

Brand-Specific Tastes for Quality

Alessandro Bonatti*

November 10, 2010

Abstract

This paper develops a model of nonlinear pricing with competition. The novel element is that each consumer's willingness to pay for quality is private information and is allowed to differ across brands. The consumer's preferences are represented by a multidimensional type containing the marginal value of quality for different products. Buyers with high willingness to pay for quality also display strong preferences for particular brands, and require higher discounts in order to switch away from their favorite product. Therefore, competition is fiercer for buyers with lower tastes for quality, and hence more elastic demands. This is in sharp contrast to earlier models in which competition is fiercer for higher-taste, more valuable buyers. In equilibrium, firms either compete intensively for the entire market (providing strictly positive rents to all consumers) or shut down the least profitable segment of the market. Quality levels are distorted downwards for all buyers, except for those with the highest type. The number of competing firms and the degree of correlation across brand preferences enhance the efficiency of the allocation.

KEYWORDS: Nonlinear pricing, multidimensional types, oligopoly, menus of contracts, product differentiation.

JEL CLASSIFICATION: D43, D86, L13

*MIT Sloan School of Management, 100 Main Street, Cambridge MA 02142. bonatti@mit.edu. This paper is based on Chapter 2 of my doctoral dissertation at Yale University. I am indebted to Dirk Bergemann and Ben Polak for their invaluable help and encouragement throughout this project. This paper benefited from comments by Rossella Argenziano, Mark Armstrong, Luigi Balletta, Dino Gerardi, Michael Grubb, Jean-Charles Rochet, Maher Said, Dick Schmalensee, and participants in the 2007 North American Summer Meeting of the Econometric Society and in the Micro Theory Lunch at Yale University. Finally, I wish to thank Bernard Caillaud and two anonymous referees whose comments helped improve both the analysis and the exposition of the paper.

1 Introduction

The empirical industrial organization literature has successfully used discrete choice models of product differentiation to analyze markets in which consumers have varying tastes for product characteristics. In these models, consumers' choices of brand are largely driven by the different features of the products offered by each firm: buyers with strong tastes for given product characteristics are more likely to purchase high quality items, and are willing to pay a larger premium for their favorite brand.

Theoretical models of competitive quality pricing have also been developed, mainly with the goal of analyzing firms' choices of product characteristics and prices simultaneously.¹ However, the existing models represent brand preferences as independent, additive shocks to the consumer's utility. An implication of this approach is that a fixed discount applied to a firm's entire product line yields an identical increase in the sales of each item. As a consequence, estimating these models may deliver unrealistic predictions about the price elasticity of demand for different products.

In this paper, we propose a screening model in which sellers offer menus of contracts (nonlinear tariffs) to buyers with private information about their preferences for both brand and quality. Brand preferences are modeled by letting each consumer's *marginal* utility of quality depend on the product's brand. Equivalently, we can interpret a consumer's idiosyncratic taste for a brand as the value of the match between her tastes for characteristics and the attributes of that brand's products. This implies that buyers who purchase high quality items also require higher discounts in order to switch away from their favorite brand.

The dependence of brand preferences on the purchased quality level best describes markets for products involving choices on both the intensive and the extensive margin, such as subscription plans. For example, consider cell phone plans: consumers' willingness to pay for a given carrier's plan depend on the desired usage intensity, and on some carrier-specific features, such as network coverage and the bundled telephones. Consumers with a higher usage intensity, who sort into plans with more free minutes, text messaging, or Internet access, assign a higher value to better network coverage. It is then reasonable to assume that consumers are willing to pay per-minute premia for their preferred carrier. As an alternative example, one might consider the market for memberships into clubs. In this case, more intensive users, who are more likely to buy higher-quality, "premium" membership packages, command a per-usage discount in order to switch clubs.²

Our formulation of brand preferences may also be derived directly from the characteristics

¹In particular, Armstrong and Vickers (2001), Rochet and Stole (2002), and Yang and Ye (2008).

²For example, see the analysis of golf clubs' rewards programs by Hartmann and Viard (2006), and that of airlines' frequent flyer programs by Lederman (2007).

approach of Lancaster (1971), in which consumers' willingness to pay for each product is given by a weighted average of several features, rather than by a one-dimensional quality measure. In this framework, consumers' tastes and product characteristics determine the equilibrium market share of each brand. With this interpretation, our model represents a theoretical contribution towards integrating endogenous product characteristics and price discrimination in an imperfectly competitive environment.

As a concrete example, consider the market for Central Processing Units (CPU), in which buyers (*e.g.* computer manufacturers) have needs for performance based on characteristics such as clock speed and cache memory. In addition, each of the two leading brands (Intel and AMD) has a comparative technological advantage at one of these two characteristics. In this context, buyers sort into different brands based on their particular needs: those who assign a higher value to clock speed are willing to pay a premium for Intel's products, and the opposite is true for those consumers with a higher valuation for cache memory. It can be useful to summarize the features of computer into a unitary measure of product quality. In this case, a consumer will naturally exhibit brand-specific tastes for quality. In fact, her willingness to pay for products of identical quality and different brand could not be represented by a fixed dollar amount. On the contrary, since brand preferences originate from differences in product characteristics, this premium should be proportional to the quality of the goods being compared (a per-gigahertz premium, for the case of clock speed).

