/cost distributions are in the same power class as unit costs for $\gamma > 0$ and in the inverse class for $\gamma < 0$;

iii) equilibrium output distribution is in the inverse class from unit cost for $\gamma < \left(\frac{1-\rho}{\rho}\right)$, and in the same power class for $\gamma > \left(\frac{1-\rho}{\rho}\right)$.

Take the example of a Pareto distribution for costs. First, prices are also Pareto distributed. Second, profits, revenue, and quality/cost are Pareto distributed for $\gamma > 0$ and power distributed

[26] Recall that one distribution is the *inverse* of another one if it is its survival function.

for $\gamma < 0$ (they are independent of cost if $\gamma = 0$). Third, output is power distributed for $\gamma < \left(\frac{1-\rho}{\rho}\right)$, and Pareto distributed for $\gamma > \left(\frac{1-\rho}{\rho}\right)$.[27] Hence, *we can deliver Pareto distributions for both prices and profits via the constant elasticity quality/cost function.* Pareto revenue and profit distributions are well documented in the literature. Coad (2009) analyzes price distributions for wine and one-week holidays in Majorca, used cars, and London house prices, situations where there is considerable quality differentiation.[28] He finds the resulting distribution is close to Pareto, though less skewed (and more skewed than the lognormal).

Proposition 4(ii) indicates that quality/cost and profit distributions are in the same power class. For example, suppose that the distribution of quality/costs is Pareto: $F_X(x) = 1-(\underline{x}/x)^{\alpha_x}$ and assume that $\alpha_x \frac{1-\rho}{\rho} > 1$. Then the size distribution of profit is Pareto with tail parameter $\alpha_\pi = \alpha_x \frac{1-\rho}{\rho}$. Our result is that the profit tail parameter is the confluence of a preference parameter and a quality/cost distribution one.[29]

The CES is special in many respects, even with quality introduced as above. First, the CES still involves at most three distribution classes (only two for the constant elasticity case, and one is the inverse class of the other). Also, prices are independent of qualities, but percentage cost increases are passed on at equal percentage rates (because $p_i = \frac{c_i}{\rho}$).

4 Consumer Surplus Distributions

So far, we have considered distributions of equilibrium variables on the supply side. We now turn to the distribution of consumer benefits. The CES model is usually presented as a Rep-

[27] If costs are power distributed, Pareto and power are reversed in the above statements.
[28] Note though that for the first two cases the products are sold by a single firm.
[29] Although why they yield the same constant across settings remains intriguing.

resentative Consumer utility which is maximized under an aggregate budget constraint. Given that individual consumers typically consume one or at most a few of the product variants available, the question has arisen whether the representative consumer can capture an aggregate relation of individuals making discrete choices. Anderson, de Palma, and Thisse (1992) have shown how to underpin the CES demand system and the representative consumer's welfare by aggregating heterogeneous individuals' discrete choices. Then, in tune with our current interest in underlying distributions, the pertinent distribution on the consumer side is the distribution of individual consumers' benefits, given that consumers make idiosyncratic choices among the variants available. Here we derive this distribution. We also micro-found how quality enters individual utilities consistent with aggregating to the formulation of the previous section.

4.1 Discrete choice roots for CES

We start with the case without quality, corresponding to Section 2. Suppose that each consumer makes a discrete choice of which variant to buy, based on a conditional (indirect) utility for variant i of the form

$$u_i = m - \ln p_i + \mu\varepsilon_i, \qquad i \in \Omega, \tag{14}$$

where m is individual income, $\mu \geq 0$ is (proportional to) the standard deviation of the individual preference shock (or idiosyncratic draw) ε_i. Each consumer chooses the option for which u_i is greatest. Assuming that the ε_i are i.i.d. Type 1 Extreme Value (or "Gumbel") distributed, the choice probabilities for any option follow the logit formula. However, applying Roy's identity to (14) indicates that the quantity consumed of option i, conditional on preferring it, is $1/p_i$.

Therefore, integrating over the distribution of the ε's, the demand for i is

$$q_i = \frac{1}{p_i} \frac{\exp\left(\frac{-\ln p_i}{\mu}\right)}{\int_{\omega\in\Omega} \exp\left(\frac{-\ln p(\omega)}{\mu}\right) d\omega}$$

or

$$q_i = \frac{1}{p_i} \frac{p_i^{-\frac{1}{\mu}-1}}{\int_{\omega\in\Omega} p(\omega)^{-\frac{1}{\mu}} d\omega}, \tag{15}$$

and hence we have the CES demand model (1) where the parameters are matched by the relation $\mu = \frac{1-\rho}{\rho} = \frac{1}{\sigma-1}$ (or $\rho = \frac{1}{1+\mu}$). Notice how the limit cases concur. If $\mu \to 0$, products are perfect substitutes from (14) and this corresponds to $\rho \to 1$. If $\mu \to \infty$, then idiosyncratic tastes are paramount, and $\rho \to 0$ (the Cobb-Douglas limit).

Because the logit discrete choice model has the IIA property, the common root of the Gumbel distribution for the ε_i implies that the CES is also an IIA demand system. That is, q_i/q_j is independent of p_k, $i \neq j \neq k$, as can be readily verified from (15).

The *aggregate consumer surplus*, $CS = E\left[\max_{i\in\Omega} u_i\right]$, is given by applying to this formulation the well-known log-sum formula of McFadden (1978), written for the continuum case. Using $p(\omega) = c(\omega)/\rho$, this gives

$$\begin{aligned} CS &= \mu \ln\left\{\int_{\omega\in\Omega} \exp\left(\frac{-\ln p(\omega)}{\mu}\right) d\omega\right\} + m \\ &= \mu \ln\left\{\int_{\omega\in\Omega} p(\omega)^{-\frac{1}{\mu}} d\omega\right\} + m, \end{aligned} \tag{16}$$

which we recognize as the consumer surplus for the quasi-linear version of the representative consumer's indirect utility function under the parameter matching $\mu = \frac{1-\rho}{\rho}$.[30] Notice that

[30]To see this, recall first that the representative consumer's utility (from the start of Section 2) is $U =$

applying Roy's identity to the second line of (16) allows us to recover the demand functions.

We are interested in the distribution of surplus for the above formulation. We therefore seek the distribution of the maximum of the u_i as given by (14), for $i \in \Omega$, given that each distribution is a Gumbel.[31] The Gumbel distribution takes the form $F(s;\beta) = \exp(-\exp(\beta - s)/\mu)$, where β is the "location" parameter, which corresponds to the mean (up to a constant), and μ is the "scale" parameter which is proportional to the standard deviation.[32]

Proposition 5 *Let conditional indirect utility for variant i be $u_i = m - \ln p_i + \mu\varepsilon_i$, $i \in \Omega$, with ε_i i.i.d. Gumbel. Then the cdf of the maximum utility in equilibrium, $\max_{i\in\Omega} u_i$, is Gumbel: $F(s;\beta) = \exp[-\exp(\beta - s)/\mu]$, where $\beta = \mu \ln\left\{\int ((1+\mu)\,c(u))^{-\frac{1}{\mu}} f_C(u)\,du\right\} + m$. This value of β also constitutes the expected maximum utility (up to a constant).*

Proof. The key property of the Gumbel is that the maximum draw from Gumbel distributions with the same scale parameter μ is also Gumbel distributed with scale parameter μ. The location parameter of the distribution of the maximum is given by $\beta = CS$ (the expression in (16)). Therefore, the maximum utility achieved is a Gumbel distribution with scale parameter μ and location parameter given by (16), where $p(\omega)$ is replaced by its equilibrium value, $(1+\mu)\,c(\omega)$. The latter equilibrium price follows from (2) with $\rho = \frac{1}{1+\mu}$. Because β is also the mean of the

$\ln\chi + m = \ln\left(\int_\Omega q(\omega)^\rho\,d\omega\right)^{1/\rho} + m$. Then, given prices $p(\omega)$ we have $q_i = \frac{p_i^{\frac{\rho}{\rho-1}-1}}{\int_{\omega\in\Omega} p(\omega)^{\frac{\rho}{\rho-1}}d\omega}$ (see (1)) substituting yields the form in the last line of (16). Indeed, $\frac{1}{\rho}\ln\left(\int_\Omega q(\omega)^\rho\,d\omega\right) = \frac{1}{\rho}\ln\left\{\int_\Omega \frac{p(\omega)^{\frac{\rho}{\rho-1}}}{\left(\int_{\omega\in\Omega} p(\omega)^{\frac{\rho}{\rho-1}}d\omega\right)^\rho}d\omega\right\} = \frac{1-\rho}{\rho}\ln\left(\int_\Omega p(\omega)^{\frac{\rho}{\rho-1}}\,d\omega\right)$.

[31] The Gumbel is a max-stable distribution: it describes the distribution of the maximum for a large class of distributions, of which one is Gumbel itself.

[32] The mean is $\beta + \mu\gamma$ (where γ is Euler's constant), and the standard deviation is $\pi\mu/\sqrt{6}$.

Gumbel distribution, β is also the expected maximum utility (up to a constant $\mu\gamma$, where γ is Euler's constant). ■

Thus β also measures consumer surplus for the population of consumers making discrete choices to maximize (14).

4.2 Quality and consumer surplus distribution

We now show how one can introduce quality into the above formulation in a manner consistent with (12). First, rewrite the individual consumer's conditional (indirect) utility for variant i (14) as

$$u_i = m + \ln v_i - \ln p_i + \mu\varepsilon_i, \qquad i \in \Omega, \tag{17}$$

with the same specification for the error term. Then the consumer surplus cum location parameter for the Gumbel distribution becomes

$$\begin{aligned} CS &= \mu \ln \left\{ \int_{\omega\in\Omega} \exp\left(\frac{\ln v(\omega) - \ln p(\omega)}{\mu} \right) d\omega \right\} + m \\ &= \mu \ln \left\{ \int_{\omega\in\Omega} \left(\frac{p(\omega)}{v(\omega)} \right)^{-\frac{1}{\mu}} d\omega \right\} + m. \end{aligned} \tag{18}$$

The second line gives demands (using Roy's identity), as well as the Representative Consumer's indirect utility.

Proposition 6 *Let conditional indirect utility for variant i be $u_i = m + \ln v(\omega) - \ln p_i + \mu\varepsilon_i$, $i \in \Omega$, with ε_i i.i.d. Gumbel. Then the cdf of the maximum utility in equilibrium, $\max_{i\in\Omega} u_i$, is Gumbel with cdf $F(s;\beta) = \exp[-\exp(\beta - s)/\mu]$, where $\beta = \mu \ln \left\{ \int \left(\frac{x(u)}{(1+\mu)} \right)^{\frac{1}{\mu}} f_X(u)\, du \right\} + m$. This value of β also constitutes the expected maximum utility (up to a constant).*

Proof. The proof parallels that of the previous Proposition. The maximum utility follows a Gumbel distribution with scale parameter μ and location parameter given by (18). Once again, $p(\omega)$ is replaced in (18) by its equilibrium value, $(1+\mu)\,c(\omega)$ (as argued in Section 3, price is independent of quality). Hence $\frac{p(\omega)}{v(\omega)}$ in (18) is given in equilibrium as $(1+\mu)/x(\omega)$, where we recall that $x(\omega) = v(\omega)/c(\omega)$. ■

The distribution depends only on the summary statistic $\int (x(u))^{\frac{1}{\mu}} f_X(u)\,du$. This statistic constitutes a "power mean." If $\mu = 1$ it is the simple mean of the distribution. Note that one distribution, $F(s;\beta_1)$, of consumer utilities first-order stochastically dominates another, $F(s;\beta_2)$, if $\int (x(u))^{\frac{1}{\mu}} f_X(u)\,du$ is lower for the second ($\beta_1 > \beta_2$). Clearly the FOSD property implies that the mean consumer surplus is lower for the second.

However, the FOSD property does not mean that each individual consumer is better off when the summary statistic rises. It means that (e.g.) the top (or bottom) 10% of consumers are better off than the top (or bottom) 10% of consumers were before the change. To see this point, the discrete choice model treats the variant draws ε_i as fixed. Then a change in the distribution $F_X(x)$ may make some consumers worse off even though the statistic goes up: those with high draws do not change their product selection, and therefore they are worse off if their preferred product suffers a fall in its x, while others rise enough so as to raise the aggregate statistic.[33]

[33] For example, suppose the values of x were high and low, and the lower one went down while the higher one went up, while raising the index β. Then those consumers still buying the low-level products would be worse off.

5 Conclusions

The CES model has been the workhorse model of monopolistic competition with asymmetric firms and the central distribution in the literature has been the Pareto. We show that all relevant distributions are Pareto if any one is (caveat: for prices and costs it is the distribution of the reciprocal that is Pareto). This result we term the *Pareto circle*. To put this another way, suppose we posit that productivity (the reciprocal of costs) is Pareto distributed (equivalently, costs have a power distribution). Then, so does the reciprocal of prices follow a Pareto distribution, and the other variables (output, revenue, and profit) are all Pareto distributed. It is not possible to have (for example) a Pareto distribution for profits and (another) Pareto distribution for prices in the CES model. The Pareto circle cannot be escaped if one element is Pareto. Similar results hold for other distributions, yielding a more general *CES circle*: the assumed distribution of productivity (cost reciprocal) is also the equilibrium distribution of outputs, profits, etc.

This analysis allows us to determine the relations between equilibrium distributions. These are simple for the core Pareto distribution, but they turn out to be surprisingly general once couched as relations between elasticities of equilibrium densities. Put another way, the Pareto results form a solid benchmark for broader distribution relations. The relations (between equilibrium distributions) tell us when a decreasing density for one variable implies a decreasing density for another, etc., and so describes how shapes of densities are all related to each other.

The CES cycle is broken by allowing for a further dimension of firm heterogeneity. Following Baldwin and Harrigan (2011) and Feenstra and Romalis (2014), we introduce "quality" heterogeneity. We link the quality and cost distributions via a function that writes quality as a function of cost. Doing this then enables us to get two linked groups of distributions. In one group are profit and revenue, and in the other are costs and prices, while output draws from both these groups. Our leading example is a quality/cost function that can deliver Pareto distributions in each group.

Our final contribution is to uncover the distribution of the welfare of heterogeneous consumers making discrete choices across variants and consistent with the CES representative consumer. This distribution is a Gumbel (extreme value, Type 1), and its equilibrium value depends on a simple statistic conflating the cost (or quality/cost, in the generalization) with the preference scale.

References

[1] Anderson, Chris (2006). The Long Tail: Why the Future of Business Is Selling Less of More. New York: Hyperion.

[2] Anderson, Simon and André de Palma (2000), From local to Global Competition, European Economic Review, 44, 423-448.

[3] Anderson, Simon P., André de Palma, and Jacques François Thisse (1992). Discrete Choice Theory of Product Differentiation. Cambridge, MA: MIT Press.

[4] Axtell, Rob (2001). Zipf distribution of U.S. firm sizes. Science 293, 1818–20.

[5] Baldwin, Richard, and James Harrigan (2011). Zeros, Quality, and Space: Trade Theory and Trade Evidence. American Economic Journal: Microeconomics 3(2), 60-88.

[6] Bertoletti, Paolo, and Federico Etro (2017). Monopolistic competition when income matters. Economic Journal, forthcoming.

[7] Blonigen, Bruce A. & Soderbery, Anson (2010). Measuring the Benefits of Product Variety with an Accurate Variety Set. Journal of International Economics, 82, 168-180.

[8] Broda, Christian and David E. Weinstein (2006). Globalization And The Gains From Variety. Quarterly Journal of Economics, 121(2), 541-585.

[9] Cabral, Luis , and Jose Mata (2003). On the evolution of the firm size distribution: Facts and theory. American Economic Review, 93(4), 1075-1090.

[10] Coad, Alex (2009). On the distribution of product price and quality. Journal of Evolutionary Economics, 19(4), 589–604.

[11] Chamberlin, Edward (1933). Theory of Monopolistic Competition. Cambridge, MA: Harvard University Press

[12] Crozet, Matthieu, Keith Head, and Thierry Mayer (2011). Quality Sorting and Trade Firm-level Evidence for French Wine. Review of Economic Studies, 79(2), 609-644.

[13] De Loecker, Jan, and Frederic Warzynski (2012). Markups and Firm-Level Export Status. American Economic Review. 102(6), 2437-71.

