



Contents lists available at ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

The properties of product line prices☆

Justin P. Johnson^{a,*}, David P. Myatt^{b,**}^a Samuel Curtis Johnson Graduate School of Management, Cornell University, United States^b London Business School, United Kingdom

ARTICLE INFO

Available online xxxx

JEL classification:

D42

D43

L11

Keywords:

Product lines

Price discrimination

Multi-product competition

ABSTRACT

We study multi-product Cournot suppliers (a monopoly is a special case) who offer menus of differentiated qualities. Under conditions that include linear and constant-elasticity demand specifications, the equilibrium price for each quality is equal to its price in a single-product Cournot market. A corollary is this: a monopolist who naively ignores the substitutability of her products, by choosing the standalone monopoly price for each quality, successfully implements the optimal discriminatory scheme. Under other reasonable specifications, the equilibrium prices are close to the single-product prices. We also evaluate the impact of minimum quality standards.

© 2015 Elsevier B.V. All rights reserved.

Students are taught that maximizing the profit from a product requires the equalization of marginal revenue and marginal cost. Advanced students are told that optimizing an entire menu of quality-differentiated products is more complex: the substitutability of the products means that a change in the supply of one quality influences the profits earned from others.

If a multi-product monopolist were to ignore this substitutability, then would she be harmed? That is, are her profits reduced by using only the rudimentary techniques of single-product profit maximization? Relatedly, do the equilibrium prices in a multi-product oligopoly differ from the prices that we would see in single-product markets?

Under certain “textbook” conditions, the answer to these questions is: No. Specifically, the equilibrium prices in a multi-product Cournot oligopoly are equal to those in single-product markets. A monopolized market is a special case: a monopolist’s optimal discriminatory prices are equal to the corresponding standalone single-product monopoly prices.

The pricing equivalence result holds when the Mills ratio (the reciprocal of the hazard rate) of the distribution of buyers’ types

(a buyer’s type is his willingness to pay for increases in quality) is linear in the type. (Formally: if a consumer of type $\theta \sim F(\cdot)$ is willing to pay θq for quality q , then $[1 - F(\theta)]/f(\theta)$ should be linear in θ .) This holds if types are either uniformly or (generalized) Pareto distributed, and is equivalent to a linear relationship between the marginal revenue of a single product and its price. This holds if the inverse demand curve for a product exhibits constant curvature, as it does when demand is linear or exhibits constant elasticity.

To develop intuition, consider a monopolist airline selling a two-step product line comprising economy and business travel classes. These products are substitutes, and so the price of one travel class influences the profits earned from both qualities. However, a business-class ticket can be considered to be a bundle of an economy ticket and a supplemental quality upgrade. This perspective is useful because the bundle’s components are neither complements nor substitutes. That is, a local change in the upgrade price influences the demand for the upgrade to business class (which is purchased in addition to an economy seat) but this change does not influence the marginal buyer of an economy ticket. Similarly, a local change to the economy price does not influence the marginal buyer of the upgrade.

Because the “baseline” economy product and the “upgrade” to business class are neither substitutes nor complements, to solve her pricing problem the monopolist may optimize separately in each of the baseline and upgrade markets. The correct product-line price for business class is the sum of the economy and upgrade prices.

Now consider the abolition of economy class. This forces the supplier to sell the baseline product and the upgrade together as a compulsory (or pure) bundle. The marginal cost of this bundle is the sum of the marginal costs of its components. At the original optimum, this sum equals the sum of the marginal revenue terms from the economy and upgrade markets. If there is a linear relationship between marginal revenue and

☆ Acknowledgments. We consulted many colleagues to discover whether the price equivalence result is widely known. In particular, we acknowledge Mark Armstrong, Heski Bar-Isaac, Simon Board, João Montez, Andrew Rhodes, Chris Wallace, and Glen Weyl, for their helpful suggestions. We also offer thanks to Ramon Casadesus-Masanell and Glenn MacDonald for their invitation to present at the CRES Foundations of Business Strategy Conference, where Dylan Minor served as an outstanding discussant, to Heski Bar-Isaac for the opportunity to participate in an EARIE 2014 invited session, and to two anonymous referees.

* Correspondence to: J.P. Johnson, Samuel Curtis Johnson Graduate School of Management, Cornell University, Ithaca NY 14853, United States.

** Correspondence to: D.P. Myatt, London Business School, Regent’s Park, London NW1 4SA, UK.

E-mail addresses: jjp25@cornell.edu (J.P. Johnson), dmyatt@london.edu (D.P. Myatt).

price then the marginal revenue for the economy-plus-upgrade combination is equal to the sum of the separate economy and upgrade marginal revenue terms, when evaluated at the sum of the original component prices. This implies that the original business-class price remains optimal in what is now a business-class-only single-product market. Similarly, abolishing business class does not change the optimal economy-class ticket price. This reasoning also holds in a Cournot oligopoly, and so the equilibrium multi-product prices are the same as the prices that emerge from separate single-product Cournot markets.