The results show that the equilibrium menus of contracts share the main qualitative features of the Mussa and Rosen (1978) monopoly allocation, namely efficient quality provision for the highest type and quality distortions for all other types. One notable difference is that firms choose either to compete intensively for the entire market (providing strictly positive rents to all consumers) or to shut down the least profitable segment of the market. The degree of correlation across buyers' brand preferences and the number of competing sellers have similar effects on the equilibrium allocation: they increase market coverage, reduce quality distortions and increase the consumers' information rents.

In sharp contrast with other competitive screening models, our formulation of brand preferences implies that quality level offered to the lowest consumer type is inefficiently low. In consequence, compared to the existing papers, our model has the following implications in terms of observable variables: (i) a wider range of offered quality levels, for a given distribution of consumers' tastes; and (ii) higher marginal prices for each quality level. We come back to the CPU market in section 7, and we use it as an example to illustrate these differences and their implications for empirical work. Compared to our model, previous models would tend to overestimate the fraction of high valuation consumers, as well as the price sensitivity of their demand. Conversely, we find that our model would predict a much

higher demand sensitivity for low valuation consumers.

This paper is mainly related to the literature on competitive nonlinear pricing. Following the classification in Stole (2007), early theoretical contributions with single dimensional models are given by Spulber (1989), Champsaur and Rochet (1989), and Stole (1995). In Ellison (2005), buyers differ both in taste for quality and in marginal utility of income. The more recent contributions of Armstrong and Vickers (2001), Rochet and Stole (2002), and Yang and Ye (2008) use multidimensional models to capture uncertainty both over brand preferences and taste for quality. In particular, the three latter papers assume that buyers' types have both a horizontal component (measuring brand preferences) and a vertical component (measuring taste for quality). We discuss these papers at length in section 7.

Our approach to brand preferences is also related to the theoretical and empirical discrete choice models of product differentiation, such as those in Perloff and Salop (1985), Anderson, De Palma, and Thisse (1992), and Berry, Levinsohn, and Pakes (1995). It is even more closely related to Berry and Pakes (2007), who develop a pure-characteristics demand model which removes the additive shocks and only defines consumers' preferences over a set of product characteristics. Song (2007) adopts a pure-characteristics framework to estimate consumer demand and welfare in the CPU market. Nonlinear pricing with competition is also the object of a growing number of papers in empirical industrial organization. This strand of the literature includes the works of Miravete and Röller (2004) and Miravete (2009) on the cellular telephone industry, the paper by Busse and Rysman (2005) on the yellow pages industry, and the work of McManus (2007) on competing coffee chains.

2 The Model

Consider an imperfectly competitive market and let $\mathcal{I} = \{1, \dots, i, \dots, I\}$ be the set of (identical) sellers. Let there also be a continuum of buyers with unit demands. Buyers differ in their valuation of the quality of the goods produced by each firm. A type for a buyer is a vector of valuations $\theta = (\theta_1, \dots, \theta_I) \in \mathbb{R}^I$. Buyer types θ are distributed on $[\theta_L, \theta_H]^I$ according to a distribution function $F(\theta)$ with a strictly positive and continuously differentiable density $f(\theta)$.

The utility of a type θ buyer, consuming a good of quality q_i , produced by firm i and sold at price p_i is given by

$$v(\theta, q_i, p_i) = \theta_i q_i - p_i. \tag{1}$$

We can view each type component θ_i as the value of the match between consumer θ and the products of firm i . Consequently, a type θ buyer is willing to pay a premium of $(\theta_i - \theta_j) q$ for

an item of quality q produced by firm i rather than by firm j . In other words, consumers are willing to pay brand-specific premia that are *proportional to product quality*. It follows that demand for high quality products is less sensitive to absolute price reductions than demand for low quality items.

It can be useful to contrast (1) with the utility function commonly adopted in the literature on competitive screening. This literature defines a buyer type as $(t, x) = (t, x_1, \dots, x_I) \in \mathbb{R}^{I+1}$. The utility of type (t, x) , when consuming a good of quality q_i produced by firm i and sold at price p_i is given by

$$v(t, x, q_i, p_i) = tq_i - p_i + x_i. \quad (2)$$

Suppose that both firms i and j offer product q at price p . Under demand specification (2), a type (t, x) buyer is willing to pay $x_i - x_j$ more for firm i 's item *independently of the quality level q* . It follows that demand functions for high quality and high price items are more sensitive to equal percentage discounts than the corresponding demand functions for low quality items.

Our utility function formulation (1) can also be obtained from a more general model, in which preferences are defined over a vector of product characteristics, as in Berry and Pakes (2007). In this context, a brand is identified by a combination of characteristics, while quality represents a scaling decision for products with similar combinations of characteristics. The rest of this paper works directly with the reduced form utility function (1), but we will come back to the product characteristics interpretation in section 6.