[14] di Giovanni, A. Levchenko, and I. Mejean (2014) Firms, Destinations, and Aggregate Fluctuations. Econometrica, 82-4, 1303-1340

[15] Dixit, Avinash K., and J. S. Stiglitz (1977). Monopolistic Competition and Optimal Product Diversity. American Economic Review, 67(3), 297-308.

[16] Dhingra, Swati and John Morrow (2017) Monopolistic competition and optimum product diversity under firm heterogeneity, Journal of Political Economy, forthcoming.

[17] Fajgelbaum, Pablo D., Gene M. Grossman, and Elhanan Helpman (2011). Income Distribution, Product Quality, and International Trade. Journal of Political Economy, 119(4), 721-765.

[18] Feenstra, Robert C. (1994), New Product Varieties and the Measurement of. International Prices, American Economic Review, 84(1), 157-177.

[19] Feenstra, Robert C. and John Romalis (2014). International Prices and Endogenous Quality, Quarterly Journal of Economics, 129(2), 477-527.

[20] Gabaix, Xavier (2016). Power Laws in Economics: An Introduction. Journal of Economic Perspectives. 30 (1), 185-206.

[21] Head, Keith, Thierry Mayer, and Mathias Thoenig (2014). Welfare and Trade without Pareto. American Economic Review 104(5): 310-16.

[22] Helpman, Elhanan and Paul R. Krugman (1985) Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy. MIT Press, Cambridge, Mass.

[23] Helpman, Elhanan, Marc J. Melitz, and Stephen R. Yeaple (2004) Export Versus FDI with Heterogeneous Firms. American Economic Review, 94(1): 300-316.

[24] Hitsch, Günter J., Ali Hortacsu, and Xiliang Lin (2017) Prices and Promotions in U.S. Retail Markets: Evidence from Big Data. Booth Discussion Paper 17-18, University of Chicago.

[25] Kaplan, Greg and Guido Menzio (2015) The Morphology Of Price Dispersion. International Economic Review, 56, 1165-1206.

[26] Mrázová, Monika, J. Peter Neary, and Mathieu Parenti (2017). Sales and Markup Dispersion: Theory and Empirics. CEPR Discussion Paper 12044.

[27] McFadden, Daniel (1978) Modelling the Choice of Residential Location. in A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull (eds.), Spatial Interaction Theory and Planning Models, Amsterdam: North Holland, 75-96, Reprinted in J. Quigley (ed.), 1997, The economics of housing, Vol. I, Edward Elgar: London, 531-552.

[28] Melitz, Marc J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. Econometrica, 71(6), 1695-1725.

[29] Nigai, Sergey (2017) A tale of two tails: Productivity distribution and the gains from trade. Journal of International Economics, 104, 44-62.

[30] Nocke, Volker and Stephen Yeaple (2008) Firm heterogeneity and the structure of US multinational activity. Review of Economic Studies, 75(2), 529-557.

[31] Pareto, Vilfredo (1965). La Courbe de la Répartition de la Richesse (Original in 1896): Oeuvres Complètes de Vilfredo. Ed. Busino G. Pareto. Geneva: Librairie Droz, 1-5.

[32] Zhelobodko, Evgeny, Sergey Kokovin, Mathieu Parenti, and Jacques-François Thisse (2012). Monopolistic competition: Beyond the constant elasticity of substitution. Econometrica 80(6), 2765-2784.

Economic distributions and primitive distributions in Monopolistic Competition

Simon P. Anderson and André de Palma*

Revised February 2018

Abstract

We link fundamental technological and taste distributions to endogenous economic distributions of prices and firm size (output, profit). We provide constructive proofs to recover the demand structure, mark-ups, and distributions of cost, price, output and profit from just two distributions (or from demand and one distribution). A continuous logit demand model illustrates: exponential (resp. normal) quality-cost distributions generate Pareto (log-normal) economic size distributions. Pareto prices and profits are reconciled through an appropriate quality-cost relation.

JEL Classification: L13, F12

Keywords: Primitive and economic distributions, monopolistic competition, pass-through and demand recovery, price and profit dispersion, Pareto and log-normal distribution, CES, Logit.

*Department of Economics, University of Virginia, USA, sa9w@virginia.edu; CES-University Paris-Saclay, FRANCE. andre.depalma@ens-cachan.fr. The first author gratefully acknowledges research funding from the NSF. We thank Isabelle Mejean, Julien Monardo, Maxim Engers, Farid Toubal, James Harrigan, and Ariell Resheff for valuable comments, and seminar participants at Hebrew University, Melbourne University, Stockholm University, KU Leuven, Laval, Vrij Universiteit Amsterdam, and Paris Dauphine. This is a revised version of CEPR DP 10748.

1 Introduction

Distributions of economic variables have attracted the interest of economists at least since Pareto (1896). In industrial organization, firm size distributions (measured by output, sales, or profit) have been analyzed, while different studies have looked at the distribution of prices within an industry. Firm sizes (profitability, say) within industries are wildly asymmetric, and frequently involve a long-tail of smaller firms. The idea of the long tail has recently been invoked prominently in studies of Internet Commerce (Anderson, 2006, Elberse and Oberholzer-Gee, 2006), and particular distributions – mainly the Pareto and log-normal – seem to fit the data well in other areas too (see Head, Mayer, and Thoenig, 2014). In international trade, recent advances have enabled studying distributions of sales revenues (see, e.g., Eaton, Kortum, and Kramarz, 2011). The distributions of these "economic" variables are (presumably) jointly determined by the fundamental underlying distributions of tastes and technologies. In this paper we determine the links between the various distributions. We link the economic ones to each other and to the primitive distributions and tastes. Moreover, the primitives can be uncovered from the observed economic distributions.

Philosophically, the paper closest (and complementary) to ours is Mrázová, Neary, and Parenti (2016). These authors also study the relations between equilibrium distributions of sales and mark-ups, the primitive productivity distribution, and (a specific) demand form (although they do not include heterogeneous quality). They are mainly interested in when distributions are in the same ("self-reflecting") class (e.g., when both productivity and sales are log-normal or Pareto). They also provide some empirical analysis of log-normal and Pareto distributions.

We start by deploying a general monopolistic competition model with a continuum of firms (see Thisse and Uschev, 2016, for a review of this literature). We first show how the demand function delivers a mark-up function, and then we show our key converse result that the mark-up (or "pass-through" function of Weyl and Fabinger, 2013) determines the form of the demand function. We next engage these results with constructive proofs to show how cost and price distributions suffice to determine the shape of the economic profit and output distributions and the demand form. Along broader lines, we show when and how any two elements (e.g., two distributions) suffice to deliver all the missing pieces.

Allowing for both quality and cost heterogeneity,[1] we show a three-way relation between two groups of distributions and the quality-to-cost relation: knowing one element from any two of these ties down the third. On one leg, we generate the relation between equilibrium profit dispersion, firm outputs, and the fundamental quality-cost distribution. On a second leg, we show the relation between the cost distribution and equilibrium price dispersion. If we know demand, then knowing any one of the distributions on one leg suffices to determine the others on that leg. Moreover, knowing a distribution from each leg allows us to determine what the relation between cost and quality must be on the third leg. If the demand form is not known, then we show that it can be deduced from observing price, output, and profit distributions (and the cost distribution and the relation between costs and quality can also be determined).

We next develop and deploy a logit model of monopolistic competition.[2] The logit is the

[1] Ironically, Chamberlin (1933) is best remembered for his symmetric monopolistic competition analysis. Yet he went to great length to point out that he believed asymmetry to be the norm, and that symmetry was a very restrictive assumption. We model both quality and production cost differences across firms.

[2] The Logit is an attractive alternative framework to the CES. Anderson, de Palma, and Thisse (1992) have shown that the CES can be viewed as a form of Logit model.

workhorse model in structural empirical IO. Some useful characterization results are that normally distributed quality-costs induce log-normal distributions for profits, and that an exponential distribution of quality-costs leads to a Pareto distribution for profit. Cost heterogeneity alone cannot induce Pareto distributions for both profits and prices. We show by construction for the logit example that the added dimension of quality (and the associated quality-cost relation) can generate Pareto distributions for both, thus allowing sufficient richness to link diverse distribution types.

2 Links between distributions

A continuum of firms produce substitute goods. Each has constant unit production costs, but these differ across firms. With a continuum of firms, each firm effectively faces a monopoly problem where the price choice is independent of the actions of rivals. In this spirit, we allow for a general demand formulation, and show how the primitive (demand and cost distribution) feed through to the endogenous economic distributions and variables.

We first give the demand model, and derive the equilibrium mark-up schedule in Lemma 1 as a function of firm unit cost, c. Our proofs here and beyond are constructive: we derive the relations between distributions and primitives. Theorem 1 shows how the economic distributions are linked to the demand form and the cost distribution. Throughout, we make explicit the appropriate monotonicity conditions.

We shall assume for the exposition that all distributions are absolutely continuous and strictly increasing. As should become apparent, any gaps in a distribution's support will correspond to gaps in supports of the other distributions; the analysis applies piecewise on the interior of the supports. Likewise, mass-points in the interior of the support pose no problem because they correspond to mass points in the other distributions.

2.1 Demand and mark-ups

Assumption 1 *Suppose that demand for a firm charging p is*

$$y = h(p), \tag{1}$$

a positive, strictly decreasing, strictly (-1)-concave, and twice differentiable function.[3]

We suppress for the present the impact of other firms' actions on demand, which would be expressed as aggregate variables in the individual demand function. Under monopolistic competition with a continuum of firms, each firm's individual action has no measurable impact on the aggregate variables.[4] Because we look at the cross-section relation between equilibrium distributions, the actions of other firms are the same across the comparison, and therefore have no bearing on our results. We return to this when we discuss specific examples. The profit for a firm with per unit cost c is $\pi = (p-c)\,h(p) = mh(m+c)$, where $m = p - c$ is its mark-up. With a continuum of firms (monopolistic competition), the equilibrium mark-up satisfies

$$m = -\frac{h(m+c)}{h'(m+c)}. \tag{2}$$

Lemma 1 *Under Assumption 1, the equilibrium mark-up, $\mu(c) > 0$ is the unique continuously differentiable solution to (2), with $\mu'(c) > -1$. $\mu'(c) \geq 0$ if $h(.)$ is log-convex and $\mu'(c) \leq 0$*

[3] This is equivalent to $\frac{1}{h(.)}$ strictly convex, and is a minimal condition ensuring a maximum to profit. See Caplin and Nalebuff (1991) and Anderson, de Palma, and Thisse (1992, p.164) for more on ρ-concave functions; and Weyl and Fabinger (2013) for the properties of pass-through as a function of demand curvature.

[4] For example, the "price index" in the CES model, or the Logit denominator.

if $h(.)$ is log-concave. The associated equilibrium demand, $h^*(c) \equiv h(\mu(c)+c)$, is strictly decreasing and continuously differentiable. The equilibrium profit function, $\pi^*(c) = \mu(c)\,h^*(c)$, is strictly convex and twice continuously differentiable with $\pi^{*\prime}(c) = -h^*(c) < 0$.

Proof. The solution to (2), denoted $\mu(c)$, is uniquely determined (and strictly positive) when the RHS of (2) has slope less than one, as is implied by $h(.)$ being strictly (-1)-concave. Applying the implicit function theorem to (2) shows that

$$\mu'(c) = \frac{-\left(\frac{h(m+c)}{h'(m+c)}\right)'}{1+\left(\frac{h(m+c)}{h'(m+c)}\right)'} > -1, \tag{3}$$

where the denominator is strictly positive under Assumption 1.[5] The numerator is (weakly) positive for h log-convex and (weakly) negative for h log-concave. Let $h^*(c) = h(\mu(c)+c)$ denote the value of $h(.)$ under the profit-maximizing mark-up. Then, $h^*(c)$ is strictly increasing, as claimed, because

$$dh^*(c)/dc = (\mu'(c)+1)\,h'(\mu(c)+c) < 0 \tag{4}$$

and given that $\mu'(c) > -1$. Finally, $\pi^*(c) = \mu(c)\,h^*(c)$ is strictly decreasing with $\pi^{*\prime}(c) = -h^*(c) < 0$ by the envelope theorem. $\pi^*(c)$ is twice continuously differentiable because $h^*(c)$ is continuously differentiable, and strictly convex because $h^*(c)$ is strictly decreasing. ■

Notice that $h^*(c)$ is the inverse marginal revenue curve. Because marginal revenue slopes down strictly, $h^*(c)$ is a continuous function. The result that $\pi^{*\prime}(c) = -h^*(c)$ is the analogue (for monopoly) to Hotelling's Lemma. Notice that the property $\mu'(c) > -1$ is just the standard property that price never goes down as costs increase. As the next Corollary stresses, continuity

[5] When $h(u)$ is strictly (-1)-concave, then $h(u)\,h''(u) - 2\left[h'(u)\right]^2 < 0$, which rearranges to $\left[\frac{h(u)}{h'(u)}\right]' > -1$.

of $\mu'(c)$ implies equilibrium price is a continuously differentiable and, because $\mu'(c) > -1$ by (3) under A1, it is a strictly increasing function of cost.

Corollary 1 *Under Assumption 1, equilibrium price is a strictly increasing and continuously differentiable function of cost, c.*

The key implication of this Corollary and Lemma 1 is that we can rely on monotonic relations between variables, which is crucial in twinning distributions (as we do below). The firms with costs higher than some value c are the same ones that have prices higher than p, an output below y and a profit below π, where the specific values satisfy $\pi = (p-c)\,h(p)$ (and the mark-up $(p-c)$ satisfies (3)). That is, letting z denote the fraction of firms with profit below some level π, we have

$$1 - F_C(c) = 1 - F_P(p) = F_Y(y) = F_\Pi(\pi) = z. \tag{5}$$

Some characterization results rely on a delineation of the degree of curvature of demand:

Corollary 2 *Under Assumption 1, if demand is strictly log-concave (resp. strictly log-convex), higher cost firms have lower (resp. higher) equilibrium markups ($\mu'(c) < 0$, resp. $\mu'(c) > 0$).*

In the log-concave case, low-cost firms use their advantage in both mark-up and output dimensions. Under log-convexity, low-cost firms exploit the opportunity to capitalize on much larger demand by setting small mark-ups. In both cases though, as per Lemma 1, profits are higher. The only demand function with constant (absolute) mark-up is the exponential (associated to the Logit), which has $h(\cdot)$ log-linear in p, and so $\frac{h(m+c)}{h'(m+c)}$ is constant. For $h(\cdot)$ strictly log-concave, $\mu'(c) < 0$, so firms with higher costs have lower mark-ups in the cross-section of firm types (price pass-through is less than 100%). They also have lower equilibrium outputs. When $h(\cdot)$ is strictly log-convex, the mark-up *increases* with c, so cost pass-through is greater than 100%, which is a hall-mark of CES demands, which have constant elasticity and hence constant relative mark-up.[6]

These properties indicate properties of the price distribution relative to the cost distribution. The price distribution is a *compression* of the cost distribution when h is log-concave, and a *magnification* when h is log-convex, in the simple sense that prices are closer together (or, respectively, farther apart) than costs. The border case (Logit / log-linear demand) has constant mark-ups, so the price distribution mirrors the cost one.

An important special case is when demand is $\rho-$linear (which means that h^ρ is linear). Suppose then that

$$h(.) = (1 + (k-p)\rho)^{1/\rho}, \tag{6}$$

where k is a constant. Then

$$\mu(c) = \frac{1+\rho(k-c)}{1+\rho},$$

which is linear in c.[7] For $\rho = 1$ demand is linear and the standard property is apparent that mark-ups fall fifty cents on the dollar with cost. Log-linearity is $\rho = 0$ (note that $\lim_{\rho\to 0} h(.) = \exp(k-p)$) and delivers a constant mark-up. For $\rho-$linear demands, equilibrium demand is $h^*(c) = \left(\frac{1+\rho(k-c)}{1+\rho}\right)^{1/\rho}$ and then (see (11) below) $\frac{dh^*(c)/dc}{h^*(c)} = \frac{-1}{1+\rho(k-c)} = -\frac{\mu'(c)+1}{\mu(c)} < 0$. Notice that

[6] So a 1% cost rise causes equilibrium price to rise by 1%.