Central to the pricing neutrality result is the property that there a linear relationship between marginal revenue and price. This corresponds to the requirement that the Mills ratio is linear. With this in mind, we also investigate the impact on prices and profits when the Mills ratio is not linear, considering numerically both beta and log-normal specifications. Although the non-linearity implies that the equilibrium multi-product prices do differ from the corresponding single-product prices, the difference is often very small. Additionally, the limited circumstances under which the price differences are larger are situations in which a supplier gains the least from selling multiple products rather than a single quality.

We contribute to the established techniques for quality-based and quantity-based price discrimination (Maskin and Riley, 1984; Mussa and Rosen, 1978). The upgrades approach (Johnson and Myatt, 2003; Johnson and Myatt, 2006a, 2006b) views higher qualities as combinations of incremental upgrades. Itoh (1983) implicitly used such an approach, and some of our findings concerning monopoly are re-discoveries of his insufficiently widely known results. We contribute by developing the implications for pricing naiveté, by extending to Cournot oligopolies, and by showing numerically that, even when we move away from common textbook specifications, there is often little difference between multi-product and single-product prices. This and other recent papers (Aguirre et al., 2010; Anderson and Çelik, 2015; Anderson and Dana, 2009; Cowan, 2007, 2012) reflect a resurgent interest in price discrimination. Demand curvature is a central feature; the treatment of this and pass-through properties (Bulow and Pfleiderer, 1983) has been developed recently by Weyl and Fabinger (2013).

1. Supply and demand for multiple qualities

1.1. Demand

A buyer with type $\theta \sim F(\theta)$ is willing to pay at most θq for a single unit of a product with quality $q \in [0, \infty)$. The continuous distribution $F(\cdot)$ has a strictly positive density $f(\cdot)$. The Mills ratio $M(\cdot)$ is the reciprocal of the hazard rate:

$$M(\theta) \equiv \frac{1 - F(\theta)}{f(\theta)}. \tag{1}$$

A buyer θ purchases a single unit of the product that offers him the greatest positive surplus.

$1 - F(p)$ is the quantity demanded at price p for a standalone product with normalized quality $q = 1$, and $P(z) \equiv F^{-1}(1 - z)$ is the associated inverse demand function at $z \in (0, 1)$. In an m -supplier Cournot single-product oligopoly selling quality $q = 1$, the marginal revenue gained from an output expansion in a symmetrically divided market is

$$\text{Marginal Revenue} = P(z) + \frac{zP'(z)}{m} = p - \frac{M(p)}{m} \quad \text{where } p = P(z). \tag{2}$$

We assume that marginal revenue is decreasing in output; equivalently, $M'(\theta) < m$.

A particular case of interest is when $M(\theta)$ is linear, so that there is a linear relationship between marginal revenue and price. For example, if $\theta \sim U[\theta_L, \theta_H]$ then

$$F(\theta) = \frac{\theta - \theta_L}{\theta_H - \theta_L} \Rightarrow M(\theta) = \theta_H - \theta \tag{3}$$

$$\Rightarrow \text{Marginal Revenue} = \frac{(m + 1)p - \theta_H}{m}.$$

Another case of interest is the generalized (type II) Pareto distribution:

$$F(\theta) = 1 - \left(\frac{\theta - \mu}{\sigma/\xi}\right)^{-1/\xi} \Rightarrow M(\theta) = \xi(\theta - \mu) \tag{4}$$

$$\Rightarrow \text{Marginal Revenue} = \frac{(m - \xi)p + \xi\mu}{m},$$

where the location, scale, and shape are determined by μ , σ , and ξ , and where $\theta \geq \mu + (\sigma/\xi)$. Setting $\mu = 0$ generates a demand function with constant elasticity.

The linearity of the Mills ratio also imposes structure on the inverse demand for a product. For example, under the generalized Pareto specification, $P(z) \equiv F^{-1}(1 - z)$ is

$$P(z) = \mu + \frac{\sigma z^{-\xi}}{\xi} \Rightarrow -\frac{zP'(z)}{P(z)} = 1 + \xi. \tag{5}$$

This final expression is the curvature (that is, the elasticity of the slope) of the inverse-demand function. Robinson (1933) called this the “adjusted concavity” of inverse demand.

Definition (Constant curvature). *The demand for quality has constant curvature if the elasticity of the slope of inverse demand for a single product is constant.*

As noted above, this property is satisfied by the linear and constant-elasticity specifications. If demand has constant curvature then, under both monopoly and Cournot oligopoly, there is a constant pass-through rate for changes in marginal cost (Bulow and Pfleiderer, 1983).

Although the class of constant-curvature demand specifications includes leading textbook examples, it can be restrictive. If a distribution has an increasing hazard rate and unbounded support, then necessarily the Mills ratio is non-linear.¹ Many common specifications (such as the unimodal beta, the normal, and gamma distributions) have a convex Mills ratio. Nevertheless, and as we show later, common specifications often exhibit approximate linearity.

1.2. Supply

The constant marginal cost of producing quality q is $c(q)$. This satisfies $c(0) = 0$, $c'(q) > 0$, and $c''(q) > 0$: there are decreasing returns to enhanced quality.