Each firm $i \in \mathcal{I}$ can produce goods of quality q_i with the same technology $c(q_i)$, which satisfies $c(0) = 0$, as well as $c'(q_i) > 0$ and $c''(q_i) > 0$ for all q_i . If each firm i chooses a nonlinear tariff $p_i(q_i)$, the indirect utility of a type θ consumer when purchasing from firm i is given by

$$U_i(\theta) = \max_{q_i} (\theta_i q_i - p_i(q_i)).$$

Let $U_0(\theta)$ denote the consumer's outside option and assume that $U_0(\theta) = 0$ for all types θ . A type θ consumer chooses to make her purchase from firm i whenever this firm provides her the highest net utility,

$$U_i(\theta) \geq U_j(\theta) \text{ for all } j = 0, 1, \dots, I.$$

Firm i 's *market segment* is defined as the set of types purchasing from i :

$$Z_i = \{\theta \in [\theta_L, \theta_H]^I : U_i(\theta) \geq U_j(\theta) \quad \forall j \in \mathcal{I} \cup \{0\}\}.$$

Denote by $q_i(\theta)$ the quality level (possibly zero) purchased by type θ from firm i . Each firm

i then seeks to maximize profits:

$$\Pi_i = \max_{q_i, p_i(q_i)} \int_{Z_i} [p_i(q_i(\theta)) - c_i(q_i(\theta))] f(\theta) d\theta.$$

3 Menu Pricing

As buyers have unit demands, we analyze competition among I sellers in an exclusive dealing framework. Throughout the paper, we maintain the assumption that firms offer deterministic menus of contracts. Under this assumption, it is without loss of generality to restrict attention to direct mechanisms. Furthermore, in these mechanisms, each firm i only conditions its price and quality offer on the buyer's relevant type component θ_i . The reason for this result lies in the separability of the agent's preferences. Under utility function (1), the agent's ranking of items within each firm's menu is unaffected by the menus offered by other firms. Therefore, following Martimort and Stole (2002), firms cannot benefit from offering out of equilibrium contracts, *i.e.* price-quality pairs that are not chosen by any type in equilibrium. Restricting the message space to the type space then entails no loss of generality. Furthermore, since all types $\theta = (\theta_i, \cdot)$ have the same preferences over the items in firm i 's menu, firm i is unable to screen over all type components $\theta_j \neq \theta_i$. We can therefore restrict attention to menu offers in which the agent only reports type component θ_i to each firm i .

This feature of our model does not require any assumption on the distribution of types (e.g. independence). In fact, if type components were correlated, then each firm i would derive additional information about the buyer's reservation utility from the revelation of θ_i . This information would certainly affect the terms of the offer the firm makes to each type θ_i . However, the firm could not possibly use this information to further screen consumers, since it lacks the instruments to discriminate among types with identical θ_i and different θ_{-i} .

Therefore, each firm's optimal nonlinear pricing problem may be solved with the traditional techniques of one-dimensional screening. In a direct mechanism, each firm's price and quality offer is a function of the buyer's reported type $\hat{\theta}_i$. The utility of a type θ_i buyer who reports $\hat{\theta}_i$ when buying from firm i may be written in the usual form:

$$U_i(\theta_i, \hat{\theta}_i) = \theta_i q_i(\hat{\theta}_i) - p_i(\hat{\theta}_i).$$

Normalize the buyers' reservation utility to zero, and write the individual rationality constraint as

$$U_i(\theta_i, \theta_i) \geq 0. \tag{3}$$

Now define each firm's utility provision schedule as

$$U_i(\theta_i) = \max_{\hat{\theta}_i} U_i(\theta_i, \hat{\theta}_i).$$

The incentive compatibility constraints for the firm's problem are then given by the consumer's first- and second-order conditions for truthful revelation of her type. By standard arguments, these are equivalent to:

$$U_i'(\theta_i) = q_i(\theta_i), \quad (4)$$

$$q_i'(\theta_i) \geq 0. \quad (5)$$

Consider a profile of incentive-compatible menus $\{q_i(\theta_i), U_i(\theta_i)\}_{i=1}^I$. The incentive compatibility constraints (4) and (5) imply the indirect utility function $U_i(\theta_i)$ is strictly increasing. Therefore, a buyer of type θ chooses firm i whenever her taste θ_i for brand i is sufficiently large, relative her other type components $\theta_{j \neq i}$. We can then characterize the market share of firm i through a vector of $I - 1$ threshold types θ_j^* . Fixing a type component θ_i , the threshold θ_j^* represents the lowest taste for brand j that would induce consumer $\theta = (\theta_i, \cdot)$ to prefer firm j over firm i . The threshold types are defined as follows:

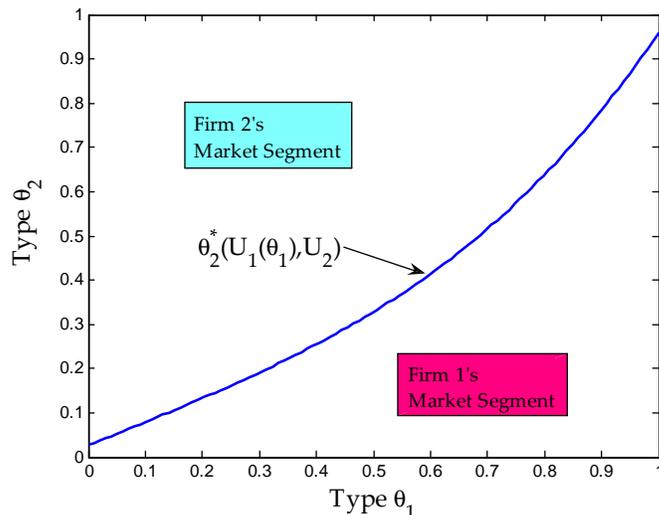
$$\theta_j^*(U_i(\theta_i), U_j) = \begin{cases} \theta_L & \text{if } U_i(\theta_i) - U_j(\theta_j) < 0 \text{ for all } \theta_j, \\ t \text{ s.t. } U_i(\theta_i) = U_j(t) & \text{if } t \in [\theta_L, \theta_H], \\ \theta_H & \text{if } U_i(\theta_i) - U_j(\theta_j) > 0 \text{ for all } \theta_j. \end{cases} \quad (6)$$