[7] More generally, $\mu'(c) \geq \frac{-\rho}{1+\rho}$ when h is ρ-convex and $\mu'(c) \leq \frac{-\rho}{1+\rho}$ when h is ρ-concave.

$h^*(c)$ is also ρ-linear.

2.2 Equilibrium distributions

The relations above already determine some links between the equilibrium price distribution and the cost distribution and demand. We now show how the other economic distributions are determined and linked in the model.

Our analysis makes extensive use of the following result.

Lemma 2 *Consider two distributions $F_{X_1}(x_1)$ and $F_{X_2}(x_2)$, which are absolutely continuous and strictly increasing on their respective domains. Let X_1 and X_2 be related by a monotone function $X_1 = \xi(X_2)$. Then $F_{X_2}(x_2) = F_{X_1}(\xi(x_2))$ for $\xi(.)$ increasing, and $F_{X_2}(x_2) = 1 - F_{X_1}(\xi(x_2))$ for $\xi(.)$ decreasing.*

Proof. For $\xi(.)$ increasing, $F_{X_1}(x_1) = \Pr(X_1 < x_1) = \Pr(\xi(X_2) < x_1) = \Pr\left(X_2 < \xi^{-1}(x_1)\right) = F_{X_2}\left(\xi^{-1}(x_1)\right)$. Equivalently, $F_{X_2}(x_2) = F_{X_1}(\xi(x_2))$. For $\xi(.)$ decreasing, $F_{X_1}(x_1) = \Pr(\xi(X_2) < x_1) = \Pr\left(X_2 > \xi^{-1}(x_1)\right) = 1 - F_{X_2}\left(\xi^{-1}(x_1)\right)$; equivalently, $F_{X_2}(x_2) = 1 - F_{X_1}(\xi(x_2))$. ■

We can now turn to the equilibrium analysis. Figure 1 illustrates. The upper right panel gives the demand curve, from which we determine the corresponding marginal revenue function. The latter is the key to finding the output distribution from the cost distribution. Notice that $h^*(c)$ defined above determines the equilibrium output (for a firm with per unit cost c) as a function of its cost. As earlier noted, the inverse function, $c = h^{*-1}(y)$ therefore traces out the marginal revenue curve.

The distribution of costs is given in the upper left panel. The negative linear relation between the cost and output distributions is given in the lower left panel: as noted in Lemma

1, higher costs are associated to lower outputs. Therefore, the $z\%$ of firms with costs below c are the $z\%$ of firms with output above $y = h^*(c)$. We hence choose some arbitrary level $z \in (0,1)$ (see (5)). This means that all firm types with cost levels above $c(z) = F_C^{-1}(1-z)$ are the firms with outputs and profits below y and π. That is, $1 - F_C(c) = F_Y(h^*(c))$ $(= z)$. The lower right panel therefore connects this relation as the output distribution, $F_Y(y)$. (Notice that in the above argument, only the marginal revenue curve was used from the demand side: as we show later in Section 5, the cost and output distribution determine the marginal revenue, but we then need to integrate up to find demand).

Figure 1 also provides information to determine the price distribution. The upper right panel gives the vertical distance between the marginal revenue and demand, which is the mark-up (which can be expressed as $\mu(c)$), and is thus the vertical shift between cost and price distributions in the upper left panel. It can be constructed simply from the information in the top two panels[8] by drawing across the demand price associated to a marginal revenue - marginal cost intersection. We could also draw in the mark-up distribution in the upper left panel, but have avoided the extra clutter. Notice that (as drawn) the price and cost distributions diverge, as is consistent with Lemma 1 for increasing $\mu(c)$, i.e., log-concave demand.

In summary then, the marginal revenue curve $h^{*-1}(y)$ together with the cost distribution ties down the output distribution (and conversely, for when we shall later be interested). The demand function then finds the price distribution, and therefore relates price and output distributions.

One relation that is missing in the Figure is the profit distribution. But, as Lemma 1 shows,

[8] Hence we were able to give results on the relationships between cost and price distributions at the start of this section without reference to the output distribution.

analogous arguments apply: $\pi^*(c)$ is a decreasing function and so the relation $1 - F_C(c) = F_\Pi(\pi^*(c))$ $(= z)$ can be used to construct the profit distribution.

The following result establishes the existence of a unique equilibrium for the monopolistic competition model. Consequently, equilibrium distributions are tied down from the primitives on costs and demand.

Theorem 1 *(existence of unique equilibrium for monopolistic competition model) Let there be a continuum of firms, with demand (1) satisfying Assumption 1. Let F_C be strictly increasing and twice differentiable on its support. Then the distributions F_P, F_Y, and F_Π are strictly increasing and twice differentiable on their supports and given by $F_P(p) = F_C(c(p))$; $F_Y(y) = 1 - F_C(h^{*-1}(y))$; and $F_\Pi(\pi) = 1 - F_C(\pi^{*-1}(\pi))$, where $c(p)$ inverts $p(c)$, $h^{*-1}(y)$ inverts $h^*(c)$, and $\pi^{*-1}(\pi)$ inverts $\pi^*(c)$.*

Proof. Let $p(c)$ denote the equilibrium price for a firm with cost c; from (3) we have $\mu'(c) > -1$ so that $p(c)$ is strictly increasing, and define the inverse relation as $c(p)$, which is strictly increasing. The relation $p(c)$ (and hence its inverse) is determined from $h(.)$ by Lemma 1.

Given F_C, then $F_P(p)$ is determined by $F_P(p) = F_C(c(p))$. Next, consider $F_Y(y)$. By result (4) we know that output $y = h^*(c)$ is a monotonic decreasing function, and so (by Lemma 2) the fraction of firms with output below $y = h^*(c)$ is the fraction of firms with cost above c, so $F_Y(h^*(c)) = 1 - F_C(c)$, or indeed

$$F_Y(y) = \Pr(h^*(C) < y) = \Pr\left(C > h^{*-1}(y)\right) = 1 - F_C\left(h^{*-1}(y)\right). \qquad (7)$$

Finally, by Lemma 1 we know that profit $\pi^*(c) = \mu(c)\,h^*(c)$ is a strictly decreasing function, and so the fraction of firms with profit below $\pi^*(c)$ is the fraction of firms with costs above c, so $F_\Pi(\pi^*(c)) = 1 - F_C(c)$, or

$$F_\Pi(\pi) = \Pr(\Pi < \pi) = \Pr(\pi^*(C) < \pi) = \Pr\left(C > \pi^{*-1}(\pi)\right) = 1 - F_C\left(\pi^{*-1}(\pi)\right). \tag{8}$$

■

The key relation underlying the twinning of distributions is the decreasing relation between cost and output, profit, and price (see Lemma 1 and Corollary 1). A specific cost distribution generates a specific output, profit, and price distribution. Conversely, as we show in the next result, this output or profit distribution could only have been generated from the initial cost distribution. These links are exploited below in Section 3, where we show how the properties of the distributions feed through to each other in terms of their shapes, and we show what are the restrictions among the admissible distributions.

Researchers often impose specific demand functions (such as CES, or logit). Here we forge the (potentially testable) empirical links that are imposed by so doing: Theorem 1 shows that when a specific functional form is imposed for h (as is done in most of the literature), then all the relevant distributions can be found from $F_C(c)$. Furthermore, all distributions can be found from just one of them.

Theorem 2 *Let there be a continuum of firms, with demand (1) satisfying Assumption 1. Consider the set of 4 distributions, $\{F_C, F_P, F_Y, F_\Pi\}$. Suppose that any one is known and is strictly increasing and twice differentiable on its support. Then all other distributions in the set are explicitly recovered and all are strictly increasing and twice differentiable functions on their* *supports.*

Proof. First, F_C was covered in Theorem 1. So consider now F_P. Then $F_C(c) = F_P(p(c))$, where $p(c)$ is the equilibrium price relation, which we showed in Corollary 1 to be continuously differentiable, and both the other distributions are determined from the steps in the proof of Theorem 1 earlier.

Next start with F_Y. Because $h(p)$ is strictly decreasing, then F_P is determined by $F_P(p) = 1 - F_Y(h(p))$. By the argument above, F_C is then determined, and hence so is $F_\Pi(\pi)$.

Finally, suppose that we start with F_Π. By Lemma 1 we know that profit $\pi^*(c) = \mu(c)\, h^*(c)$ is a strictly decreasing function. Therefore $F_C(c)$ is recovered from $F_C(c) = 1 - F_\Pi(\pi^*(c))$. From Theorem 1, F_P is recovered, and so is F_Y. ■

The Theorem says that for any (-1)- concave demand function and any potential economic distribution, there is only one cost distribution that is consistent with the economic distribution. The other economic distributions are likewise pinned down.

Later on we turn our attention to pairs of distributions that are not consistent with the monopolistic competition model; that is, which pairs would indicate violation of the model. Conversely, for admissible pairs, we show how the implicit demand function is determined.

2.3 Atoms and gaps

Some remarks are in order about relaxing the assumptions made in the last two Theorems. There are two main issues with distributions; gaps in the support, and spikes. On the demand side, we address failures of (-1)-concavity.

If $F_C(c)$ has an atom, then $F(p)$ and the other two economic distributions have corresponding atoms of the same size. Likewise, if $F_C(c)$ has a gap, then the three economic distributions have corresponding gaps.

If $h(p)$ has a kink down at some price, while $F_C(c)$ remains continuous, then $F_P(p)$ and $F_Y(y)$ have atoms corresponding to the kink (a range of costs are associated to the same price and output) while the the profit distribution remains continuous.

If $h(p)$ is not (-1)-concave over some range, the corresponding marginal revenue curve slopes up. As a function of c, equilibrium price jumps down (and equilibrium output jumps up) so that $F_P(p)$ and $F_Y(y)$ have corresponding gaps, while $F_\Pi(\pi)$ does not.

Conversely, $F_P(p)$ and $F_Y(y)$ have gaps, while $F_C(c)$ does not, then $h(p)$ is not (-1)-concave over some part of the intervening range, etc. Therefore, such behavior of the distributions can still be consistent with the monopolistic competition model, although not under Assumption 1 and a continuous $F_C(c)$.

3 shapes of things (and inheritance properties)

We take two complementary perspectives on describing how distribution shapes are related to each other. The first is in terms of the degree of concavity that is inherited from other distributions. The second is the relationship between elasticities of distributions and densities. These are crisply expressed via elasticities of the other pertinent economic variables. The latter are expressed as various demand-side statistics.

3.1 Distribution ρ-concavity/convexity inheritance properties

The import of the next result is that we can determine how curvature properties of one distribution carry over to a related one, and vice versa.

Lemma 3 *Consider two functions $F_{X_1}(x_1)$ and $F_{X_2}(x_2)$, which are absolutely continuous and strictly increasing on their respective domains. Let X_1 and X_2 be related by an increasing function $X_1 = \xi(X_2)$. Then:*

a) if $\xi(X_2)$ is concave, $F_{X_2}(x_2) = F_{X_1}(\xi(x_2))$ is a ρ-concave function if $F_{X_1}(x_1)$ is ρ-concave.

b) if $\xi(X_2)$ is convex, $F_{X_2}(x_2) = F_{X_1}(\xi(x_2))$ is a ρ-convex function if $F_{X_1}(x_1)$ is ρ-convex.

If $F_{X_1}(x_1)$ is strictly decreasing on its domain, then

c) if $\xi(X_2)$ is convex, $F_{X_2}(x_2) = F_{X_1}(\xi(x_2))$ is a ρ-concave function if $F_{X_1}(x_1)$ is ρ-concave.

d) if $\xi(X_2)$ is concave, $F_{X_2}(x_2) = F_{X_1}(\xi(x_2))$ is a ρ-convex function if $F_{X_1}(x_1)$ is ρ-convex.

Proof. Lemma 2 shows for $\xi(.)$ increasing that $F_{X_2}(x_2) = F_{X_1}(\xi(x_2))$. Hence $F^{\rho}_{X_2}(x_2) = F^{\rho}_{X_1}(\xi(x_2))$. If $F_{X_1}(x_1)$ is ρ-concave, then $F^{\rho}_{X_1}(\xi(x_2))$ is concave, because it is a concave function of an increasing and concave function. Therefore $F^{\rho}_{X_2}(x_2)$ is concave, and so $F_{X_2}(x_2)$ is ρ-concave. The other relations are proved in a similar manner.[9] ■

The reason for dealing with decreasing F for the two last statements is that we can formulate properties for survivor functions.

The Lemma enables us to bound curvatures, as long as the inside argument obeys the requisite concavity/convexity. For example, take the relation $F_C(c) = F_P(p(c))$, and recall

[9] Succinctly, notice that $F''_{X_2}(x_2) = F''_{X_1}\xi'^2(x_2) + F'_{X_1}\xi''(x_2)$.

that $p(c)$ is increasing. If $p(c)$ is concave, then the Lemma (case (a)) tells us that if F_P is ρ-concave, then so is F_C (and conversely). Case (b) tells us the analogous property for $p(c)$ convex. So, say, if $p(c)$ is linear, as behooves linear demand, and F_P lies between the limits of ρ-concavity and ρ'-convexity (with $\rho \geq \rho'$) then F_C also must lie within the same concavity-convexity limits.[10] Or, indeed, if F_P were log-concave ($\rho = 0$), then so too would be F_C if price were concave in cost. If the relevant $\rho \geq 1$, then a decreasing density f_P implies a decreasing density f_C, and vice versa.

The more interesting relations concern the cases when the distributions are negatively related, e.g., $F_C(c) = 1 - F_\Pi(\pi(c))$. Then we can work with the survivor function $G_\Pi(\pi) = 1 - F_\Pi(\pi)$, which is a decreasing function so that cases (c) and (d) apply. Taking the cost-profit example, we recall $\pi(c)$ is convex, so that case (c) applies here. Then $F_C(c)$ is ρ-concave if $G_\Pi(\pi)$ is ρ-concave. A log-concave profit survivor function implies the cost distribution that would generate such a pattern must be log-concave.

Or take the output distribution and the price survivor function, $G_P(p) = 1 - F_P(p)$. Since $y = h(p)$, the shape of the output distribution is related to the price distribution (and its survivor function) via the concavity or convexity of demand. If demand is concave, we have case (d): the output distribution is ρ-convex if the price survivor function is ρ-convex. Uniform prices are associated to a convex output distribution, which means an increasing output density. This makes sense: half the prices exceed the average one, while the output at the average price is above half the average output for concave demand.

[10] See Anderson and Renault (2005) for more on bounds of ρ-concave and ρ'-convex functions.

3.2 Density and distribution elasticity relations

Economics has several key relations involving elasticities, notably the inverse-elasticity of demand relation with the Lerner index, and the Dorfman-Steiner relations for optimal advertising. The ones we provide below (in the Lemma) are straightforward to derive, but they are quite fundamental for the monopolistic competition setting.

Elasticities of distributions (or survivor functions) are readily calculated from (5) using the relations between the various variables. There are also clean and useful conditions that relate the elasticities of equilibrium densities. These simple formulae show which other elasticities connect the distributions (or survivor functions) and densities, and they are all different aspects of the demand side. For example, the profit density elasticity is related to the cost density elasticity via the elasticities of profit and (inverse) marginal revenue (with respect to unit cost, c), both of which are derived from the fundamental demand form. When the demand side delivers constants for the various elasticities, as it does for constant elasticity demand, then density elasticities are just affine functions of each other (this result is presented in Anderson and de Palma, 2018, for the CES demand model). Moreover, when in turn one of the density elasticities is constant, then they all are constant, as an implication of this property.

The fundamental relations are derived from the following elasticity lemma.

Lemma 4 *Consider two distributions $F_{X_1}(x_1)$ and $F_{X_2}(x_2)$, which are absolutely continuous and strictly increasing on their respective domains. Let X_1 and X_2 be related by a monotone function $X_1 = \xi(X_2)$. Then we have the elasticity relation between densities as $\eta_{f_{X_2}} = \eta_{f_{X_1}} \eta_\xi + \eta_{\xi'}$.*

Proof. For $\xi(.)$ increasing, from Lemma 2, we have $F_{X_2}(x_2) = F_{X_1}(\xi(x_2))$. Differentiation yields $f_{X_2}(x_2) = f_{X_1}(\xi(x_2))\,\xi'(x_2)$; differentiating again gives $f'_{X_2}(x_2) = f'_{X_1}(\xi(x_2))\,(\xi'(x_2))^2 + f_{X_1}(\xi(x_2))\,\xi''(x_2)$. Dividing through the second expression by the first and multiplying both sides by x_2 delivers

$$x_2 \frac{f'_{X_2}(x_2)}{f_{X_2}(x_2)} = \xi'(x_2)\,\frac{f'_{X_1}(\xi(x_2))}{f_{X_1}(\xi(x_2))}\frac{\xi'(x_2)}{\xi(x_2)}x_2 + \frac{\xi''(x_2)}{\xi'(x_2)}x_2;$$

which is the expression given in the Lemma.