There are n quality levels $0 < q_1 < \dots < q_n$.² $c_i \equiv c(q_i)$ is the constant marginal cost of quality q_i , $\Delta q_i \equiv q_i - q_{i-1}$, $\Delta c_i \equiv c_i - c_{i-1}$, $q_0 \equiv 0$ and $c_0 \equiv 0$. The properties of $c(q)$ ensure that $0 < \Delta c_1/\Delta q_1 < \dots < \Delta c_n/\Delta q_n$. Finally, $0 < F(\Delta c_1/\Delta q_1)$ (it is inefficient to supply everyone) and $F(\Delta c_n/\Delta q_n) < 1$ (the efficient supply of the highest quality is positive).

¹ An increasing hazard implies a decreasing Mills ratio. If $M(\theta)$ were linear then it would eventually cross zero.

² There is no limit to n , and our results can also be applied to a continuum-of-qualities specification. Recent complementary work by Anderson and Çelik (2015) elegantly characterizes optimal product lines.

The market is served by m multi-product Cournot oligopolists. To keep our notation simple, we mainly focus on symmetric suppliers with the same production technologies. However, our results are robust to some differences in their technological capabilities (Proposition 2).

2. Equilibrium prices in a multi-product Cournot oligopoly

2.1. Benchmark: a single-product equilibrium

Before characterizing the equilibrium with multiple products, as a benchmark we consider a Cournot game in which the suppliers sell a single product with quality q_i . Given our assumptions, the equilibrium is symmetric. The total industry output z_i^\dagger equates marginal revenue to marginal cost in the usual way:

$$q_i \left[P(z_i^\dagger) + \frac{z_i^\dagger P'(z_i^\dagger)}{m} \right] = c_i, \tag{6}$$

where we note that $P(z)$ is the inverse-demand curve for a product with quality $q = 1$, and so in general marginal revenue is proportional to the product's quality. Equivalently, the marginal buyer's type is θ_i^\dagger where $z_i^\dagger = 1 - F(\theta_i^\dagger)$. He is indifferent, and so

$$p_i^\dagger = \theta_i^\dagger q_i \quad \text{where} \quad \theta_i^\dagger - \frac{M(\theta_i^\dagger)}{m} = \frac{c_i}{q_i}. \tag{7}$$

If the Mills ratio is linear then this yields a simple closed-form solution to the Cournot price.³

Lemma 1 (Single-product Cournot prices). *If demand has constant curvature, so that $M(\theta) = \alpha + \beta\theta$ for some α and $\beta < m$, then the equilibrium price in a Cournot market for quality q_i is*

$$p_i^\dagger = \frac{mc_i + \alpha q_i}{m - \beta}. \tag{8}$$

2.2. Multi-Product Equilibrium

In a multi-product Cournot oligopoly, the m symmetric suppliers simultaneously choose their outputs of the n different qualities, where z_{ik} is supplier k 's output of quality q_i . At the industry level, z_i is the total output of quality q_i , p_i is its price, and we define $\Delta p_i \equiv p_i - p_{i-1}$.

Market-clearing prices ensure that higher types purchase higher qualities. The marginal buyer of quality q_i is the type θ_i with $Z_i \equiv \sum_{j=i}^n z_j$ others above him, and so $Z_i = 1 - F(\theta_i)$. He is indifferent between qualities q_i and q_{i-1} , and so $\Delta p_i \equiv p_i - p_{i-1} = \theta_i \Delta q_i$.

Z_i is the fraction of products sold with quality q_i or higher. This can be interpreted as the supply of upgrades to this quality level (and beyond). Similarly, Δp_i is the price of an upgrade from quality q_{i-1} to the next step in the product line.

Supplying n different qualities is equivalent to selling n upgrades. From an upgrades perspective, quality q_i is a combination of the i upgrades from Δq_1 to Δq_i . An upgrade's price Δp_i depends only on its own supply, and so the advantage of thinking in terms of upgrades is that they are neither substitutes nor complements.

We write Z_{ik} for the quantity of the i th upgrade offered by supplier k . Supplier k must choose her upgrades to satisfy $Z_{1k} \geq \dots \geq Z_{nk} \geq 0 \equiv Z_{(n+1)k}$. She seeks to maximize:

$$\begin{aligned} \text{Multi-Product Profit of Supplier } k &= \sum_{i=1}^n Z_{ik}(p_i - c_i) \\ &= \sum_{i=1}^n (Z_{ik} - Z_{(i+1)k}) \left(\sum_{j=1}^i (\Delta p_j - \Delta c_j) \right) \\ &= \sum_{i=1}^n Z_{ik} (\Delta p_i - \Delta c_i) \\ &= \sum_{i=1}^n \Delta q_i Z_{ik} \left(P(Z_i) - \frac{\Delta c_i}{\Delta q_i} \right). \end{aligned} \tag{9}$$

The i th element of this summation depends only upon Z_{ik} and Z_i . Hence, the Nash equilibrium of a multi-product Cournot game replicates n single-product Nash equilibria in each upgrade market, so long as these upgrade supplies, which we denote Z_{ik}^* , satisfy the monotonicity constraints $Z_{1k}^* \geq \dots \geq Z_{nk}^*$ (so that $z_{ik}^* = Z_{ik}^* - Z_{(i+1)k}^* \geq 0$). These constraints are satisfied when there are decreasing returns to quality (Johnson and Myatt, 2003, 2006a, 2006b).