Notice that the set of types $\theta = (\theta_i, \cdot)$ that choose to purchase from firm i depends on the utility level $U(\theta_i)$ assigned to type θ_i and on the entire utility schedules $U_j(\theta_j)$ offered by firms $j \neq i$. If we let $f_i(\theta_i)$ denote the marginal density of type component θ_i , we can then use equation (6) to define firm i 's market share function $M_i(U_i(\theta_i), U_{-i}, \theta_i)$ as

$$\begin{aligned} M_i(U_i(\theta_i), U_{-i}, \theta_i) &= \Pr(U_j(\theta_j) \leq U_i(\theta_i) \ \forall j \neq i \mid \theta_i) \cdot f_i(\theta_i) \\ &= \Pr(\theta_j \leq \theta_j^*(U_i(\theta_i), U_j) \ \forall j \neq i \mid \theta_i) \cdot f_i(\theta_i). \end{aligned} \quad (7)$$

Figure 1 provides an illustration of the market segments in the case of a duopoly, when types θ are distributed on the unit square.

Figure 1: Threshold Types and Market Segments



Finally, define the total surplus generated by providing quality q_i to type θ_i as

$$S_i(q_i, \theta_i) = \theta_i q_i(\theta_i) - c_i(q_i(\theta_i)).$$

Given the strategy profile $\{q_j(\theta_j), U_j(\theta_j)\}_{j \neq i}$ of all firms $j \neq i$, each seller i solves the following problem:

$$\begin{aligned} \max_{q_i(\theta_i), U_i(\theta_i)} \int_{\theta_L}^{\theta_H} (S_i(q_i, \theta_i) - U_i(\theta_i)) M_i(U_i(\theta_i), U_{-i}, \theta_i) d\theta_i \\ \text{s.t. (3), (4), and (5).} \end{aligned} \quad (8)$$

The firm's objective function differs substantially from the Mussa and Rosen (1978) monopoly problem. The differences are mainly due to the effect of buyers' brand-specific tastes on the competition among sellers. In particular, each firm's allocation of buyer types is endogenously determined, meaning that firms can acquire larger market shares by offering higher utility levels. For a given strategy profile $\{q_i(\theta_i), U_i(\theta_i)\}_{i=1}^I$, the utility levels offered by firms $j \neq i$ influence the fraction of types (θ_i, \cdot) served in equilibrium by firm i . In other words, market shares measure sales of products of quality $q_i(\theta_i)$ by firm i . The market share function $M_i(U_i(\theta_i), U_{-i}, \theta_i)$ given in (7) may then be viewed as a weighted average of the initial marginal distribution $f_i(\theta_i)$ that places a higher weight on high θ_i types. Clearly, because firms split the market, the weights do not sum to one.

4 Symmetric Equilibrium

We now present the necessary conditions for a symmetric equilibrium, and their implications for the properties of the solution. To simplify the notation, we now drop the subscript i from the firms' strategies, and (with a slight abuse of notation) we let type components θ_i be identically and independently distributed according to a univariate distribution $F(\theta_i)$. We extend the analysis to the case of correlation in section 6.

The necessary conditions for a symmetric equilibrium can be expressed as a second-order differential equation in $U(\theta_i)$. For greater clarity, Proposition 1 presents them in terms of both functions $q(\theta_i)$ and $U(\theta_i)$, which are linked by incentive compatibility constraint (4), or equivalently by condition (10) in what follows.

Proposition 1 (Necessary Conditions)

The necessary first-order conditions for quality and utility provision at a differentiable symmetric equilibrium are given by:

$$c''(q(\theta_i))q'(\theta_i) = 2 + \left(\frac{c(q(\theta_i)) + U(\theta_i)}{q(\theta_i)} - c'(q(\theta_i)) \right) (I - 1) \frac{f(\theta_i)}{F(\theta_i)} + (\theta_i - c'(q(\theta_i))) \frac{f'(\theta_i)}{f(\theta_i)}, \quad (9)$$

where

$$U(\theta_i) = U(\theta_L) + \int_{\theta_L}^{\theta_i} q(t) dt, \quad (10)$$

with the boundary condition

$$c'(q(\theta_H)) = \theta_H. \quad (11)$$

The transversality condition (11) delivers the well-known “no distortion at the top” result. In the Mussa and Rosen (1978) monopolistic price discrimination model, the firm extracts all the rent from the lowest type in the distribution. The equilibrium nonlinear tariff may therefore be found using the boundary conditions (11) and $U(\theta_L) = 0$.

In a competitive environment, this second boundary condition is no longer valid, because information rents have a positive impact on market shares and may therefore increase profits. In particular, as shown in condition (10), the utility of the lowest type $U(\theta_L)$ shifts the entire rent function $U(\theta_i)$. As such, it represents a free endpoint for the firm's problem. This immediately allows us to conclude that the shadow value of the incentive compatibility constraint (4) must be zero at θ_L . In the Rochet and Stole (2002) model, in the absence of bunching, this condition delivers the “no distortion at the bottom” result.³

³This feature is also present in other contexts where the objective function is not everywhere decreasing in information rents. For instance, see the optimal taxation model of Seade (1977).