For $\xi(.)$ decreasing, from Lemma 2, we have $F_{X_2}(x_2) = 1 - F_{X_1}(\xi(x_2))$. Differentiating, $f_{X_2}(x_2) = -f_{X_1}(\xi(x_2))\,\xi'(x_2)$; and the steps above then again imply the expression given. ■

Assumption 1 imposes several restrictions on the various demand-side elasticities that appear in the density elasticity relations below. In particular, $\eta_h < -1$ is the property that demand must be elastic at in a monopolistic competition setting, mirroring the standard monopoly property. Furthermore, $\eta_{h^*} < 0$ is the property that marginal revenue slopes down. The elasticity of the demand curve slope, $\eta_{h'} = \frac{h"p}{h'}$, has the sign of $-h"$ and so is positive for concave demand, and negative for convex demand. The elasticity of the inverse marginal revenue slope, $\eta_{h^{*\prime}}$, involves third derivatives of demand, though notable benchmarks are that it is zero for linear demand (because marginal revenue is linear) and the constant elasticity case discussed in the next paragraph. Finally, the elasticity of maximized profit (with respect to c), η_π, is particularly interesting. Write this as

$$\eta_\pi = \frac{\pi'(c)}{\pi(c)}c = -\frac{ch^*(c)}{\mu(c)\,h^*(c)} = -\frac{c}{\mu(c)} < 0.$$

The third expression is the ratio of total cost to total profit;[11] and the last one is a variant on the Lerner index (which is $\frac{p-c}{p}$) and here we have the statistic $\frac{p-c}{c}$, which is just the relative mark-up over cost.

In the analysis that follows, the reference point of the constant elasticity demand is useful. Write then $h(p) = p^k$ with $k < -1$[12] so $\eta_h = k < -1$ and $\eta_{h'} = k-1 < -2$. Furthermore, inverse marginal revenue has the same elasticity as demand,[13] so $\eta_{h^*} = k$; and the profit elasticity is $\eta_{\pi^*} = k + 1 < 0$, namely the demand elasticity plus the (unit) mark-up elasticity.[14]

We can track six pairs of relations, those between price, profit, output, and cost. The latter is the fundamental one, the others are the economic ones induced through the taste side encapsulated in the demand-side elasticities. However, any three of the pairs of relations tie down the rest. We describe the more interesting ones.

3.2.1 price and output

The defining relation between these two is $F_P(p) = 1 - F_Y(h(p))$, Lemma 4 tells us that

$$\eta_{f_P} = \eta_{f_Y}\eta_h + \eta_{h'}. \tag{9}$$

[11]Else $\eta_\pi = -\frac{TC}{TR-TC} = \frac{1}{\frac{TR}{TC}-1}$.

[12]Demand must be elastic or else Assumption 1 is violated: with inelastic demand a firm would produce infinitesimal output.

[13]Because the MR=MC condition yields $\left(\frac{1}{k}+1\right)y^{1/k} = c$.

[14]For the CES demand model, with demand parameter ρ, and equilibrium price $p = c/\rho$, then $k = \frac{1}{\rho-1}$ so $\eta_h = k = \frac{1}{\rho-1} < 0$; $\eta_{h'} = k - 1 = \frac{2-\rho}{\rho-1} < 0$; $\eta_{h^*} = \eta_h = \frac{1}{\rho-1} < 0$; and $\eta_\pi = \frac{\rho}{\rho-1} < 0$. The latter can also be expressed (as per the preceding text) as $-\frac{c}{p-c}$ and recall that $p = c/\rho$. These ρ values concur with those used in the analysis of Anderson and de Palma (2018). [check!]

Recall that the elasticity of the demand slope has shown up elsewhere in pricing formulae (e.g., in Helpman and Krugman, 1985). Note that the demand elasticity on the RHS is negative, so that, ceteris paribus, the effect of the first term on the RHS is to deliver a negative relation between price density elasticity and output density elasticity. Now consider the other term on the RHS. As noted above, $\eta_{h'} = \frac{h"p}{h'}$ has the sign of $-h"$ and so is positive for concave demand, and negative for convex demand. For linear demand, the term disappears. Then we have a benchmark that the price and output densities have *opposite* signs. But, this is apparently not always true otherwise. But we can say, for example, that concave demand implies that decreasing output density drives increasing price density. For convex demand, increasing price density drives decreasing output density. To interpret the negative relation in the benchmark, recall that the low price firms are the high output ones, so we are looking at opposite ends of the distributions/densities effectively. Think about an increasing price density. Then there are more firms with higher prices: translating to the output density, there are fewer firms with lower outputs.

For the constant elasticity demand, (9) reduces to

$$\eta_{f_P} = k\eta_{f_Y} + (k-1)$$

(where we recall that $k < -1$), which implies that an output density $\eta_{f_Y} > -1$ entails an decreasing price density. That is, the price density tends to be decreasing even when the output density is decreasing (though it should not be decreasing by too much). So CES suggests strongly tending for price density to decrease (you would need $\eta_{f_Y} < \rho - 2$ to overturn this).

The linear benchmark suggests the outcome of price and output densities sloping opposite ways, and this is tempered or reinforced by the demand curvature.

As regards the distribution elasticity relation, we have

$$\eta_{F_P} = \eta_{G_Y} \eta_h$$

where η_{G_Y} denotes the elasticity of the survivor function of output (and not the output distribution per se). The equilibrium relation shows how the price distribution is more elastic when equilibrium demand is more elastic, ceteris paribus.

3.2.2 costs and profit; profit and price

The defining distribution relation between profit and cost distributions is $F_\Pi(\pi(c)) = 1 - F_C(c)$. This delivers the density relation as

$$\begin{aligned} f_\Pi(\pi(c))\, \pi'(c) &= -f_C(c) \\ \text{or } f_\Pi(\pi(c))\, h^*(c) &= f_C(c) \end{aligned}$$

and then the elasticity relation immediately follows (or else use Lemma 4) as

$$\eta_{f_\Pi} \eta_\pi = \eta_{f_C} - \eta_{h^*},$$

or, equivalently,

$$\eta_{f_\Pi} = \frac{\mu(c)}{c} \left(-\eta_{f_C} + \eta_{h^*} \right),$$

which tells us that the profit density elasticity is proportional to the difference of the demand-side inverse MR and cost density elasticities. Recalling that η_{h^*} is negative, if $\eta_{f_C} > 0$, then necessarily $\eta_{f_\Pi} < 0$. The contra-positive is that $\eta_{f_\Pi} > 0$ implies $\eta_{f_C} < 0$. Then the profit density is falling quite naturally if the cost density is increasing (or not decreasing too much). It would need a strongly falling cost density (strongly increasing productivity) to overturn the effect.

Now consider the CES version, for which $\mu(c) = c\left(\frac{1}{\rho} - 1\right)$, so that $\eta_{f_\Pi} = -\frac{1-\rho}{\rho}\left(\eta_{f_C} + \eta_{h^*}\right) = -\frac{1-\rho}{\rho}\left(\eta_{f_C} - \frac{1}{1-\rho}\right)$,[15] and that therefore the cost tail parameter gets powered into the profit one by ρ.

In terms of distributions, the elasticity relation is

$$\eta_{F_C} = \eta_{G_\Pi}\eta_{\pi(c)},$$

which can be rewritten as

$$\frac{\mu(c)}{c} = -\frac{\eta_{G_\Pi}}{\eta_{F_C}}.$$

This indicates how the relative mark-up in equilibrium can be found from the distribution shapes.

For profit and price, the analogous density expression is $\eta_{f_\Pi}\eta_\pi = \eta_{f_P} - \eta_h$, or $\eta_{f_\Pi} = \frac{\mu(c)}{c}\left(-\eta_{f_P} + \eta_{h^*}\right)$. The distributional elasticity relation writes $\eta_{F_C} = \eta_{G_\Pi}\eta_{\pi(p)}$.

[15] See also Anderson and de Palma (2018).

3.2.3 output and cost

The defining relation for this pair is $F_Y(h^*(c)) = 1 - F_C(c)$. Then Lemma 4 tells us that[16]

$$\eta_{f_Y}\eta_{h^*} = \eta_{f_C} - \eta_{h^{*\prime}}$$

This is directly comparable to the price-output relation described above (namely $\eta_{f_Y}\eta_h = \eta_{f_P} - \eta_{h'}$). Drawing on that analysis, a linear marginal revenue is a useful benchmark,[17] for which output and cost densities necessarily go in opposite directions. The constant elasticity of demand case is just like the output-price case, given that the parameters are the same for both cases.

Another general link (which is also useful for the next case) is that $\eta_{h^*} = \eta_h \eta_{p(c)}$ where $\eta_{p(c)} = \frac{p'(c)c}{p(c)}$. When the elasticity of demand is constant, $\eta_{f_C} = \eta_{f_P}$ so that $\eta_{p(c)}$ has unit elasticity, concurring with the claim made earlier that then $\eta_{h^*} = \eta_h$.

3.2.4 price and cost; profit and cost; output and cost;

These are all similar relations. First, $F_P(p(c)) = F_C(c)$, so $f_P(p)\, p'(c) = f_C(c)$ and thence

$$\eta_{f_C} = \eta_{p(c)}\eta_{f_P} + \eta_{p'}.$$

The analogous expressions for the other pairs noted in the header are simply given by just replacing the p's by the other variables. For constant elasticity demand, we see the property

[16] Or, indeed, the density relation is $f_Y(h^*(c))\, h^{*\prime}(c) = -f_C(c)$; write this in log form and the elasticity relation follows directly.

[17] This comes from linear demand, but is not limited to that – we can add a rectangular hyperbola to demand and still get a linear marginal revenue.

noted above, $\eta_{f_C} = \eta_{f_P}$, because $\eta_{p(c)} = 1$ and $\eta_{p'} = 0$ (because $p'(c)$ is constant).

4 Equivalences

In the sequel in the following sections, we shall determine how to recover demand (and other distributions) from any pair of distributions, and what restrictions on distributions (if any) must be obeyed in order to satisfy Assumption 1 (that demand is twice continuously differentiable and strictly (-1)-concave). To do so, we shall make use of the results of this section, which form the converse properties to the results of Lemma 1. These are properties that form a stand-alone contribution to the theory of monopolistic competition, and of monopoly, so we collect them together here.

Specifically, Lemma 1 and Corollary 1 show that the demand Assumption 1 (statement (i)) implies the properties (ii) through (iv):

(i) demand is twice continuously differentiable and strictly (-1)-concave;

(ii) the equilibrium mark-up, $\mu(c) > 0$, is a continuously differentiable function with $\mu'(c) > -1$;

(iii) the equilibrium price, $p(c)$, is a continuously differentiable function with $p'(c) > 0$;

(iv) the equilibrium demand, $h^*(c)$, is a continuously differentiable function with $h^{*\prime}(c) < 0$;

(v) the equilibrium profit, $\pi^*(c)$, is strictly convex and twice continuously differentiable, with $\pi^{*\prime}(c) = -h^*(c) < 0$.

Here we show that these are all equivalent statements, so that any one implies the others. Indeed, Corollary 1 already indicates that (ii) and (iii) are equivalent. Likewise, (iv) and (v) are equivalent given the envelope theorem result (the monopoly analogue to Hotelling's Lemma),

$\pi^{*\prime}(c) = -h^*(c) < 0$, shown in Lemma 1. Therefore it remains to prove that (ii) implies (i), and (iv) implies (i). We treat these in turn (in reverse order).

4.1 strictly decreasing marginal revenue implies strictly (-1)-concave demand

First note that $h*(c)$ is strictly decreasing if and only if marginal revenue, $MR(y) > 0$, is strictly decreasing, with both continuously differentiable. This is because these are inverse functions. Next, integrating $MR(y)$ yields total revenue, $TR(y)$, which is therefore twice continuously differentiable (and it is strictly quasi-concave, and monotone increasing for $MR(y)$). Average revenue, $AR(y)$, is then $TR(y)/y$, and this twice continuously differentiable function is inverse demand, $p(y)$. Inverting it yields $h(p)$ as a twice continuously differentiable function. It remains to show that $h(p)$ is strictly (-1)-concave. The proof following the next result concludes the issue.

Lemma 5 *If inverse marginal revenue, $h^*(c)$, is strictly decreasing and continuously differentiable, then demand, $h(p)$, is strictly (-1)-concave and twice continuously differentiable.*

Proof. First note that $h(p)$ is strictly (-1)-concave if and only if $h''h - 2(h')^2 < 0$. Write the inverse demand as $p(y)$ so that $h'(p) = \frac{1}{p'(y)}$ and $h''(p) = -\frac{p''(y)}{(p'(y))^3}$. Then the strict (-1)-concavity condition we are to show becomes

$$p''y + 2p' < 0. \tag{10}$$

Now we want to find $p(y)$, using the steps explained before the Lemma. Let $MR(y)$ denote $h^{*-1}(c)$, i.e., marginal revenue. So then Total Revenue, $TR(y)$ is the integral of $MR(y)$ and equilibrium inverse demand, $p(y)$, is

$$p(y) = \frac{TR(y)}{y} = \frac{\int_0^y MR(u)\,du}{y},$$

and the inverse is $h(p)$. Hence $p'(y) = \frac{yc(y) - \int_0^y c(u)du}{y^2}$ and $p''(y) = \frac{c'(y) - 2\left(yc(y) - \int_0^y c(u)du\right)}{y^3}$. Using these expressions in (10) gives

$$MR'(y) < 0,$$

which follows because $h^{*\prime}(c)$ is continuously differentiable and negative. Q.E.D. ■

Notice that the demand $h(p)$ is only determined up to a constant (from the step where $MR(.)$ is integrated): intuitively, one can always add a rectangular hyperbola to any inverse demand (the rectangular hyperbola has a zero Marginal Revenue) and get the same Marginal Revenue function.

4.2 constructing demand from the mark-up function

Here we show how $\mu(c)$ (with $\mu'(c) > -1$) can be used to find the associated equilibrium demand and demand function, $h(p)$. Equivalently, we could start with a continuously differentiable and strictly increasing relation between equilibrium price and cost, $p(c)$. Our converse result to Lemma 1 indicates how the mark-up function $\mu(c)$ determines the form of inverse marginal revenue, $h^*(c)$, and hence determine the form of $h(p)$.

Lemma 6 *Consider any positive mark-up function $\mu(c)$ for $c \in [\underline{c}, \bar{c}]$ with $\mu'(c) > -1$. Then there exists an equilibrium demand function $h^*(c)$ with $h^{*\prime}(c) < 0$, defined on its support $[\underline{c}, \bar{c}]$ and given by (12), which is unique up to a positive multiplicative factor. The associated primitive* *demand function* $h(p)$*, given by (13), satisfies Assumption 1 on its support* $[\mu(\underline{c}) + \underline{c}, \mu(\bar{c}) + \bar{c}]$*.* $h(p)$ *is log-convex if* $\mu'(c) \geq 0$ *and log-concave if* $\mu'(c) \leq 0$*.*

Proof. First note from (2) and (4) that

$$\frac{dh^*(c)/dc}{h^*(c)} = \frac{(\mu'(c)+1)\,h'(\mu(c)+c)}{h(\mu(c)+c)} = -\frac{\mu'(c)+1}{\mu(c)} \equiv g(c) < 0, \tag{11}$$

because $\mu'(c) > -1$ by assumption. Thus $[\ln h^*(c)]' = g(c)$, and so $\ln\left(\frac{h^*(c)}{h^*(\underline{c})}\right) = \int_{\underline{c}}^{c} g(v)\,dv$, or

$$h^*(c) = h^*(\underline{c}) \exp\left(\int_{\underline{c}}^{c} g(v)\,dv\right), \quad c \geq \underline{c}, \tag{12}$$

which determines $h^*(c)$ up to the positive factor $h^*(\underline{c})$; it is strictly decreasing because $g(c) < 0$.