Lemma 2 (from Johnson and Myatt 2006a). *There is a unique Nash equilibrium of the multi-product Cournot oligopoly game. The total industry output Z_i^* of the i th upgrade satisfies*

$$\Delta q_i \left[P(Z_i^*) + \frac{Z_i^* P'(Z_i^*)}{m} \right] = \Delta c_i, \tag{10}$$

and the equilibrium industry output of quality q_i itself is $z_i^* = Z_i^* - Z_{i+1}^*$.

2.3. Pricing Equivalence

In equilibrium the marginal consumer of quality q_i (and hence of the i th upgrade) has a type $\theta_i^* = P(Z_i^*)$. The equilibrium price of that upgrade is

$$\Delta p_i^* = \theta_i^* \Delta q_i \quad \text{where} \quad \theta_i^* - \frac{M(\theta_i^*)}{m} = \frac{\Delta c_i}{\Delta q_i}. \tag{11}$$

A special case is when demand exhibits constant curvature, so that $M(\theta) = \alpha + \beta\theta$ for some α and $\beta < m$. In this case, the equilibrium price for the i th upgrade is

$$\Delta p_i^* = \frac{m\Delta c_i + \alpha\Delta q_i}{m - \beta}. \tag{12}$$

To obtain the equilibrium price of quality q_i^* , we simply sum the equilibrium upgrade prices. If the demand for quality has constant curvature, then

$$p_i^* = \sum_{j=1}^i \Delta p_j^* = \frac{\sum_{j=1}^i (m\Delta c_j + \alpha\Delta q_j)}{m - \beta} = \frac{mc_i + \alpha q_i}{m - \beta} = p_i^\dagger, \tag{13}$$

which (from Eq. (8) of Lemma 1) is the equilibrium price in a single-product market.

Proposition 1 (Pricing Equivalence). *If the demand for quality has constant curvature, then the equilibrium price for a particular quality in a multi-product Cournot oligopoly coincides with the corresponding equilibrium price in a standalone single-product Cournot market.*

This proposition holds for any product line. Thus, for any two different product lines containing the same quality, the price of that common quality must be the same.

³ Equivalently: there is a linear relationship between marginal revenue and price.

Corollary (to Proposition 1). *If the demand for quality has constant curvature, then the equilibrium price for each quality in a product line is independent of the other qualities that are available.*

A special case is when $m = 1$. If a monopolist naively specifies single-product monopoly prices for her qualities and if her demand is linear or exhibits constant elasticity, then she replicates the fully optimal prices from *Mussa and Rosen (1978)*.⁴ This simple result is not widely known, even though it follows straightforwardly from Proposition 5 of *Itoh (1983)*.

2.4. Asymmetric oligopolies

Proposition 1 also holds when suppliers have asymmetric capabilities, so long as they do not differ too much. Suppose that cost asymmetries are such that in equilibrium all suppliers offer complete product lines (so that $z_{ik} > 0$ for all i and k). This is true if the suppliers' production technologies are not too different. With complete product lines, each supplier satisfies the relevant first-order condition in each upgrade market. Writing Δc_{ik} for the marginal cost of upgrade i to supplier k , equilibrium prices satisfy

$$\Delta p_i^* = \theta_i^* \Delta q_i \quad \text{where} \quad \theta_i^* - \frac{M(\theta_i^*)}{m} = \frac{\sum_{k=1}^m \Delta c_{ik}}{m \Delta q_i} \tag{14}$$

We can sum over the upgrades, just as we did in the derivation of Eq. (13).

Proposition 2 (Pricing equivalence in an asymmetric oligopoly). *If, in equilibrium, all suppliers offer complete product lines then the claim of Proposition 1 holds in an asymmetric oligopoly.*

If competing suppliers offer different product lines, then typically the pricing equivalence result does not hold. If quality q_i is omitted from supplier k 's product line then the constraint $Z_{ik} \geq Z_{(i+1)k}$ binds. This means that the first-order condition for the i th upgrade no longer holds for this supplier. A consequence is that Eq. (14) no longer applies.

3. The effect of product-line changes

We now study the response of equilibrium prices to changes in the menu of available qualities. For the monopoly case, such an exercise was conducted by *Itoh (1983)*. A contribution here is to show that his results also hold for a Cournot oligopoly.

3.1. Bundling upgrades

Consider two neighboring upgrades $j \in \{i, i + 1\}$ in the product line. Their equilibrium prices satisfy Eq. (11). Combining the solutions,

$$\frac{\Delta q_i \theta_i^* + \Delta q_{i+1} \theta_{i+1}^*}{\Delta q_i + \Delta q_{i+1}} - \frac{\Delta q_i M(\theta_i^*) + \Delta q_{i+1} M(\theta_{i+1}^*)}{m(\Delta q_i + \Delta q_{i+1})} = \frac{\Delta c_i + \Delta c_{i+1}}{\Delta q_i + \Delta q_{i+1}} \tag{15}$$

Now suppose that the two upgrades are forcibly bundled to form a two-step upgrade. The equilibrium price for this bundle is $\Delta \hat{p} = \hat{\theta}(\Delta q_i + \Delta q_{i+1})$, where $\hat{\theta}$ is the unique solution to

$$\hat{\theta} - \frac{M(\hat{\theta})}{m} = \frac{\Delta c_i + \Delta c_{i+1}}{\Delta q_i + \Delta q_{i+1}} \tag{16}$$

⁴ The problem differs slightly from *Mussa and Rosen (1978)* because the product line is finite. However, the qualities intervals can be arbitrarily narrow, and our approach does work in the continuum-of-qualities case.