This is not the case in our model, where the differential equation (9) has a singularity at θ_L . This occurs because firm i 's market share of types $\theta_i = \theta_L$ is equal to zero, since the probability of another type component $\theta_{j \neq i}$ being larger than θ_L is equal to one. As we show in the Appendix, this implies that shadow value of the incentive compatibility constraint (4) is equal to zero independently of the level of q_L . Therefore, efficient quality provision at the bottom of the distribution is not an equilibrium feature in the brand-specific tastes model.

Indeed, as we show in Proposition 2, the quality level served to the lowest type in equilibrium must be distorted downwards from the efficient level. This feature of the equilibrium brings the monopoly and oligopoly results closer together, and represents the key novel implication of the brand-specific tastes model. We stress that this feature of our model extends to the case of correlation across type components, as long as $\Pr(\theta_j > \theta_L \mid \theta_i = \theta_L) = 1$ for all i and j .

Proposition 2 (Lowest Type)

1. *If the market is covered in equilibrium, the utility level of the lowest type is given by*

$$U(\theta_L) = c'(q(\theta_L))q(\theta_L) - c(q(\theta_L)). \tag{12}$$

2. *The quality level provided to the lowest type is distorted downward ($\theta_L > c'(q(\theta_L))$).*

An implication of Proposition 2 is that a symmetric equilibrium cannot involve both positive quality provision and full surplus extraction at the bottom of the distribution. To gain some intuition about this result, observe that the provision of positive quality levels to any type θ_i requires the firms to incur in the shadow cost of the incentive constraints for all higher types. Not providing a type with a strictly positive utility level effectively means making zero sales. Therefore, offering positive quality and zero utility (as in the monopoly case) bears only costs and no benefits to the firm.

Perhaps more importantly, condition (12) in Proposition 2 allows us to rule out efficient quality provision at the bottom ($\theta_L = c'(q_L)$). If that were the case, direct substitution into (12) would immediately imply that firms leave the entire surplus $S(q(\theta_L), \theta_L)$ to the buyer, thereby obtaining zero profit margins on the lowest type.⁴ *Per se*, this is not sufficient to rule out efficiency at the bottom, because each firm i never actually serves any consumer with $\theta_i = \theta_L$. Furthermore, we know from condition (10) that $U(\theta_L)$ is chosen optimally, taking into account its effect on *the entire* rent function $U(\theta_i)$, and that competition is fiercer for the lower types, which are more price-elastic. Therefore, it is not a priori clear that lowering the utility of the lowest type (*i.e.* shifting the rent function down) can increase

⁴My thanks to an anonymous referee for pointing this out.

the firm's profits. However, as we show in the Appendix, efficient quality provision at the bottom would imply that profit margins on types θ_i in a neighborhood of θ_L vanish faster than firm i 's equilibrium market share. Intuitively, the firm can then profitably raise prices (at the expense of market shares) and obtain positive profits in a neighborhood of θ_L .

The positive distortions result brings our model closer to the findings of Yang and Ye (2008). The main difference with this paper is that Yang and Ye (2008) assume vertical types t are uniformly distributed on $[0, 1]$. As a result, a set of lowest types are inevitably shut down, because they provide a cost in terms of incentives, and no surplus. Quality distortions are then a consequence of shut down. In other words, the main novel element of our findings is that quality distortions arise even under full market coverage.

In addition to these economic insights, Proposition 2 allows us to use (12) as boundary condition, together with (11), to solve for the symmetric equilibrium under the hypothesis of full market coverage. In the alternative case, in which all types θ_i below a lower threshold θ_0 are excluded in equilibrium, we can use (11) together with boundary conditions $U(\theta_0) = 0$ and $q(\theta_0) = 0$, in order to solve (9).

In the Appendix, we provide two different algorithms to compute the solution in the cases of full and partial market coverage, respectively.⁵ These algorithms include a simple procedure, based on a result by Seierstad and Sydsaeter (1977), to verify the sufficiency of our system (9)-(11). There are two main difficulties associated with analytically checking these conditions: the first one arises because a closed form solution to the first-order conditions can only be obtained in a few special cases; the second one originates from the equilibrium functional form of market shares, which (when all firms $j \neq i$ adopt the same strategy) is given by $F^{I-1}(\theta_j^*(U_i(\theta_i), U_j))$. This makes it difficult to show the concavity of the objective function in the information rents. The Appendix contains a duopoly example, with uniformly distributed types and quadratic costs, which admits a closed-form solution and for which we verify the second order conditions analytically. Both for this reason, and for tractability, we now specialize the model by introducing these two assumptions.

⁵Algorithms such as these ones are also used by Rochet and Stole (2002) and Yang and Ye (2008).

5 Uniform-Quadratic Model

Throughout this section, we maintain the following assumptions:

$$c(q_i) = \frac{1}{2}q_i^2, \quad (13)$$

$$F(\theta_i) = \frac{\theta_i - \theta_L}{\theta_H - \theta_L} \text{ i.i.d. } \forall i. \quad (14)$$

These restrictions allow us to clearly separate our results with those of Armstrong and Vickers (2001) and Rochet and Stole (2002), and to provide some insights into the comparative statics of the equilibrium nonlinear prices with respect to the number of competing firms.