We can now use the inverse marginal revenue function, $h^*(c)$, to back out the demand function, $h(m+c)$, via the following steps. First, define $u \equiv \phi(c) = \mu(c) + c$, which is strictly increasing because $\mu'(c) + 1 > 0$, so the inverse function $\phi^{-1}(\cdot)$ is strictly increasing. Now, $h(u) = h^*\left(\phi^{-1}(u)\right)$ and thus the function $h(\cdot)$ is recovered on the support $u \in [\mu(\underline{c}) + \underline{c}, \mu(\bar{c}) + \bar{c}]$ (cf. Lemma 1). Using (12) with $h(u) = h^*\left(\phi^{-1}(u)\right)$,

$$h(u) = h^*(\underline{c}) \exp\left(\int_{\underline{c}}^{\phi^{-1}(u)} g(v)\,dv\right), \tag{13}$$

and so

$$\frac{h(u)}{h'(u)} = \frac{1}{g\left(\phi^{-1}(u)\right)\left[\phi^{-1}(u)\right]'} = \frac{\phi'(c)}{-\frac{\mu'(c)+1}{\mu(c)}} = -\mu(c),$$

where the middle step follows from (11) with $u = \phi(c)$ and the last step follows because $\phi'(c) = \mu'(c) + 1$. Thus

$$\left[\frac{h(u)}{h'(u)}\right]' = -\frac{\mu'(c)}{\mu'(c)+1} > -1,$$

and so $h(u)$ is strictly (-1)-concave (as shown in footnote 5). Note that $h(.)$ is twice differentiable because $\mu(.)$ was assumed differentiable. ■

Recalling that $\mu(c) = p(c) - c$ for $c \in [\underline{c}, \bar{c}]$, the restriction used in the Lemma ($\mu'(c) > -1$) is that $p'(c) > 0$ so that *any* arbitrary (differentiable) increasing price function of costs can be associated to a unique demand function that could generate it (up to the multiplicative factor).

The reason that demand is only determined up to a positive factor is simply that multiplying demand by a positive constant does not change the optimal mark-up (when marginal costs are constant, as here). The mark-up function can only determine the demand shape, but not its scale. In conjunction, Lemmas 1 and 6 indicate the property that $\mu'(c) > -1$ if and only if $h(u)$ is strictly (-1)-concave.

The steps in the proof are readily confirmed for the $\rho-$linear example given after Lemma 6. Along with Lemma 1, the results of this section indicate that knowing any of $\mu(c)$, $h^*(c)$, or $h(.)$ suffices to determine them all (up to constants in the first two cases). This constitutes a strong characterization result for monopoly pass-through (see Weyl and Fabinger, 2013, for the state of the art, which deeply engages ρ-concave functions).

Notice that the function $h(\cdot)$ is tied down only on the support corresponding to the domain on which we have information about the equilibrium mark-up value in the market. Outside that support, we know only that $h(\cdot)$ must be consistent with the maximizer $\mu(c)$, which restricts the shape of $h(\cdot)$ to be not "too" convex.

5 Rationalizability of distributions via demand

An old question in consumer theory is whether a demand system can be generated from a set of underlying preferences (see Antonelli, 1898, and the discussion in Mas-Collel, Whinston, and Green, 20xx, p.?). Here we look at whether any arbitrary pair of economic/primitive distributions could be consistent with the monopolistic competition model with demand satisfying A1. Surprisingly, for 4 of the possible pairs of distributions, the answer is affirmative, so that the model places no restrictions (above the twice continuously differentiable assumption we retain for simplicity.) For the other two, we derive the conditions the distributions must satisfy, and in all cases we recover the implied demand function. We start with the key result that enables us to recover demand from the mark-up function.

5.1 Deriving demand from price and profit distributions

We now use Lemma 6 to find a unique demand function satisfying A1 from these distributions. This is quite a surprising result. For example, there exists a demand function that squares Pareto distributions for both prices and profits. In the next sub-sections we do likewise for other distribution pairs.

Note that all other distributions are determined (along with demand) from the original pair.

Theorem 3 *Let the price and profit distributions, F_P and F_Π, be two arbitrary strictly increasing and twice continuously differentiable functions on their supports. Then there exists a strictly (-1)-concave demand function (unique up to a positive constant) that rationalizes these distributions in the monopolistic competition model.*

Proof. Applying the techniques above (see (5)), first write $1 - F_C(c) = 1 - F_P(p) = F_Y(y) = F_\Pi(\pi) = z$. Then we can write $\pi = F_\Pi^{-1}(1 - F_P(p)) = h(p)\mu(p) \equiv \tilde{\pi}(p)$, where $\tilde{\pi}(p)$ therefore denotes the relation between the maximized profit level observed and the value of the corresponding maximizing price (i.e., $\tilde{\pi}(p)$ is the optimizing price delivering the profit level). Recall from the optimal choice of mark-up that $h(p)$ and $\mu(p)$ are related by $\mu(p) = -h(p)/h'(p)$ (see (2)), and so $\tilde{\pi}(p) = -h^2(p)/h'(p)$. Integrating,

$$h(p) = \frac{1}{\int_{\underline{p}}^{p} \frac{dr}{F_\Pi^{-1}(1-F_P(r))} + k}. \tag{14}$$

This determines the demand form up to the positive constant $k = 1/h(\underline{p})$ (in the position in the above formula). Finally, (14) is decreasing in p, and is twice continuously differentiable. Furthermore,

$$\left(\frac{1}{h(p)}\right)' = \frac{1}{F_\Pi^{-1}(1 - F_P(p))}$$

which is strictly increasing because both distributions are strictly increasing. That is $1/h(p)$ is convex and so, equivalently, $h(p)$ is (-1)-concave.

By Theorem 2 all the other distributions are determined. ■

Therefore, after using the price and profit distributions we can define the function (14) and the proof shows that the resulting demand function satisfies Assumption 1 *without any further restrictions*. This means, for example, that a decreasing price density is consistent with an increasing profit density (very many high profit firms and yet very few high price ones). The underlying cost distribution along with demand is what renders these features compatible. As regards the constant k, knowing the demand level at any one point ties down the whole demand function.

We have just shown that there are no restrictions on price and profit distribution shapes, though below we have restrictions on some other pairs of distribution functions that can be combined and be consistent with the monopolistic competition model.

5.2 price and cost

We now determine distributions from each other when there are monotone relations between two variables. Suppose first that price and cost distributions, F_P and F_C, are known. Because mark-ups are necessarily positive, it must be that the price distribution first-order stochastically dominates the cost one. However, we will show that this is the *only* restriction on the distributions. The demand function will ensure that they are compatible, even despite it being (-1)-concave.

We show how to find the implied other economic distributions as well as the demand form and mark-up function: we can find all other elements in the market from just the two distributions. This strong result relies on the monotonic relations between all pairs of variables from Corollary 1. We now show how this works. Because price strictly increases with cost, the price and cost distributions are matched: the fraction of firms with costs below some level c equals the fraction of firms with prices below the price charged by a firm with cost c. This enables us to back out the corresponding mark-up function $\mu(c)$ and then access Lemma 6.

Theorem 4 *Let the cost and price distributions, F_C and F_P be two arbitrary strictly increasing and twice continuously differentiable functions on their supports with $F_C(c) > F_P(c)$. Then there exists a strictly (-1)-concave demand function (unique up to a positive factor) that ratio-*

nalizes these distributions in the monopolistic competition model. Then the mark-up function $\mu(c)$ (with $\mu'(c) > 0$) is found from (15); inverse marginal revenue is found from (12) and the demand function is given from (13), up to a positive multiplicative factor, $h^*(\underline{c})$. The output and profit distributions are determined, up to $h^*(\underline{c})$, by (7) and (8).

Proof. Consider a distribution of costs, F_C and a distribution of prices, F_P satisfying $F_C(c) > F_P(c)$ (so that the price distribution is right of the cost one: note that $F_C(c) > F_P(c) = 0$ for c below the lower bound of the support of the price distribution). We wish to find a demand function satisfying A1. Define $p(c) = F_P^{-1}(F_C(c))$ which is an increasing function. Then Lemma 6 implies that there exists an $h(.)$ satisfying A1 (consistent with Corollary 1 that the price charged by a firm with cost c is a strictly increasing continuously differentiable function $p(c)$).

Hence we can write the price-cost margin, as a function of c, as

$$\mu(c) = F_P^{-1}(F_C(c)) - c, \tag{15}$$

with $\mu(c) > 0$ because $F_C(c) > F_P(c)$ and $\mu'(c) > -1$. Hence a unique such mark-up function $\mu(c)$ exists given the cost and price distributions. With the function $\mu(c)$ thus determined, we can invoke Lemma 6 to uncover the equilibrium demand function $h^*(\cdot)$ (unique up to a positive multiplicative factor) as given by (11) and (12), and the demand function is given from (13). By Lemma 6, this demand function satisfies A1, as postulated. ■

The idea behind the result is as follows. Given the first key property that prices rise with costs, we know that the $z\%$ of firms with cost below c are the $z\%$ of firms with an equilibrium price below p. This links the mark-up and the cost level, so we can use Lemma 6 to uncover the demand form and equilibrium output of the zth percentile firm, due to the second key property that equilibrium output is a decreasing function of cost. We hence uncover the output distribution. The profit distribution then follows immediately from knowing the output and mark-up distributions. The latter two distributions are only determined up to a positive factor because the mark-up function is consistent with any multiple of the demand (under the maintained hypothesis of constant returns to scale).

The construction of the demand function is illustrated in Figure 1 above. The only restriction we use here is that the cost distribution first-order stochastically dominates the price one. Given this property, any pair of (twice continuously differentiable) price and cost functions is consistent with the monopolistic competition model. We next show that the price and output distributions are restricted if they are to be consistent.

5.3 price and output

We now suppose that price and output distributions, F_P and F_Y, are known.

Theorem 5 *Let the price and output distributions, F_P and F_Y be two arbitrary strictly increasing and twice continuously differentiable functions on their supports. Then there exists a unique strictly (-1)-concave demand function $h(p) = F_Y^{-1}(1 - F_P(p))$ that rationalizes these distributions in the monopolistic competition model if and only if $F_Y^{-1}(1 - F_P(p))$ is (-1)-concave.*

Proof. Because from the two distributions $y = F_Y^{-1}(1 - F_P(p)) = h(p)$ this is the unique candidate demand function. While this is decreasing in p, as desired, we also require that the function $F_Y^{-1}(1 - F_P(p))$ is (-1)-concave to be consistent with the monopolistic competition model. ■

Therefore if the implied demand shape does not satisfy the given condition, the purported demand relation would not have a downward-sloping marginal revenue curve everywhere, and any price-output pair with an upward sloping marginal revenue could not be consistent with profit maximization by a firm.

5.4 cost and output

Although price and output distributions are jointly restricted, surprisingly, cost and output distributions are not. Suppose that F_C and F_Y are known.

Theorem 6 *Let the cost and output distributions, F_C and F_Y be two arbitrary strictly increasing and twice continuously differentiable functions on their supports. Then there exists a strictly (-1)-concave demand function (unique up to a positive constant) that rationalizes these distributions in the monopolistic competition model.*

Proof. From the two distributions $y = F_Y^{-1}(1 - F_C(c)) = h^*(c)$ is the candidate function for optimized demand. The only restriction is that it slope down, which is satisfied, and that it be continuous, which is also immediately satisfied. Hence it is rationalizable, and we can use earlier results (Lemma 5) to back up to the implied demand function, $h(p)$, which is therefore determined up to a positive constant. ■

5.5 cost and profit

This is another case where monopolistic competition restricts the distributions. Suppose that F_C and F_Π are known.

Theorem 7 *Let the cost and profit distributions, F_C and F_Π be two arbitrary strictly increasing and twice continuously differentiable functions on their supports. Then there exists a strictly (-1)-concave demand function (unique up to a positive constant) that rationalizes these distributions in the monopolistic competition model if and only if f_C/f_Π is increasing.*

Proof. From the two distributions, $\pi(c) = F_\Pi^{-1}(1 - F_C(c))$ is the candidate profit function. This is decreasing in c, as desired, but it also needs to be convex. The convexity condition is that f_C/f_Π is increasing in c in order to be consistent with the monopolistic competition model.

■

5.6 output and profit

The final case returns to no restrictions. Consider a distribution of output, F_Y and a distribution of profit, F_Π.

Theorem 8 *Let the output and profit distributions, F_Y and F_Π be two arbitrary strictly increasing and twice continuously differentiable functions on their supports. Then there exists a strictly (-1)-concave demand function (unique up to a positive constant) that rationalizes these distributions in the monopolistic competition model. This unique net demand function, and the other distributions, are determined explicitly in the proof.*

Proof. The assumption that F_Y and F_Π are continuously differentiable distributions means that we can invert them and write each of them as a function of the counter z. Both output and profit are increasing functions of cost, c. Therefore we can match the distributions: the firms with the highest $z\%$ of the costs are those with the lowest $z\%$ of the outputs and profits. Furthermore, because the distribution functions are differentiable, then we should have that z is a differentiable function of the underlying cost, and invert it. Call this inverted relation $c(z)$, with $c'(z) < 0$. Then we can write

$$\pi(c(z)) - \mu(c(z))\,h^*(c(z)) = 0.$$

Differentiating this identity,

$$\pi'(c) - \mu'(c)\,h^*(c) - \mu(c)\,h^{*'}(c) = 0.$$

But, by Lemma 1 $\pi'(c) = -h^*(c)$ so that

$$\mu'(c) = -\mu(c)\frac{h^{*'}(c)}{h^*(c)} - 1 > -1$$

which means that A1 holds.

We now determine the net inverse demand, the mark-up function, and the other distributions. We know that $h^*(c)$ is strictly decreasing in c, and so too is $\pi^*(c) = \mu(c)\,h^*(c)$ (by Lemma 1). We hence choose some arbitrary level $z \in (0,1)$ such that $1 - F_C(c) = F_Y(y) = F_\Pi(\pi) = z$. This means that all firm types with cost levels above $c(z) = F_C^{-1}(1-z)$ are the firms with outputs and profits below y and π. For this proof, we introduce z as an argument into the various outcome variables to track the dependence of the variables on the level of $z(c) = 1 - F_C(c)$. From (**??**) we can write $y(z) = F_Y^{-1}(z)$ and demand is

$$h^*(c) = y(z(c)) = F_Y^{-1}(1 - F_C(c)). \tag{16}$$

Because $\pi^*(z) = m(z)\, y(z) = F_\Pi^{-1}(z)$ then

$$m(z(c)) = \frac{F_\Pi^{-1}(z(c))}{F_Y^{-1}(z(c))} = \mu(c), \tag{17}$$

and equilibrium profit is $\pi^*(c) = \mu(c)\, h^*(c) = F_\Pi^{-1}(z(c))$.

It remains to find the relation $z(c)$. From (17), $\mu'(c) = m'(z(c))\, z'(c)$ and similarly $h^{*\prime}(c) = y'(z(c))\, z'(c)$ (where $m'(z(c)) = \frac{dm(z(c))}{dz(c)}$, etc.). The two functions $\mu(c)$ and $h^*(c)$, which are to be determined, satisfy condition (11), which implies

$$\frac{h^{*\prime}(c)}{h^*(c)} = \frac{y'(z(c))\, z'(c)}{y(z(c))} = -\frac{m'(z(c))\, z'(c) + 1}{m(z(c))}.$$

Rearranging the last equality to solve out for $z'(c)$ gives[18]

$$z'(c) = -\frac{y(z(c))}{\left[m(z(c))\, y(z(c))\right]'} = -\frac{F_Y^{-1}(z(c))}{\left[F_\Pi^{-1}(z(c))\right]'}. \tag{18}$$

Thus:

$$\int_0^z \frac{\left[F_\Pi^{-1}(r)\right]'}{F_Y^{-1}(r)} dr = -\int_{\bar{c}}^c dv = \bar{c} - c,$$

[18] An alternative derivation is to use Theorem 1 to write $\pi^{*\prime}(c) = -h^*(c)$, so the relation between the counter z and the cost level c is $dz/dc = -h^*(c) / [\pi^*(z(c))]'$, which is (18).

or $c(z) = \bar{c} - \Psi(z)$, where $\Psi(z)$ is the key transformation between z and c given by

$$\Psi(z) = \int_0^z \frac{\left[F_\Pi^{-1}(r)\right]'}{F_Y^{-1}(r)} dr. \tag{19}$$

Because $\Psi'(z) = \frac{\left[F_\Pi^{-1}(z)\right]'}{F_Y^{-1}(z)} > 0$, the required relation between z and c is $z(c) = \Psi^{-1}(\bar{c} - c)$.