Now suppose that $M(\cdot)$ is strictly convex. Using Eqs. (15) and (16),

$$\begin{aligned} \hat{\theta} - \frac{M(\hat{\theta})}{m} &= \frac{\Delta q_i \theta_i^* + \Delta q_{i+1} \theta_{i+1}^*}{\Delta q_i + \Delta q_{i+1}} - \frac{\Delta q_i M(\theta_i^*) + \Delta q_{i+1} M(\theta_{i+1}^*)}{m(\Delta q_i + \Delta q_{i+1})} \\ &< \frac{\Delta q_i \theta_i^* + \Delta q_{i+1} \theta_{i+1}^*}{\Delta q_i + \Delta q_{i+1}} - \frac{1}{m} M\left(\frac{\Delta q_i \theta_i^* + \Delta q_{i+1} \theta_{i+1}^*}{\Delta q_i + \Delta q_{i+1}}\right) \\ &\Rightarrow \hat{\theta} < \frac{\Delta q_i \theta_i^* + \Delta q_{i+1} \theta_{i+1}^*}{\Delta q_i + \Delta q_{i+1}} \Rightarrow \Delta \hat{p} < \Delta p_i^* + \Delta p_{i+1}^* \end{aligned} \tag{17}$$

The first inequality follows from the convexity of $M(\cdot)$, the implication holds because $\theta - \frac{M(\theta)}{m}$ is increasing in θ , and the conclusion follows from the definitions of $\Delta \hat{p}$, Δp_i^* , and Δp_{i+1}^* .

Lemma 3 (Combining upgrades). *If the Mills ratio is convex (respectively, concave) then the price for a bundle of upgrades is less than (respectively, more than) the sum of the upgrade prices.*

The upgrade bundle combines commonly shaped demand curves. The sum of the original marginal revenue terms equals the sum of the marginal costs. However, for the bundle the marginal revenue terms are now evaluated at an appropriately weighted average type. If marginal revenue is a concave function of price then, at this average point, marginal revenue is greater; this induces an output expansion and so a lower price.

3.2. Product line changes

The elimination of quality q_i forces the bundling of the upgrades Δq_i and Δq_{i+1} . The prices of qualities below q_i are unaffected; such lower qualities are combinations of upgrades Δq_j for $j < i$. However, if the price of the bundle $\Delta q_i + \Delta q_{i+1}$ falls then the price of q_{i+1} must fall. There are no changes to upgrades above this and so other complete qualities above q_{i+1} must experience the same change in price seen for p_{i+1} .⁵

Proposition 3 (The effect of product line changes). *The elimination of a quality does not change the equilibrium prices of lower qualities, but results in a common price change for higher qualities. If the Mills ratio is convex (respectively, concave) then prices of higher qualities fall (respectively, rise).*

This proposition also reveals the relationship between equilibrium product-line prices and the corresponding single-product prices when the curvature of demand is non-constant. Consider the price of quality q_i . If all higher qualities are removed then its price is unaffected; however, if all lower qualities are removed and if the Mills ratio is convex then its price falls. Once all other qualities are removed then its price is the single-product price.

Corollary (to Proposition 3). *If the Mills ratio is convex (respectively, concave) then the equilibrium price for a quality in a product line is higher (respectively, lower) than its standalone price. In both cases the gap between the multi-product and standalone prices is higher for higher qualities.*

The final claim holds because higher qualities are the sum of a larger number of upgrade steps, and so their prices are influenced by more bundling operations.

⁵ This implies that the claims of *Proposition 3* from *Itoh (1983)* apply in a Cournot environment. This is in the context of a multiplicative specification for consumers' preferences, so that the willingness of type θ to pay for quality q is $u(\theta, q) = \theta q$. However, for general $u(\theta, q)$ it is also true that the prices of qualities below q_i are unaffected, and the prices of qualities q_{i+1} change by the same amount. However, the size of the change depends upon the properties of $u(\theta, q)$ as well as the shape of the Mills ratio $M(\theta)$.

3.3. Quality regulations

The regulation of a multi-product market might take several forms, including bounds to the qualities offered, or the requirement to offer only one quality.

A restriction to a single quality directly removes choice from buyers. However, if the price of the remaining product falls, and falls by enough, then some buyers may benefit. If the demand for quality exhibits constant curvature then (applying Proposition 1) the price of the remaining product is unchanged, and so nobody is helped. If the Mills ratio is concave, then the remaining single product has a higher price, and so everyone is harmed. The only situation in which some (but not necessarily all) buyers benefit is when the Mills ratio is convex. A regulatory restriction to a single quality helps an interval of buyer types that includes all of the original purchasers of the surviving quality.

Next consider a cap on the maximum quality that is offered. This is equivalent to removing the highest qualities in sequence, and has no effect on the prices of remaining qualities (Proposition 3). A quality cap, then, can only harm buyers.