Assessing the effects of competition requires some attention. First, because our model lacks a general existence and uniqueness result, we need to rely on numerical solutions to ensure that the comparative statics exercise is well-founded. Second, increasing the number of firms enhances competition, but also multiplies the preference space of the consumer and the potential social surplus. This effect is similar to what happens in discrete choice models, where consumers have preferences over product characteristics. In that context, the effect of introducing a new product depends on the degree of similarity with the existing ones. In section 6, we relate the degree of similarity of the underlying product characteristics to the correlation among brand-specific taste parameters. In the case of positive correlation, the effects of enlarging the preference space are then clearly dampened. Indeed, the effects of introducing new brands are most stark in the independent types case. As such, the analysis under the independence assumption provides a benchmark, and a lower bound, on the competitive effects of increasing the number of firms.

The effect of a higher number of competitors need not a priori dominate the stochastic improvement in the distribution of consumers' valuations.⁶ Therefore, in the uniform-quadratic model, we vary the number of products I , and we contrast the case of competition with the multiproduct monopolist's solution. The following results assume (and numerically validate this assumption) that there is a unique solution to the first order conditions, and that the second order conditions are satisfied. In particular, Proposition 3 compares the effect of an increase in the number of brands I in the cases of competition and multiproduct monopoly. In order to facilitate this comparison of the equilibrium menus, we solve the monopolist's problem under the assumptions of anonymous pricing. This corresponds to forcing the monopolist to literally posting I menus of contracts. Formally, this means we restrict the direct mechanisms offered by the monopolist to offer the same set of price-quality options to all

⁶The work of Chen and Riordan (2008) addresses exactly this point in the context of a duopoly model with price competition.

buyers, regardless of their reported types.

Proposition 3 (Comparative Statics)

Assume (9)-(11) admit a unique solution, which is in fact an equilibrium.

1. *In the competitive model, as the number of firms I increases:*
 - (a) *market coverage (weakly) increases;*
 - (b) *utility provision (weakly) increases for all buyer types;*
 - (c) *quality provision (weakly) increases for all buyer types.*

2. *In the monopoly model, as the number of products I increases, the quality and utility levels $q_i(\theta_i)$ and $U_i(\theta_i)$ decrease for all θ_i .*

Consider first the multiproduct monopolist’s problem. Since consumers have unit demand, restricting attention to anonymous pricing schemes reduces the monopolist’s problem to one of one-dimensional screening. Each product i ’s market share of consumers is given by $M_i = \{\theta : \theta_i = \max_{j \in \mathcal{I}}(\theta_j)\}$. In other words, each product i is sold as if the monopolist were facing a population of consumers distributed according to $F^I(\theta_i)$. Therefore the quality schedules are given by

$$q_i^m(\theta_i, I) = \theta_i - \frac{1 - F^I(\theta_i)}{IF^{I-1}(\theta_i)f(\theta_i)}.$$

Proposition 3 shows that the number of products yields an increase in the magnitude of the quality distortions. In terms of indirect utility, buyers may either gain or lose. In particular, each consumer draws a taste parameter for the new product. If this new type component is not high enough, her utility will decrease. Similarly, the effect on overall market coverage is ambiguous, as each individual product sells to fewer buyer types, but new products capture a share of the market as well.

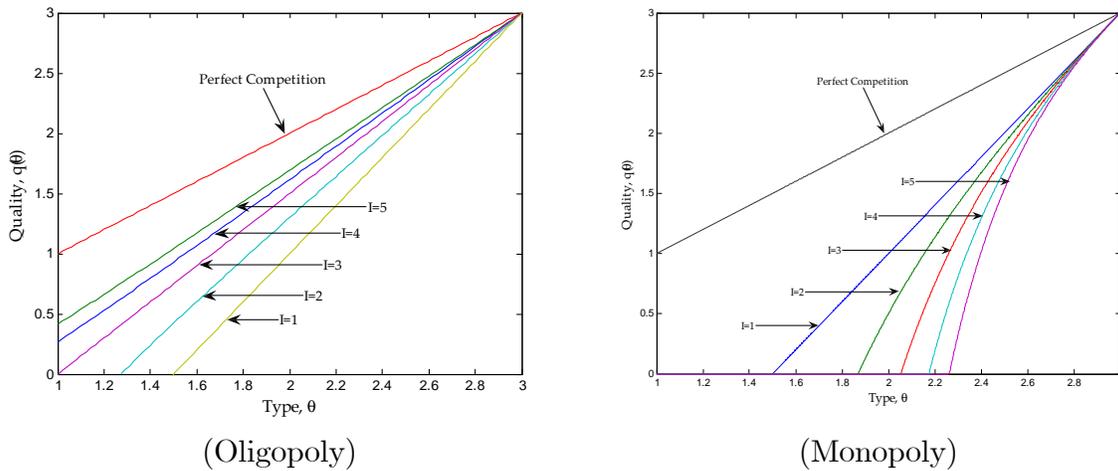
The reasons for the different comparative statics results in the competitive and monopoly models lie in the opportunity cost of providing utility. In the monopoly case, there are no gains in terms of market shares. An increase in the number of products only enlarges the preference space. This creates the possibility for the firm to exploit the gains from additional product variety. In the competitive case, an increase in the number of firms I affects both the size and the composition of each firm’s equilibrium market shares. The size of the overall market shares obviously decreases, reducing both the costs and the benefits of providing extra utility in a similar fashion. However, an increase in competition also increases the proportion consumers with a higher θ_i served by each firm. This means that the provision of utility at the top of the distribution is now more rewarding. As a consequence, the incentive compatibility requirements of distorting quality for low θ_i types are less stringent, resulting in higher utility everywhere. In other words, as I increases, the equilibrium quality levels move towards the

efficient levels and the agents' rents increase. Qualities remain distorted downward at the bottom of the distribution, but the range of quality levels offered in equilibrium decreases.