Observe that $h(p) = h(\mu(c) + c)$ so that inverse demand is $p = \frac{F_\Pi^{-1}(z(c))}{F_Y^{-1}(z(c))} + c = \frac{F_\Pi^{-1}\left(\Psi^{-1}(\bar{c}-c)\right)}{F_Y^{-1}(\Psi^{-1}(\bar{c}-c))} + c$. This makes clear that a shift up in all costs by Δ and a corresponding shift up in the inverse demand by Δ (so the support of the cost distribution shifts up by Δ, i.e., $\bar{c}$ becomes $\bar{c} + \Delta$) keeps both the firm's output choice and mark-up constant so output and profit are not changed. This means that these two distributions can only pin down *net* (inverse) demand.

The distribution of cost is thus given by

$$F_C(c) = 1 - z(c) = 1 - \Psi^{-1}(\bar{c} - c). \tag{20}$$

The remaining unknowns can be backed out now knowing $z(c)$: equilibrium demand is $h^*(c) = F_Y^{-1}(\Psi^{-1}(\bar{c} - c))$ from (16), and the mark-up function is $\mu(c) = \frac{F_\Pi^{-1}\left(\Psi^{-1}(\bar{c}-c)\right)}{F_Y^{-1}(\Psi^{-1}(\bar{c}-c))}$ from (17). ■

What the Theorem ties down is net demand (inverse demand minus cost): if both demand price and cost shift by the same amount then equilibrium quantity (output) and mark-up are unaffected, so profit is unchanged too. Thus output and profit distributions tie down the shape of the inverse demand and the shape of the other distributions, but not the inverse demand curve height. As we saw above, price and cost distributions alone do not tie down the demand scale, and nor do price and profit distributions. But the other pairs of distribution combinations fully determine the demand function and all distributions.

5.7 examples

We illustrate the theorem above with distributions that generate ρ-linear demand.

Example 1: *ρ-linear demands and uniform cost distribution.*

Suppose that $F_Y(y) = \frac{(1+\rho)y^{\rho}-1}{\rho}$, $y \in \left[\frac{1}{(1+\rho)^{1/\rho}}, 1\right]$, *and* $F_\Pi(\pi) = \frac{(1+\rho)\pi^{\rho/(1+\rho)}-1}{\rho}$, $\pi \in \left[\frac{1}{(1+\rho)^{(1+\rho)/\rho}}, 1\right]$, *with* $\rho > -1$.

Hence $F_Y^{-1}(z) = \left(\frac{\rho z+1}{1+\rho}\right)^{1/\rho}$ *and* $F_\Pi^{-1}(z) = \left(\frac{\rho z+1}{1+\rho}\right)^{(1+\rho)/\rho}$. *By (17), the ratio of these two yields the mark-up,* $m(z) = \frac{\rho z+1}{1+\rho} > 0$. *Because* $\left[F_\Pi^{-1}(z)\right]' = \left(\frac{\rho z+1}{1+\rho}\right)^{1/\rho}$, *we can write* $\Psi(z) = \int_0^z \frac{\left[F_\Pi^{-1}(r)\right]'}{F_Y^{-1}(r)}dr = z$, *and, because* $c(z) = F_C^{-1}(1-z)$, *then* $c(z) = \bar{c} - z$ *(with* $c'(z) = -1$*), so that* $\underline{c} = \bar{c} - 1$. *Now,* $F_C(c) = 1 - \Psi^{-1}(\bar{c} - c) = c - \underline{c}$. *Hence* $\mu(c) = \frac{\rho(\bar{c}-c)+1}{1+\rho}$. *Then* $y(c) = F_Y^{-1}(z(c)) = \left(\frac{\rho(\bar{c}-c)+1}{1+\rho}\right)^{1/\rho}$, *and* $h^*(c) = y(c)$. *We now want to find the associated demand,* $h(p)$. *We use the fact that* $p = \mu(c) + c = \frac{1+c+\rho\bar{c}}{1+\rho}$ *to write* $h(p) = (1+\rho(\bar{c}-p))^{1/\rho}$, *which is therefore a* $\rho-$*linear demand function (see (6)) with the parameter* k *set at* $k = \bar{c}$, *and* $\rho > -1$ *implies* $h(.)$ *is* (-1)*-concave.*

Note that $y(\bar{c}) = \left(\frac{1}{1+\rho}\right)^{1/\rho}$, *as verified by the upper bound,* $\bar{c}$, *while the lower bound condition* $\underline{c} = \bar{c} - 1$ *implies that* $y(\underline{c}) = 1$, *so costs are uniformly distributed on* $[\underline{c}, \bar{c}]$. *Lastly,* $\lim_{\rho\to 0} y(c) = \exp(\bar{c} - c)$ *gives the logit equilibrium demand (see Section 6).*

The uniform cost example gives a useful benchmark for some important properties relating cost distribution to profit distribution. For the example above, we have $f_\Pi(\pi) = \pi^{-1/(1+\rho)}$, so that the density of the profit distribution is decreasing, despite the underlying cost distribution that generates it being flat. This property indicates how profit density "piles up" at the low end. The output density shape is also interesting. For linear demand ($\rho = 1$), it is clearly flat

– equilibrium quantity is a linear function of cost. For convex demand ($\rho < 1$), it is decreasing, but for concave demand it is *increasing*, despite the property just noted that the profit density is decreasing. This suggests that (for concave demand), a decreasing output density requires an increasing cost density, which *a fortiori* entails a decreasing profit density.

As per Theorem 8, the (output, profit) distribution pair does not tie down the value of the constant. Theorems 5, 6, and 7 show which distribution pairs do tie down the full model, and knowing the demand form plus any distribution ties down everything (Theorem 1.) We return to the above example to illustrate, with the same parameters as Example 1, that knowing the profit and cost distributions ties down the full model.

Example 2: *ρ-linear demands from uniform cost distribution.*

Suppose that it is known that $F_C(c) = c$ for $c \in [0,1]$ and (as above) $F_\Pi(\pi) = \frac{(1+\rho)\pi^{\rho/(1+\rho)}-1}{\rho}$, $\pi \in \left[\frac{1}{(1+\rho)^{(1+\rho)/\rho}}, 1\right]$. We first write $\pi^*(c)$ to find $h^*(c) = -\pi^{*\prime}(c)$. Matching the distribution levels, $1 - c = \frac{(1+\rho)\pi^{\rho/(1+\rho)}-1}{\rho}$, or $\pi^*(c) = \left(\frac{\rho(1-c)+1}{1+\rho}\right)^{(1+\rho)/\rho}$ and hence $y(c) = h^*(c) = \left(\frac{\rho(1-c)+1}{1+\rho}\right)^{1/\rho}$, so both output and profit are power functions. Then we use $c = 1 - F_Y(y)$ with $F_Y(y) = \frac{(1+\rho)y^\rho - 1}{\rho}$ to get $\mu(c) = \frac{\pi^*(c)}{h^*(c)} = \left(\frac{\rho(1-c)+1}{1+\rho}\right) = [h^*(c)]^\rho$. Now use $p = \mu(c) + c$ to find $h(p) = (1 + \rho(1-p))^{1/\rho}$ and hence the (ρ-linear) demand form is tied down, including the value of the constant ($k = 1$: see (6), and consistent with the specification $\overline{c} = 1$).

6 The Logit model of monopolistic competition

We here derive specific results for the Logit model with quality-cost heterogeneity and a continuum of active firms, using a log-linear demand form. Total demand is normalized to 1, so output for Firm i is a Logit function of active firms' qualities and prices:

$$y_i = \hat{h}\left(v_i - p_i\right) = \frac{\exp\left(\frac{v_i - p_i}{\mu}\right)}{\int_{\omega\in\Omega} \exp\left(\frac{v(\omega)-p(\omega)}{\mu}\right) d\omega + \exp\left(\frac{v_0}{\mu}\right)}, \quad i \in \Omega, \tag{21}$$

where $\mu > 0$ measures the degree of product heterogeneity and $v_0 \in (-\infty, \infty)$ measures the attractivity of the outside option (which could also be a competitive sector). We thus adapt the continuous Logit model (see Ben-Akiva and Watanada, 1981) to monopolistic competition.[19]

As before, the (gross) profit for Firm i is $\pi_i = (p_i - c_i)\, y_i$, $i \in \Omega$. Since there is a continuum of firms, the own-demand derivative is $\frac{dy_i}{dp_i} = \frac{-y_i}{\mu}$, $i \in \Omega$, so that $\frac{d\pi_i}{dp_i} = y_i \left[1 - \frac{(p_i - c_i)}{\mu}\right]$, $i \in \Omega$: the term inside the square brackets is strictly decreasing in p_i, so the profit function is strictly quasi-concave and the profit-maximizing price of Firm i is[20]

$$p_i = c_i + \mu, \quad i \in \Omega. \tag{22}$$

As we showed earlier for log-linear demand, the absolute mark-up is the same for all firms.[21] The corresponding equilibrium outputs are

$$y_i = \frac{\exp\left(\frac{x_i}{\mu}\right)}{\int_{\omega\in\Omega} \exp\left(\frac{x(\omega)}{\mu}\right) d\omega + \mathcal{V}_0}, \quad i \in \Omega, \tag{23}$$

where $\mathcal{V}_0 \equiv \exp\left(\frac{v_0}{\mu} + 1\right) \geq 0$, and recall that $x(\omega) = v(\omega) - c(\omega)$ is a one-dimensional parameterization of quality-cost. (23) verifies the output ranking over firms seen before: $y_i > y_j$

[19] Anderson et al. (1992) show that logit demands can be generated from an entropic representative consumer utility function as well as the traditional discrete choice theoretic root (see McFadden, 1978).

[20] For oligopoly with n firms, the equilibrium prices are (implicit) solutions to $p_i = c_i + \frac{\mu}{1-y_i}$, $i = 1...n$. Under symmetry, $p = c + \frac{\mu n}{n-1}$, which converges to $c + \mu$ as $n \to \infty$ (Anderson, de Palma, and Thisse, 1992, Ch.7).

[21] The CES model gives a constant *relative* mark-up property, $p_i^* = c_i(1+\mu)$, regardless of quality (see Section **??**). The similarity between the Logit and CES is not fortuitous: μ is related to ρ in CES models by $\mu = \frac{1-\rho}{\rho}$. Both models can be construed as sharing their individual discrete choice roots (Anderson et al. 1992).

if and only if $x_i > x_j$, for $i, j \in \Omega$. Equilibrium (gross) profit is μy_i, $i \in \Omega$, so outputs and profits are fully characterized by quality-cost levels, yielding the following special case of Theorem **??**:

Proposition 1 *In the Logit Monopolistic Competition model, all firms set the same absolute mark-up, μ. Higher quality-cost entails higher equilibrium output and profit.*

As seen in Section **??**, insofar as higher qualities also bear higher costs then they are also higher priced, but output and profit may well be highest for medium-quality products.

6.1 Quality-cost, output, and profit distributions

Recall that the distribution of quality-cost is $F_X(x) = \Pr(X < x)$, with density $f_X(\cdot)$ and support $[\underline{x}, \infty)$. The corresponding distribution of equilibrium output, $F_Y(y)$, and the relation between x and y is[22]

$$y = \frac{1}{D} \exp\left(\frac{x}{\mu}\right), \quad y \geq \underline{y} = \frac{1}{D} \exp\left(\frac{\underline{x}}{\mu}\right), \tag{24}$$

where we assume henceforth that $f_X(\cdot)$ ensures the output denominator D (which is the aggregate variable) is finite, as is true for any finite support and the examples below:

$$D = M \int_{u \geq \underline{x}} \exp(u/\mu) f_X(u)\, du + \mathcal{V}_0, \tag{25}$$

and $M = \|\Omega\|$ is the total mass of firms. Equilibrium (gross) profit is $\pi = \mu y \geq \underline{\pi} = \mu \underline{y}$. The following uses results from Section **??**.

Proposition 2 *For the Logit Monopolistic Competition model, the quality-cost distribution, $F_X(x)$, generates the equilibrium output distribution $F_Y(y) = F_X(\mu \ln(yD))$ and the equilib-*

[22] Here all firms are active. Section **??** introduces fixed costs to render endogenous the set of active firms.

rium profit distribution $F_\Pi(\pi) = F_X(\mu \ln(\pi D/\mu))$*, where* D *is given by (25). Conversely,* $F_X(x)$ *can be derived from the equilibrium output distribution as* $F_Y\left(\frac{1}{D_Y}\exp\left(\frac{x}{\mu}\right)\right)$ *with* $D_Y = \frac{\mathcal{V}_0}{(1-My_{av})}$*, or from the equilibrium profit distribution as* $F_\Pi\left(\frac{1}{D_\Pi}\exp\left(\frac{x}{\mu}\right)\right)$ *with* $D_\Pi = \frac{\mathcal{V}_0}{(\mu-M\pi_{av})}$*, where* y_{av} *and* π_{av} *denote average output and profit, respectively.*

(Proof in online Appendix 2).

Finally, the distribution of costs $F_C(c)$ and the distribution of prices $F_P(p)$ are related by the mark-up shift, so $F_C(c) = F_P(c+\mu)$, with $\underline{p} = \underline{c} + \mu$. Conversely, knowing the price distribution ties down the cost distribution when the mark-up level, μ, is known. Note some special case results. First, there is no price dispersion if and only if there is no cost dispersion. Second, there is no profit dispersion if and only if there is no quality-cost dispersion: then there is only cost dispersion, which price dispersion mirrors. The "classic" symmetry assumption often analyzed in the literature since Chamberlin (1933) has neither cost nor profit dispersion.

6.2 Specific distributions

We derive the equilibrium profit distributions for the Logit model.[23] Proofs are in Appendix 2.

The *normal* distribution is perhaps the most natural primitive assumption to take for quality-costs. Then profit $\Pi \in (0,\infty)$ is log-normally distributed. The log-normal has sometimes been fitted to firm size distribution (see Cabral and Mata, 2003, for a well-cited study of Portuguese firms). Note that a truncated normal begets a truncated log-normal (which is therefore important once we consider free entry equilibria below).

The simplest text-book case is the *uniform* distribution. Then the equilibrium profit Π has

[23] The reverse relations hold by Proposition 2. Equilibrium output distributions are analogous (as $\pi = \mu y$).

distribution $F_\Pi(\pi) = \mu \ln\left(\frac{\pi D}{\mu}\right)$ and its density is unit elastic. A *truncated Pareto* distribution leads to a truncated Log-Pareto for profit (or output).

At a simplistic level, Proposition 2 indicates that we just need to find the log-distribution of the seed distribution. However, we still need to match parameters, as done in Appendix 2 for the examples, and we also need to find the corresponding expression for D and ensure it is defined. Notice too that the methods described above work for more general demands under monopolistic competition (see Section 2).

The most successful function to fit the distribution of firm size has been the Pareto. We reverse-engineer using Proposition 2 to find the distribution of quality-cost. This gives:

Proposition 3 *Let quality-cost be exponentially distributed:* $F_X(x) = 1 - \exp(-\lambda(x - \underline{x}))$, $\lambda > 0$, $\underline{x} > 0$, $x \in [\underline{x}, \infty)$, *with* $\lambda\mu > 1$. *Then equilibrium output and profit are Pareto distributed:* $F_Y(y) = 1 - \left(\frac{\underline{y}}{y}\right)^{\alpha_y}$ *and* $F_\Pi(\pi) = 1 - \left(\frac{\underline{\pi}}{\pi}\right)^{\alpha_\pi}$, *where* $\alpha_y = \alpha_\pi = \lambda\mu > 1$. *Conversely, a Pareto distribution for equilibrium output or profit can only be generated by an exponential distribution of quality-costs.*

Thus the shape parameter, $\alpha_y = \alpha_\pi$, for the endogenous economic distributions depends just on the product of the taste heterogeneity and the technology shape parameter.[24]

If the price distribution follows the Pareto form (suggested as empirically viable) $F_P(p) = 1 - \left(\frac{\underline{p}}{p}\right)^{\alpha_p}$, with $p \in [\underline{p}, \infty)$ and $\alpha_p > 1$, the corresponding cost distribution is also Pareto:

$$F_C(c) = 1 - \left(\frac{\underline{c} + \mu}{c + \mu}\right)^{\alpha_p}, \quad c \in [\underline{c}, \infty), \ \alpha_p > 1. \tag{26}$$

[24] D is bounded if $\mu > 1/\lambda$: which requires that taste heterogeneity exceeds average quality-cost.