Finally, consider a minimum quality regulation. As with other product removals, this can only be helpful if the Mills ratio is convex. The prices of all remaining higher qualities fall, and so higher types must benefit; lower types, however, are harmed.

Proposition 4 (Minimum quality standards). *If the Mills ratio is concave, then all consumers are harmed by a minimum quality standard. If the ratio is convex, then there is a critical consumer type $\bar{\theta}$ such that types above $\bar{\theta}$ are helped by such a standard, whereas types below $\bar{\theta}$ are harmed.*

A minimum quality may well be expected to lower suppliers' profits. Given a free-entry condition, this induces exit and so raises prices. Thus, if the Mills ratio is close to linear (as it is for specifications studied in the next section) the net effect of the standard will be to raise all prices. This is (in such a free-entry world) unambiguously bad for welfare.

4. The impact of naive pricing

Pricing neutrality (Proposition 1) does not hold if demand has variable curvature. Here we consider (as robustness checks) the prices and profits of a monopolist for two flexible specifications with non-linear Mills ratios: the beta distribution and the log-normal distribution.⁶

4.1. Beta distributed types

Suppose that types and qualities lie in [0, 1]. Furthermore:

$$c(q) = \frac{q^2}{2} \quad \text{and} \quad f(\theta) \propto \theta^{\gamma-1} (1-\theta)^{(1-\gamma)\delta-1} \quad \text{for } \gamma \in (0, 1) \text{ and } \delta > 0. \quad (18)$$

Hence θ follows a beta distribution with mean $E[\theta] = \gamma$ and variance $var[\theta] = \gamma(1-\gamma)/(\delta+1)$. Three cases are illustrated in Fig. 1. The relevant region is where $M(\theta) < \theta$, so that marginal revenue is positive. The examples exhibit substantial convexity in that region.

Fig. 1C reports the discriminatory and standalone prices for $\gamma = 0.5$, $\delta = 20$, and $n = 100$ equally spaced qualities. The price sequences are different but close. The price difference is greatest for the highest quality (the Corollary to Proposition 3 predicts this). The standalone price p_n^\dagger is 7% adrift of the optimal price p_n^* , and at the mid-point quality it is 2.3% below. Overall, the incorrect use of single-product prices causes a loss of 1.7% of potential (gross) profit. Note that $\delta = 20$ yields modest but

non-negligible dispersion in buyers' types: the 95% confidence interval for θ is approximately [0.32, 68].

Now we explore what happens as consumer heterogeneity varies, via a change in δ . For fixed mean γ , when δ is very small, the distribution of types is approximately uniform. In contrast, when δ becomes very large, buyers clump tightly around the mean.

Fig. 1D compares the profit from optimal pricing to the profit from naive pricing, for the case of $\gamma = \frac{1}{2}$ and $n = 20$ equally spaced qualities. For this figure, $var[\theta]$ is smallest for $\delta = 100$, which yields a 95% confidence interval of approximately [0.42, 0.58].⁷

For small δ , there is no meaningful penalty from the use of naive single-product prices. The reason is that the Mills ratio becomes approximately linear for such δ , because the distribution becomes approximately uniform. In contrast, precisely because there is significant heterogeneity amongst consumers, restricting to only a single product can lead to larger losses. For example, even with $\delta = 10$, the reduction in profits exceeds 10%.

As buyers become homogeneous, the convex Mills ratio means that naiveté is more costly. Interestingly, for large δ , the monopolist is better off if she sells only a single (optimally chosen and priced) product, rather than selling the full range at naive prices.

Selling a range of products priced naively rather than a single product priced optimally presents a trade-off. On the one hand, offering multiple products increases the potential to extract revenue from consumers. On the other hand, the sub-optimal pricing means that consumers are offered substitution possibilities that limit the monopolist's profit. The magnitude of these effects depends on the underlying dispersion of consumers.

For high dispersion (low δ), there is a larger gain from offering a menu of products. The Mills ratio is approximately linear here, and so the naive prices are close to the optimal ones. For low dispersion (high δ) there is less to be gained: a single product extracts nearly all surplus from the market if it is the only product offered. In this situation, offering additional mis-priced options gives consumers unwelcome substitution opportunities. Offering the most-profitable single product (optimally priced) is (heuristically, at least) safer.

4.2. Log-normally distributed types

We now consider a log-normal specification for types, so that $\log \theta \sim N(\mu, \sigma^2)$. This satisfies $\log E[\theta] = \mu + \frac{\sigma^2}{2}$ with a Gini coefficient equal to $2\Phi(\sigma/\sqrt{2}) - 1$ where $\Phi(\cdot)$ is the standard normal distribution function. Equivalently,

$$\sigma = \sqrt{2}\Phi^{-1}\left(\frac{\text{Gini} + 1}{2}\right) \quad \text{and} \quad \mu = \log E[\theta] - \frac{\sigma^2}{2}. \quad (19)$$

We fix the mean valuation $E[\theta]$ (it does not effect the relative performance of optimal and naive prices) and use the Gini coefficient as a scale-free measure of dispersion.⁸ Fig. 2 uses the normalization $E[\theta] = 1$. In Fig. 2A each line reports the Mills ratio $M(\theta)$ across the interval of types $[\theta_L, \theta_H]$ where $F(\theta_L) = 0.02$ and $F(\theta_H) = 0.98$.⁹

For most cases the Mills ratio is approximately linear in the region where $M(\theta) < \theta$. Only when types are rather homogeneous (Gini = 0.15) does a significant non-linearity exist. This is more

⁷ Profit is quasi-convex, or "U-Shaped" in δ , and so is maximized for very homogeneous (high δ) or very heterogeneous (low δ) buyers. This is a prediction from Johnson and Myatt (2006b, Proposition 1).