In Figure 2, we show the equilibrium quality levels when $\theta_i \in [1, 3]$ for the case of competition (left panel) and multiproduct monopoly (right panel).

The forces underlying the comparative statics results in our model differ sharply from those in other models of nonlinear pricing with spatial competition.⁷ These models adopt a linear-city or circular-city framework. As a consequence, firms directly compete only with their “neighbors.” By allowing a more general formulation of brand preferences, our model allows for “global” competition among several firms, as is the case in the applied discrete-choice literature. The extent to which two rival firms are competing is then determined by the specific distribution of consumers' brand preferences.

Figure 2: Equilibrium Quality Supply, Uniform Distribution



Finally, we point out that in the uniform quadratic model, the entire market is served, and the lowest type $\theta_i = \theta_L$ receives exactly zero quality when the following relationship is satisfied:

$$\theta_H = \frac{3 + I}{2} \theta_L.$$

As we show in the Appendix, in this case the system of first-order conditions for a symmetric equilibrium has an analytic solution characterized by linear quality provision:

$$q(\theta_i) = \frac{(I + 3) \theta_i - 2\theta_H}{I + 1}.$$

For example, in Figure 2 (left), the analytic solution is obtained for $I = 3$, $\theta_H = 3$, $\theta_L = 1$. The following corollary can then be derived from Proposition 3.

⁷For example, Spulber (1989), Stole (1995), and Yang and Ye (2008).

Corollary 1 (Market Coverage)

Assume (9)-(11) admit a unique solution, which is in fact an equilibrium. Then the market is fully covered if and only if $I \geq 2(\theta_H/\theta_L) - 3$.

An equivalent interpretation of this result suggests that whenever the ratio of θ_L to θ_H exceeds a critical value the equilibrium involves the provision of strictly positive utility to all types. Conversely, whenever the ratio of θ_L to θ_H falls short of the critical value, $(3/2 + I/2)$, the equilibrium involves the shutdown of the lowest types. For the case of a duopoly, the threshold is equal to $2/5$, which for example is lower than the threshold of $1/2$ obtained in the Mussa and Rosen (1978) monopoly model.

6 Correlated Types and Product Characteristics

Correlation across brand preferences directly affects the intensity of competition by influencing the distribution of buyers' outside options. In particular, positive and negative correlation may be expected to respectively increase and relax competition between firms.⁸

In this section, we derive the necessary conditions for a symmetric equilibrium. These conditions generalize those in Proposition 1. They also serve as a building block for the analysis of the link between product characteristics and brand preferences. We show that brands selling products with similar characteristics generate positively correlated tastes for quality, and provide an illustration of this link through an example with a bivariate normal type distribution.

6.1 Symmetric Equilibrium Conditions

Assume that types θ is distributed over $[\theta_L, \theta_H]^I$ according to a symmetric distribution $F(\theta)$ with density $f(\theta)$. Analogously to the independence case, fix a profile $\{q_i(\theta_i), U_i(\theta_i)\}_{i=1}^I$ of incentive-compatible menus and a seller i . The indifferent types $\theta_j^*(U_i(\theta_i), U_j)$ are defined as in (6). Firm i 's market share of participating types (θ_i, \cdot) may be written as

$$M_i(U_i(\theta_i), U_{-i}, \theta_i) = \int_{\theta_L}^{\theta_j^*(U_i(\theta_i), U_j)} f(\theta) d\theta_{-i}.$$

As noted in section 3, the distribution of types does not affect the buyer's ranking of the items within each firm's menu. The firms' incentive compatibility and individual rationality

⁸The case of perfect negative correlation has been analyzed by Spulber (1989), who finds that firms operate in a local monopolies regime. The case of perfect positive correlation has been analyzed by Champsaur and Rochet (1989), who introduce vertical differentiation by letting firms choose quality ranges first, and then compete in nonlinear prices.

constraints are therefore unchanged and firm i 's objective function may be formulated as follows:

$$\max_{q_i(\theta_i), U_i(\theta_i)} \int_{\theta_L}^{\theta_H} \int_{\theta_L}^{\theta_j^*(U_i(\theta_i), U_j)} (S_i(q_i, \theta_i) - U_i(\theta_i)) f(\theta_i, \theta_{-i}) d\theta_{-i} d\theta_i.$$

The first order conditions for a symmetric equilibrium are then stated in the following proposition. Let $\lambda(\theta_i)$ denote the multiplier on the incentive compatibility constraint (4) for each firm i .