Suppose for illustration that prices are Pareto distributed so that $F_C(c)$ is given by (26). If there were no quality heterogeneity, then we would find a power distribution for quality-cost,[25] which would therefore be inconsistent with the required exponential function predicated in Proposition 3. It is the extra relation that we are afforded via $\beta(c)$ that decouples the allowable distributions, and therefore can enable us to fit (for example) both Pareto prices and profits. The next Proposition gives the underlying $\beta(c)$ function.

Proposition 4 *Let $F_P(p)$ be Pareto distributed with shape parameter α_p and let $F_\Pi(\pi)$ be Pareto distributed with shape parameter $\alpha_\pi > 1$ and $\underline{\pi} < \frac{\mu}{M}\frac{\alpha_\pi - 1}{\alpha_\pi}$, and suppose that $x = \beta(c)$ is an increasing function. Then $\beta(c) = \beta(\underline{c}) + \mu\frac{\alpha_p}{\alpha_\pi}\ln\left(\frac{c+\mu}{\underline{c}+\mu}\right)$, where $\beta(\underline{c}) = -\mu\ln\left[\frac{1}{\mathcal{V}_0}\left(\frac{\mu}{\underline{\pi}} - \frac{\alpha_\pi M}{\alpha_\pi - 1}\right)\right]$.*

Proof. Let $F_X(x)$ be the exponential function given in Proposition 3. Because $\beta(\cdot)$ is increasing, $F_C(c) = F_X(\beta(c)) = 1 - \exp[-\lambda(\beta(c) - \beta(\underline{c}))]$, $\lambda = \frac{\alpha_\pi}{\mu} > 0$, $\underline{c} > 0$, $c \in [\underline{c}, \infty)$. A Pareto price distribution with shape parameter α_p delivers the cost distribution (26). Equating these two expressions gives

$$\beta(c) = \beta(\underline{c}) + \mu\frac{\alpha_p}{\alpha_\pi}\ln\left(\frac{c+\mu}{\underline{c}+\mu}\right), \quad \beta(c) \in [\beta(\underline{c}), \infty].$$

Thus $\beta'(c) > 0$ so that valuations rise faster than costs. The lower bound of the distribution $\beta(\underline{c}) = \underline{x}$ is given from Proposition 3 and is given in the Proposition statement. ■

Thus we can close the loop and deliver a model consistent with Pareto distributions (for example) for both profit and price, once we allow for a quality-cost relation. Figure 1 above is based on just such a logit example (with parameters $\alpha_p = \alpha_\pi = \mu = 2$, $\beta(\underline{c}) = 0$, $\underline{c} = 0$, and

[25] I.e., calling the common quality level $\bar{v}$, $F_X(x) = \left(\frac{\bar{v}-\underline{x}+\mu}{\bar{v}-x+\mu}\right)^{\alpha_p}$, $x \in (-\infty, \underline{x}]$.

$\underline{x} = 1$). So far, we have assumed the population of firm types is fixed: we next use the logit as a vehicle to endogenize the surviving form types when there are entry costs.

6.3 Comparative statics of distributions

So far in the paper we have considered cross-section properties of distributions and how they can be uncovered. Although it is not our main theme, we can also give some indication of comparative static properties across equilibria. How do distributions change with underlying preference and technology changes? This is a more tricky question because the values of the aggregate variables in demand, which we were able to suppress in the cross section analysis because they do not change, are now endogenously determined. While a full treatment is beyond our current scope, we can give some pointers (and deliver results) for the current Logit model, and determine how economic distributions change with fundamentals.

We here therefore briefly consider the comparative static properties for the Logit model. Because we are dealing with distributions, the natural way of doing so is to engage first order stochastic dominance (fosd). Proofs for this sub-section are in the CEPR Discussion Paper.

Proposition 5 *A fosd increase or a mean-preserving spread in quality-cost raises mean output and mean profit, and strictly so if the market is not fully covered (i.e., if $\mathcal{V}_0 > 0$).*

While moving up quality-cost mass will move up output mass ceteris paribus, it also increases competition for all the other firms (a D effect), which ceteris paribus reduces their output. Mean output does not necessarily rise if mean quality-cost rises.[26]

[26] Mean output rises with a mean-preserving spread but then if the mean of quality-cost is reduced slightly, mean output can still go up overall, so the two means can move in opposite directions.

Because the relation between output and profit distributions does not involve D, a fosd increase in output implies an increase in profit, and vice versa. However, a fosd increase in quality-cost does not necessarily lead to a fosd increase in output. Suppose for example that the increase in quality-cost is small for low quality-costs, but large for high ones. Then competition is intensified (an increase in D), and output at the bottom end goes down, while rising at the top end. So then there can be a rotation of $F_Y(.)$ (in the sense of Johnson and Myatt, 2006) without fosd (a similar rotation is delivered in Proposition 6 below). Nevertheless, specific examples do deliver stronger relations, as we show in Section 6.2.

We next determine how taste parameters feed through into the endogenous economic distributions.

Proposition 6 *A more attractive outside option ($\mathcal{V}_0$) fosd decreases outputs and profits. More product differentiation (μ) fosd increases outputs and profits for low quality-cost and fosd decreases them for high quality-cost; a lower profit implies a firm has a lower output.*

The first result is quite obvious, but the impact of higher product heterogeneity is more subtle.[27] When μ goes up, weak (low quality-cost) firms are helped and good ones are hurt. The intuition is as follows. With little product differentiation, consumers tend to buy the best quality-cost products. With more product differentiation (which increases the mark-up), consumers tend to buy more of the low quality-cost goods (which have lower outputs) and less of the high quality-cost goods (which have higher outputs). Hence, higher μ evens out demands across options. The fact that output may decrease and profit increase with μ follows because

[27] The proof proceeds by showing that the derivative of consumer surplus is given by the Shannon (1948) measure of information (entropy), which is positive.

$\pi_i = \mu y_i$. Thus it can happen that doubling μ does not double the profit of the top quality-cost firms, but it may more than double the profit of the lowest quality-cost firms. Whether high or low qualities are most profitable depends on whether quality-costs rise or fall with quality.

7 Conclusions

The basic ideas here are simple. Market performance depends on the economic fundamentals of tastes and technologies, and how these interact in the market-place.[28] The fundamental distribution of tastes and technologies feeds through the economic process to generate the endogenous distribution of economic variables, such as prices, outputs, and profits. Invoking the monopolistic competition assumption delivers a straight feed-through from fundamental distributions to performance distributions.

Conversely, the derived economic distributions can be reverse engineered to back out the model's primitives. Lemma 6 inverts the mark-ups to deliver both the equilibrium output choices and the form of the demand curve. This analysis constitutes an stand-alone contribution to the theory of monopoly pass-through, extending Weyl and Fabinger (2013) by working from pass-through *back* to implied demands. Theorem 4 shows how to use price and cost distributions to find the shape of the profit and output distributions and demand form (up to a positive factor). Theorem 8 shows how to invert the (potentially observable) output and profit distributions to find the underlying net inverse demand form (i.e., demand up to a cost shift), and the underlying primitive cost distribution, $F_C(c)$. Theorem 3 does likewise with price and profit distributions as the starting point (again up to a positive constant). Theorem 5, 6,

[28] Firm size distributions came to the fore in Chris Anderson's (2006) work on the Long Tail of internet sales.

and 7 shows that all other distribution pairs tie down all primitives and outcome distributions (including constants).

The Logit model gives some similar properties to the CES, while differing on others. For example, the simple CES has constant percentage mark-ups while the Logit has constant absolute mark-ups.[29] The Logit can be deployed for similar purposes as the CES, and has an established pedigree in its micro-economic underpinnings. It has a strong econometric backdrop which is at the heart of much of the structural empirical industrial organization revolution. For Logit, a normal quality-cost distribution leads to a log-normal distribution of firm size, and an exponential quality-cost distribution generates a Pareto distribution. But, our main cross-sectional analysis extends far beyond the restrictive IIA property of Logit and CES.

Indeed, our main results in the heart of the paper show how to back out the demand form from distributions. Surprisingly, this can be done just from profit and price distributions (for example) if firms differ only by production costs.

Another future research direction we only initiated here (in the context of Logit) is the comparative statics of distributions. Hopefully, we have given some pointers how to find how equilibrium distributions change with changes in underlying fundamentals, namely the quality-cost relation and consumer tastes. More can be contributed on understanding the inheritance properties of distributions.

[29] The latter property is perhaps quite descriptive for cinema movies, DVDs, and CDs.

References

[1] Anderson, Chris (2006). The Long Tail: Why the Future of Business Is Selling Less of More. New York: Hyperion.

[2] Anderson, Simon and André de Palma (2000), "From local to Global Competition," European Economic Review, 44, 423-448.

[3] Anderson, Simon P., André de Palma, and Jacques François Thisse (1992). Discrete Choice Theory of Product Differentiation. Cambridge, MA: MIT Press.

[4] Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit Khandelwal, and Eric Verhoogen (2015). Mark-up and Cost Dispersion Across Firms: Direct Evidence from Producer Surveys in Pakistan. American Economic Review Papers & Proceedings, 105(5), 537–544.

[5] Baldwin, Richard, and James Harrigan (2011). "Zeros, Quality, and Space: Trade Theory and Trade Evidence." American Economic Journal: Microeconomics 3(2), 60-88.

[6] Ben-Akiva, Moshe, and Thawat Watanatada (1981). "Application of a Continuous Spatial Choice Logit Model." Structural Analysis of Discrete Data with Econometric Application. Ed. Charles F. Manski and Daniel McFadden. Cambridge, MA: MIT.

[7] Bertoletti, Paolo, and Federico Etro (2016). "Monopolistic competition when income matters." Economic Journal, forthcoming.

[8] Cabral, Luis , and Jose Mata (2003). "On the evolution of the firm size distribution: Facts and theory." American Economic Review, 93(4), 1075-1090.

[9] Caplin, Andrew, and Barry Nalebuff (1991). "Aggregation and imperfect competition: On the existence of equilibrium." Econometrica, 59(1), 25–59.

[10] Chamberlin, Edward (1933). Theory of Monopolistic Competition. Cambridge, MA: Harvard University Press

[11] Crozet, Matthieu, Keith Head, and Thierry Mayer (2011). "Quality Sorting and Trade Firm-level Evidence for French Wine." Review of Economic Studies, 79(2), 609-644.

[12] Dixit, Avinash K., and J. S. Stiglitz (1977). "Monopolistic Competition and Optimal Product Diversity." American Economic Review, 67(3), 297-308.

[13] Dhingra Swati and John Morrow (2013) Monopolistic competition and optimum product diversity under firm heterogeneity, LSE, mimeo.

[14] Eaton, Jonathan, Samuel Kortum, and Francis Kramarz (2011). "An Anatomy of International Trade from French Firms." Econometrica, 79(5), 1453-1498.

[15] Elberse, Anita, and Felix Oberholzer-Gee (2006). Superstars and underdogs: An examination of the long tail phenomenon in video sales. Harvard Business School.

[16] Fajgelbaum, Pablo D., Gene M. Grossman, and Elhanan Helpman (2011). "Income Distribution, Product Quality, and International Trade." Journal of Political Economy, 119(4), 721-765.

[17] Feenstra, Robert C. and John Romalis (2014). International Prices and Endogenous Quality, Quarterly Journal of Economics, 129(2), 477-527.

[18] Fisk, P. (1961). “The Graduation of Income Distributions.” Econometrica, 29(2), 171-185.

[19] Head, Keith, Thierry Mayer, and Mathias Thoenig (2014). “Welfare and Trade without Pareto.” American Economic Review 104(5): 310-16.

[20] Helpman, E., & Krugman, P. R. (1985). Market structure and foreign trade: Increasing returns, imperfect competition, and the international economy. MIT press.

[21] Johnson, Justin P. and David, P. Myatt (2006). “Multiproduct Cournot oligopoly”. The RAND Journal of Economics, 37(3), 583-601.

[22] Mrázová, Monika, J. Peter Neary, and Mathieu Parenti (2016). "Sales and Markup Dispersion: Theory and Empirics," mimeo, University of Oxford.

[23] McFadden, Daniel (1978) “Modelling the Choice of Residential Location.” in A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull (eds.), Spatial Interaction Theory and Planning Models, Amsterdam: North Holland, 75-96, Reprinted in J. Quigley (ed.), 1997, The economics of housing, Vol. I, Edward Elgar: London, 531-552.

[24] Melitz, Marc J. (2003). “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.” Econometrica, 71(6), 1695-1725.

[25] Pareto, Vilfredo (1965). “La Courbe de la Répartition de la Richesse” (Original in 1896): Oeuvres Complètes de Vilfredo. Ed. Busino G. Pareto. Geneva: Librairie Droz, 1-5.

[26] Shannon, Claude. E. (1948). “A Mathematical Theory of Communication,” Bell System Technical Journal, 27(3), 379–423.

[27] Thisse, Jacques-François, and Philip Ushchev (2016). "Monopolistic competition without apology." Mimeo, Higher School of Economics, St. Petersburg.

[28] Weyl, Glen, and Michal Fabinger (2013). "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." Journal of Political Economy, 121(3), 528-583.

[29] Zhelobodko, Evgeny, Sergey Kokovin, Mathieu Parenti, and Jacques-François Thisse (2012). "Monopolistic competition: Beyond the constant elasticity of substitution." Econometrica 80(6), 2765-2784.

Appendix 1

Proof of Proposition 3

We first calculate the logit denominator, D, from (25), using the exponential CDF $F_X(x) = 1-\exp(-\lambda(x-\underline{x}))$ with density $f_X(x) = \lambda \exp(-\lambda(x-\underline{x}))$ and $\lambda > 0$, $\underline{x} > 0$, and $x \in [\underline{x}, \infty)$. Integrating,

$$D = \frac{M\lambda\mu}{\lambda\mu - 1} \exp\left(\frac{\underline{x}}{\mu}\right) + \mathcal{V}_0,$$

which is positive and bounded for any $\mathcal{V}_0$ under the assumption that $\lambda\mu > 1$. Now, from Proposition 2,

$$F_\Pi(\pi) = 1 - \exp\left(-\lambda\left(\mu \ln\left(\frac{\pi D}{\mu}\right) - \underline{x}\right)\right) = 1 - \left(\frac{\pi D}{\mu}\right)^{-\lambda\mu} \exp(\lambda \underline{x}).$$

The profit, $\underline{\pi}$, of the lowest quality-cost firm solves $F_\Pi(\underline{\pi}) = 0$, and thus verifies the expected property $\underline{\pi} = \frac{\mu}{D}\exp\left(\frac{\underline{x}}{\mu}\right)$. Inserting this value back into $F_\Pi(\pi)$ gives the expression in Proposition 3. The output distribution follows from the profit distribution:

$$F_Y(y) = \Pr(Y < y) = \Pr\left(\frac{\Pi}{\mu} < y\right) = F_\Pi(\mu y) = 1 - \left(\frac{\underline{\pi}}{\mu y}\right)^{\lambda\mu} = 1 - \left(\frac{\underline{y}}{y}\right)^{\lambda\mu},$$

where the lowest output, $\underline{y}$, is associated to the lowest profit, $\underline{\pi} = \mu\underline{y}$.

The last statement follows from Theorem 2: starting with a Pareto distribution for output or profit implies an underlying exponential distribution for quality-cost. The lowest quality-cost is given by the condition $\underline{y} = \frac{1}{D}\exp\left(\frac{\underline{x}}{\mu}\right)$, so

$$\underline{\pi} = \mu\underline{y} = \frac{\mu}{\frac{M\lambda\mu}{\lambda\mu-1}\exp\left(\frac{\underline{x}}{\mu}\right) + \mathcal{V}_0} \exp\left(\frac{\underline{x}}{\mu}\right) < \frac{1}{M}\left(\mu - \frac{1}{\lambda}\right). \tag{27}$$

Inverting (27) gives $\underline{x}$.