⁸ The Gini coefficient is half of the expected absolute difference between two random draws (that is, half of the Gini mean difference) expressed as a fraction of the mean: $\text{Gini} = E[|\theta - \theta'|]/2E[\theta]$.

⁹ The highest four Gini coefficients correspond roughly to income inequality in countries such as Denmark (0.25), the United Kingdom (0.35), the United States (0.45), and Brazil (0.55). If the willingness-to-pay for a product is a fixed fraction of income, then this range contains reasonable values for the distribution of types.

⁶ Although we focus on a monopolist here, the findings readily extend to the properties of Cournot prices.

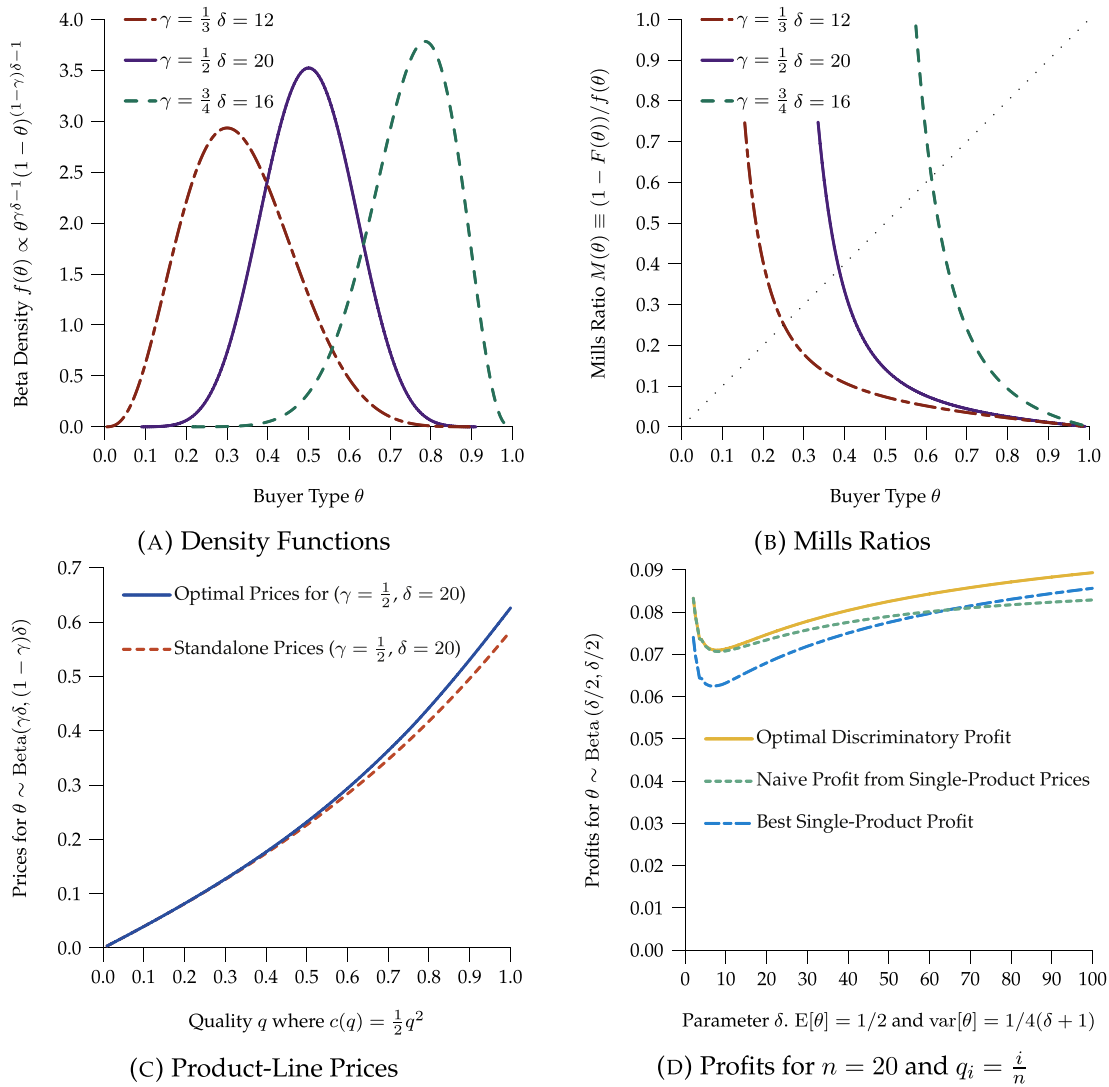


Fig. 1. Illustrations for the beta distribution.

pronounced when θ is small, corresponding to a situation in which most of the market is served. Hence, the divergence between optimal multi-product prices and single-product prices will be small unless very low production costs induce high output.

When consumer valuations are very similar, the gain from offering multiple products is smaller (Fig. 2B). We see a repetition of the lesson

from the beta distribution case. If the Gini coefficient is small, so that potential buyers are similar, then there is a significant difference between the multi-product and single-product prices, but in these cases it is optimal to serve most consumers the same product. This suggests that a monopolist does quite well by selling only a single product, so long as she selects and prices that product correctly.

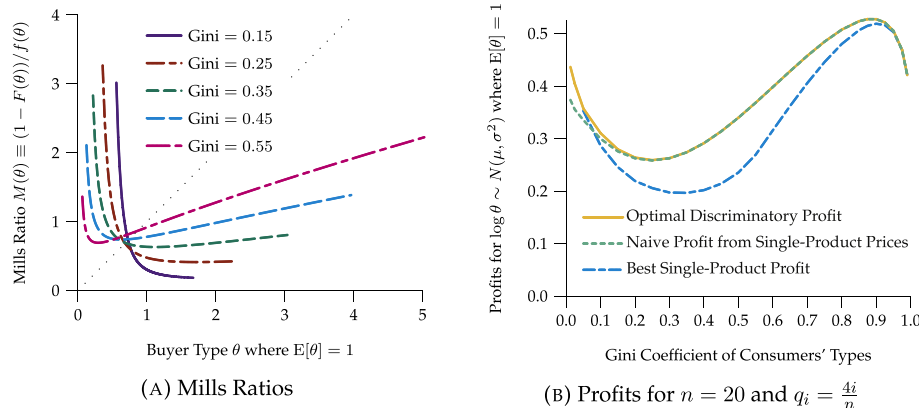


Fig. 2. Illustrations for the log-normal distribution.

A final observation from Fig. 2B is that profit is no longer purely U-shaped in the dispersion of types: for very high inequality, profit drops with further increases in inequality.¹⁰ For very high Gini coefficients most mass shifts toward zero, but the distribution has a very long upper tail. To capture surplus from that tail requires products of higher quality. However, for Fig. 2B we limit to a grid of $n = 20$ qualities on the interval $q_i \in [0, 4]$. The upper bound prevents the monopolist from reacting optimally to high Gini coefficients.

5. Concluding remarks

As teachers of economics, we emphasize the conceptual differences between single-product and multi-product (discriminatory) pricing. But we do not typically explain whether more sophisticated tools are associated with significant changes to prices or final profits.

The early literature established how the removal of an option from a monopolist's product line affects the prices of higher qualities; the shape of the Mills ratio is determinative (Itoh, 1983). From here, it is a simple step to observe that, in situations corresponding to common "textbook" specifications, discriminatory prices must be equal to single-product prices. A role of this note is to highlight this observation; we suggest that it is something that students of economics should know. Furthermore, our results show that this message carries over to a Cournot environment, and it holds in spirit under a wider range of demand specifications. Finally, the circumstances in which the price

equivalence result breaks down seem to be those in which there is relatively little to be gained by offering a multi-product menu.

References

- Aguirre, I., Cowan, S.G.B., Vickers, J., 2010. Monopoly price discrimination and demand curvature. *Am. Econ. Rev.* 100 (4), 1601–1615.
- Anderson, S.P., Çelik, L., 2015. Product line design. *J. Econ. Theory* 157, 517–526.
- Anderson, E.T., Dana, J.D., 2009. When is price discrimination profitable? *Manag. Sci.* 55 (6), 980–989.
- Bulow, J.I., Pfleiderer, P., 1983. A note on the effect of cost changes on prices. *J. Polit. Econ.* 91 (1), 182–185.
- Cowan, S.G.B., 2007. The welfare effects of third-degree price discrimination with nonlinear demand functions. *RAND J. Econ.* 38 (2), 419–428.
- Cowan, S.G.B., 2012. Third-degree price discrimination and consumer surplus. *J. Ind. Econ.* 60 (2), 333–345.
- Itoh, M., 1983. Monopoly, product differentiation and economic welfare. *J. Econ. Theory* 31 (1), 88–104.
- Johnson, J.P., Myatt, D.P., 2003. Multiproduct quality competition: fighting brands and product line pruning. *Am. Econ. Rev.* 93 (3), 748–774.
- Johnson, J.P., Myatt, D.P., 2006a. Multiproduct Cournot oligopoly. *RAND J. Econ.* 37 (3), 583–601.
- Johnson, J.P., Myatt, D.P., 2006b. On the simple economics of advertising, marketing, and product design. *Am. Econ. Rev.* 96 (3), 756–784.
- Maskin, E., Riley, J.G., 1984. Monopoly with incomplete information. *RAND J. Econ.* 15 (2), 171–196.
- Mussa, M., Rosen, S., 1978. Monopoly and product quality. *J. Econ. Theory* 18 (2), 301–317.
- Robinson, J., 1933. *The Economics of Imperfect Competition*. Macmillan, London.
- Weyl, E.G., Fabinger, M., 2013. Pass-through as an economic tool: principles of incidence under imperfect competition. *J. Polit. Econ.* 121 (3), 528–583.