Proof of Theorem ??

Because $F_X(x)$ is continuous and increasing on support $[\underline{x}, \bar{x}]$ and because $\beta^{-1}(x)$ is continuous and increasing on support $[\underline{x}, \bar{x}]$ with $\beta^{-1}(\underline{x}) > 0$. Then $F_C(c) = \Pr(C < c)$ is uniquely defined and continuous and increasing on support $[\underline{c}, \bar{c}]$:

$$F_C(c) = \Pr\left(\beta^{-1}(X) < c\right) = \Pr(X < \beta(c)) = F_X(\beta(c)).$$

The last term is a continuous and increasing function of a continuous and increasing function, so $F_C(c)$ is recovered. Constructing $F_X(x)$ from $F_C(c)$ and $\beta(c)$ is completely analogous.

We now show how to construct a unique increasing $\beta(c)$ from the two distributions: let $F_X(x) = \Pr(X < x)$ and we postulate that there exists a continuous increasing function $\beta(C) = X$ and so $F_X(x) = \Pr(\beta(C) < x) = \Pr\left(C < \beta^{-1}(x)\right)$ which is then equal to $F_C\left(\beta^{-1}(x)\right)$. Now, since $F_X(x) = F_C\left(\beta^{-1}(x)\right)$, then $\beta^{-1}(x) = F_C^{-1}(F_X(x))$, so $\beta(x) = \left[F_C^{-1}(F_X(x))\right]^{-1}$ and $\beta(x) = F_X^{-1}(F_C(x))$. This is clearly increasing and continuous in x as desired.

The claim in the Theorem is shown because $F_X(x)$ can be used to construct the other distributions on its leg, and can be constructed from them; and likewise for $F_C(c)$.

Appendix 2 Distribution details (NOT FOR PUBLICATION)

Proof of Proposition 2

We first seek the distribution of outputs, $F_Y(y) = \Pr(Y < y)$, that is generated from the primitive distribution of quality-cost. First note from (24) that $Y = \frac{\exp\left(\frac{X}{\mu}\right)}{D}$, so:

$$F_Y(y) = \Pr\left(\frac{\exp\left(\frac{X}{\mu}\right)}{D} < y\right) = F_X(\mu \ln(yD)),$$

where D is given by (25). Because equilibrium profit is proportional to output ($\pi = \mu y$), we have a similar relation for the distribution of profit, $F_\Pi(\pi) = \Pr(\Pi < \pi)$:

$$F_\Pi(\pi) = \Pr\left(\mu\frac{\exp\left(\frac{X}{\mu}\right)}{D} < \pi\right) = F_X\left(\mu \ln\left(\frac{\pi D}{\mu}\right)\right),$$

where D is given by (25).

We next prove the converse result. We first determine the distribution of quality-costs consistent with a given observed distribution of output. Suppose that output has a distribution $F_Y(y)$. Applying the increasing transformation $y = \frac{1}{D}\exp\left(\frac{x}{\mu}\right)$, and $Y = \frac{1}{D}\exp\left(\frac{X}{\mu}\right)$, we get:

$$\begin{aligned} F_X(x) &= \Pr\left(\frac{1}{D}\exp\left(\frac{X}{\mu}\right) < \frac{1}{D}\exp\left(\frac{x}{\mu}\right)\right) \\ &= \Pr\left(Y < \frac{1}{D}\exp\left(\frac{x}{\mu}\right)\right) = F_Y\left(\frac{1}{D}\exp\left(\frac{x}{\mu}\right)\right). \end{aligned}$$

However, D is written in terms of $f_X(x)$, and we want to find the distribution solely in terms of $F_Y(y)$: this means writing D in terms of $f_Y(y)$. The corresponding expression, denoted D_y

is derived below as (28). Similar reasoning gives the profit expression:

$$F_X(x) = \Pr(X < x) = \Pr\left(\Pi < \frac{\mu}{D}\exp\left(\frac{x}{\mu}\right)\right) = F_\Pi\left(\frac{\mu}{D}\exp\left(\frac{x}{\mu}\right)\right),$$

where the expression for D in terms of $f_\Pi(\pi)$ (i.e., D_π) is given in the Theorem and derived below as (29).

We now show here how to write the function D as a function of $f_Y(.)$ or $f_\pi(.)$.

We first find the value of D in terms of the distribution of Y. Recall (25):

$$D = M\int_{u\geq \underline{x}} \exp\left(\frac{u}{\mu}\right) f_X(u)\, du + \mathcal{V}_0.$$

Now, $\Pr(x < X) = \Pr(y < Y)$, so $y(x) = \frac{\exp\left(\frac{x}{\mu}\right)}{M\int\exp\left(\frac{u}{\mu}\right)f_X(u)du+\mathcal{V}_0}$; hence $D\int y(u)\, f_Y(u)\, du = (D - \mathcal{V}_0)/M$, and thus:

$$D_y = \frac{\mathcal{V}_0}{1 - My_{av}}. \tag{28}$$

The denominator in the last expression is necessarily positive because My_{av} is total output, which is less than one when $\mathcal{V}_0 > 0$ because the market is not fully covered. Similarly, $F_X(x) = F_\Pi\left(\frac{\mu}{D}\exp\left(\frac{x}{\mu}\right)\right)$, so that

$$D_\pi = \frac{\mathcal{V}_0}{1 - \frac{M}{\mu}\pi_{av}}, \tag{29}$$

which is now expressed as a function of $f_\Pi(.)$, and where π_{av} is average firm profit. (29) is positive because total profit, $M\pi_{av}$, is less than μ because the market is not fully covered ($\mathcal{V}_0 > 0$).

Study of specific distributions

We now derive the distributions in Section 6.2: these involve parameter matching for the distribution examples.

Normal: For the normal, $F_X(x) = \frac{1}{\sigma\sqrt{2\tilde{\pi}}}\int_{-\infty}^{x}\exp\left(-\frac{(u-m)^2}{2\sigma^2}\right)du$, where $\tilde{\pi} = 3.1415...$. From Theorem 2, we have

$$F_\Pi(\pi) = F_X\left(\mu\ln\left(\frac{\pi D}{\mu}\right)\right),$$

where $\mu\ln\left(\frac{\pi D}{\mu}\right) \in (-\infty,\infty)$, so

$$F_\Pi(\pi) = F_X\left(\mu\ln\left(\frac{\pi D}{\mu}\right)\right) = \frac{1}{\sigma\sqrt{2\tilde{\pi}}}\int_{-\infty}^{\mu\ln\left(\frac{\pi D}{\mu}\right)}\exp\left(-\frac{(u-m)^2}{2\sigma^2}\right)du.$$

Using the change of variable $\Pi = \frac{\mu}{D}\exp\left(\frac{u}{\mu}\right)$ (so $u = \mu\ln\left(\frac{\Pi D}{\mu}\right)$ and $du = \frac{\mu}{\Pi}d\Pi$) we obtain

$$F_\Pi(\pi) = \frac{\mu}{\sigma\sqrt{2\tilde{\pi}}}\int_0^{\pi}\exp\left(-\frac{\left(\mu\ln\left(\frac{\Pi D}{\mu}\right)-m\right)^2}{2\sigma^2}\right)\frac{d\Pi}{\Pi},$$

which can be written in a standard form as:

$$F_\Pi(\pi) = \frac{1}{\left(\frac{\sigma}{\mu}\right)\sqrt{2\tilde{\pi}}}\int_0^{\pi}\frac{1}{\Pi}\exp\left(-\frac{\left(\ln\Pi-\left(\ln\left(\frac{D}{\mu}\right)-\frac{m}{\mu}\right)\right)^2}{2\left(\frac{\sigma}{\mu}\right)^2}\right)d\Pi.$$

Hence profits are log-normally distributed with parameters $\left[\ln\left(\frac{D}{\mu}\right)-\frac{m}{\mu}\right]$ and $\left(\frac{\sigma}{\mu}\right)$.

Recall: $D = M\int_{u\geq \underline{x}}\exp\left(\frac{u}{\mu}\right)f_X(u)\,du + \mathcal{V}_0$. Then:

$$D = \frac{M}{\sigma\sqrt{2\tilde{\pi}}}\int_{-\infty}^{\infty}\exp\left(\frac{x}{\mu}\right)\exp\left(\frac{-(x-m)^2}{2\sigma^2}\right)dx + \mathcal{V}_0.$$

Routine computation shows that:

$$D = M\exp\left(\frac{m}{\mu}+\frac{\sigma^2}{2\mu^2}\right)+\mathcal{V}_0.$$

Logistic.

A *logistic* distribution for quality-cost has a CDF given by $F_X(x) = \left(1 + \exp\left(-\frac{x-m}{s}\right)\right)^{-1}$, $x \in (\underline{x}, \infty)$, with mean m and variance $s^2\pi^2/3$. The PDF is similar in shape to the normal, but it has thicker tails (see the discussion in Fisk, 1961, and the comparison with the Weibull distribution). Hence, for $\mu > s$, profit $\Pi \in (0, \infty)$ is log-logistically distributed with parameters $\left(\frac{D}{\mu}\right) \exp\left(-\frac{m}{\mu}\right)$ and $\frac{\mu}{s}$:

$$F_\Pi(\pi) = \left(1 + \left(\frac{\pi}{\left(\frac{D}{\mu}\right) \exp\left(-\frac{m}{\mu}\right)}\right)^{-\frac{\mu}{s}}\right)^{-1}, \qquad \pi \in [0, \infty).$$

There is no closed form expression for D in this case. However, it can be shown that the condition $\mu > s$ guarantees that the output denominator D exists. The Log-logistic distribution (which provides a one parameter model for survival analysis) is very similar in shape to the log-normal distribution, but it has fatter tails. It has an explicit functional form, in contrast to the Log-normal distribution.

The logistic distribution (with mean m and standard deviation $s\tilde{\pi}/\sqrt{3}$) is given by:

$$F_X(x) = \frac{1}{1 + \exp\left(-\frac{x-m}{s}\right)}, x \in (\underline{x}, \infty).$$

From Theorem 2,

$$F_\Pi(\pi) = F_X\left(\mu \ln\left(\frac{\pi D}{\mu}\right)\right) = \frac{1}{1 + \exp\left(-\frac{\mu \ln\left(\frac{\pi D}{\mu}\right) - m}{s}\right)},$$

where $F_{\Pi}(0) = 0$ and $F_{\Pi}(\infty) = 1$. Thus

$$F_{\Pi}(\pi) = \frac{1}{1 + \exp\left(\frac{m}{s}\right)\exp\left(\ln\left(\frac{\pi D}{\mu}\right)^{-\frac{\mu}{s}}\right)} = \frac{1}{1 + \left(\frac{\pi}{\left(\frac{D}{\mu}\right)\exp\left(-\frac{m}{\mu}\right)}\right)^{-\frac{\mu}{s}}}.$$

Recall the log-logistic distribution is defined as:

$$F^{LL}(x;\alpha,\beta) = \frac{1}{1 + \left(\frac{x}{\alpha}\right)^{-\beta}}, \quad x > 0.$$

Thus, the parameter matching is:

$$F_{\Pi}(\pi) = F^{LL}\left(x; \left(\frac{D}{\mu}\right)\exp\left(-\frac{m}{\mu}\right), \frac{\mu}{s}\right).$$

We need to check when D converges, i.e., when $\int_{-\infty}^{\infty} \exp\left(\frac{x}{\mu}\right) f_X(x)\, dx$ converges. Because

$$f_X(x) = \frac{1}{s}\frac{\exp\left(-\frac{x-m}{s}\right)}{\left(1 + \exp\left(-\frac{x-m}{s}\right)\right)^2},$$

we need to ensure the convergence of the expression

$$\int_{-\infty}^{\infty} \frac{\exp\left(-x\left(\frac{1}{s} - \frac{1}{\mu}\right)\right)}{\left(1 + \exp\left(-\frac{x-m}{s}\right)\right)^2} dx.$$

Convergence is guaranteed if and only if $\mu > s$.

Pareto: The Pareto distribution is given by: $F_X(x) = \frac{1-\left(\frac{\underline{x}}{x}\right)^{\alpha}}{1-\left(\frac{\overline{x}}{\underline{x}}\right)^{-\alpha}}$. From Theorem 2

$$F_{\Pi}(\pi) = \frac{1 - \left(\frac{\underline{x}}{\mu \ln\left(\frac{\pi D}{\mu}\right)}\right)^{\alpha}}{1 - \left(\frac{\overline{x}}{\underline{x}}\right)^{-\alpha}}.$$

Recall that $\pi = \frac{\mu \exp\left(\frac{x}{\mu}\right)}{D}$ or $x = \mu \ln\left(\frac{\pi D}{\mu}\right)$, so that $\underline{x} = \mu \ln\left(\frac{\underline{\pi} D}{\mu}\right)$ and $\overline{x} = \mu \ln\left(\frac{\overline{\pi} D}{\mu}\right)$, so that $F_{\Pi}(\underline{\pi}) = 0$ and $F_{\Pi}(\overline{\pi}) = 1$. D is bounded because the quality-cost distribution is bounded.

Consider a log-Pareto distribution with scale parameter σ and shape parameters γ and β:

$$F^{LP}(\pi; \gamma, \beta, \sigma) = \frac{1 - \left(1 + \frac{1}{\beta}\ln\left(1 + \frac{\pi - \underline{\pi}}{\sigma}\right)\right)^{-\frac{1}{\gamma}}}{1 - \left(1 + \frac{1}{\beta}\ln\left(1 + \frac{\overline{\pi} - \underline{\pi}}{\sigma}\right)\right)^{-\frac{1}{\gamma}}}, \qquad \pi > \underline{\pi}.$$

In order to match parameters, of $F_{\Pi}(\pi)$ with $F^{LP}(\pi; \gamma, \beta, \sigma)$ observe that

$$\mu \ln\left(\frac{\pi D}{\mu}\right) = \mu \ln\left(\frac{\pi D}{\mu}\frac{\underline{\pi}}{\underline{\pi}}\right) = \underline{x} + \mu \ln\left(\frac{\pi}{\underline{\pi}}\right) = \underline{x} + \mu \ln\left(1 + \frac{\pi - \underline{\pi}}{\underline{\pi}}\right).$$

Therefore:

$$F_{\Pi}(\pi) = \frac{1 - \underline{x}^{\alpha}\left(\mu \ln\left(\frac{\pi D}{\mu}\right)\right)^{-\alpha}}{1 - \left(\frac{\overline{x}}{\underline{x}}\right)^{-\alpha}} = \frac{1 - \left(1 + \frac{\mu}{\underline{x}}\ln\left(1 + \frac{\pi - \underline{\pi}}{\underline{\pi}}\right)\right)^{-\alpha}}{1 - \left(\frac{\overline{x}}{\underline{x}}\right)^{-\alpha}}.$$

Thus π obeys a Log-Pareto distribution $F^{LP}\left(\pi; \frac{1}{\alpha}, \frac{\underline{x}}{\mu}, \underline{\pi}\right)$, i.e., $\gamma = \frac{1}{\alpha}, \beta = \frac{\underline{x}}{\mu}, \sigma = \underline{\pi}$.

It remains to check that the normalization factors are equal. Recall that $\overline{\pi}/\underline{\pi} = \exp\frac{\overline{x} - \underline{x}}{\mu}$. Using the specification $\gamma = \frac{1}{\alpha}, \beta = \frac{\underline{x}}{\mu}, \sigma = \underline{\pi}$, we get:

$$\begin{aligned}
\left(1 + \frac{\mu}{\overline{x}}\ln\left(1 + \frac{\overline{\pi} - \underline{\pi}}{\underline{\pi}}\right)\right)^{-\alpha} &= \left(1 + \frac{\mu}{\overline{x}}\ln\left(\frac{\overline{\pi}}{\underline{\pi}}\right)\right)^{-\alpha} \\
&= \left(1 + \left(\frac{\overline{x} - \underline{x}}{\overline{x}}\right)\right)^{-\alpha} = \left(\frac{\overline{x}}{\underline{x}}\right)^{-\alpha}.
\end{aligned}$$