Proposition 4 (Necessary Conditions with Correlation)

The necessary first-order conditions for quality and utility provision at a differentiable symmetric equilibrium are given by:

$$(\theta_i - c'(q(\theta_i))) \int_{\theta_L}^{\theta_i} f(\theta) d\theta_{-i} + \lambda(\theta_i) = 0, \tag{15}$$

$$\int_{\theta_L}^{\theta_i} f(\theta) d\theta_{-i} - \frac{S_i(q_i, \theta_i) - U(\theta_i)}{q(\theta_i)} \sum_{j \neq i} \left(\int_{\theta_L}^{\theta_i} \dots \int_{\theta_L}^{\theta_i} f(\theta) d\theta_{-i-j} \right) = \lambda'(\theta_i), \tag{16}$$

$$\theta_H - c'(q(\theta_H)) = 0. \tag{17}$$

In the Appendix, we provide a duopoly example based on the Farlie-Gumbel-Morgenstern copula (see Nelsen (2006) for more details), in which the degree of correlation is small. We show that the qualitative properties of the equilibrium under independence carry over to the case of correlation. One drawback of using the FMG copula is that it only allows for low values of correlation. We consider higher degrees of correlation in the next example. Before we do so, we demonstrate how brand-specific tastes for quality may be derived from a pure characteristics demand model. We then apply the equilibrium conditions from this subsection to this new “micro-founded” model.

6.2 Product Characteristics

Consider a hedonic oligopoly model in which products are defined as bundles of characteristics $x \in \mathbb{R}^K$ with $K \geq 2$. Let β_k be the consumer's marginal utility of consuming characteristic k . The utility of a buyer consuming product $x = (x_1, \dots, x_K)$ is then given by:

$$u(\beta, x) = \sum_{k=1}^K \beta_k x_k - p(x).$$

This specification is virtually identical to the one adopted in the pure characteristics demand model of Berry and Pakes (2007). In contrast with the earlier discrete-choice em-

pirical literature, which takes all product characteristics as given, we assume that each firm can produce several different quality levels of its variant of the product. We define firm i 's variant of the product as a bundle of characteristics $\alpha_i = (\alpha_{i1}, \dots, \alpha_{iK})$. In other words, firm i 's production possibility set is given by the α_i ray in \mathbb{R}^K .

$$Y_i = \{x \in \mathbb{R}^K : x = (q\alpha_{i1}, \dots, q\alpha_{iK}), q \in \mathbb{R}_+\}.$$

Production costs are expressed in terms of quality levels q only. This formulation allows one to identify a brand with an exogenously determined variant of the product, and quality with an endogenously chosen scaling parameter that distinguishes products of the same brand. The utility of a consumer with preferences β is then given by

$$u(\beta, \alpha_i, q) = q \sum_{k=1}^K \beta_k \alpha_{ik} - p(q).$$

We can now define

$$\theta_i(\beta) = \sum_{k=1}^K \alpha_{ik} \beta_k$$

and obtain the original formulation of brand-specific tastes for quality (1) from preferences over product characteristics. The degree of heterogeneity between firms' variants α_i is an index of product differentiation and, as such, measures the intensity of competition in the market. When adopting the original formulation (1), given an initial distribution of buyers' tastes for characteristics, the degree of similarity between firms' variants is directly related to the degree of correlation between brand-specific taste parameters in the population. In order to obtain a tractable model with a symmetric equilibrium, we only need to ensure that all type components θ_i are distributed over the same support. This can be achieved by assuming that $\sum_k \alpha_{ik} = a$ for all i and that the distribution of each component β_k has the same support. To summarize, the brand-specific model represents a concise way of summarizing tastes for characteristics into a single dimensional brand-specific index. We now illustrate an example based on the multivariate normal distribution.

6.3 The Multivariate Normal Case

Let tastes for each characteristic k be identically and independently distributed according to the following normal density:

$$\beta_k \sim N(\mu, \sigma^2) \text{ i.i.d. for } k = 1, 2.$$

Define firm i 's variant of the product by $\alpha_i = (\alpha_{i1}, \alpha_{i2})$. The two firms' production possibility sets are then given by:

$$X_i = \{x \in \mathbb{R}^2 : x = (q\alpha_{i1}, q\alpha_{i2}), q \in \mathbb{R}_+\}, \forall i = 1, 2.$$

Brand-specific tastes can therefore be defined as:

$$\theta_i(\beta) = \alpha_{i1}\beta_1 + \alpha_{i2}\beta_2.$$

Under the normal distribution assumption for β_k , the brand specific tastes (θ_1, θ_2) are jointly normally distributed with mean vector

$$m = ((\alpha_{11} + \alpha_{12})\mu, (\alpha_{21} + \alpha_{22})\mu),$$

and dispersion matrix

$$\Sigma = \begin{bmatrix} v_1 & \rho\sqrt{v_1v_2} \\ \rho\sqrt{v_1v_2} & v_2 \end{bmatrix},$$

where $v_i = \sigma^2((\alpha_{i1})^2 + (\alpha_{i2})^2)$ denotes the variance of each type component. The correlation coefficient is given by $\rho = \frac{\alpha_1 \cdot \alpha_2}{\|\alpha_1\| \|\alpha_2\|} = \cos(t)$, where t denotes the angle between the two characteristics vectors $(\alpha_{11}, \alpha_{12})$ and $(\alpha_{21}, \alpha_{22})$. Therefore, in this model, the degree of similarity between the two firms' variants of the product determines the correlation coefficient between the consumers' tastes, and hence the intensity of competition on the market. For example, if firms characteristics are collinear, the brand-specific tastes are perfectly correlated, and Bertrand equilibria are obtained. If characteristics are orthogonal, types are independent, and the analysis is identical to that of section 4.

To apply these findings, consider an example in which types θ_i follow a symmetric bivariate normal distribution. In particular, we make the following assumptions on consumers' preferences (β) and firms' variants (α_i) :

$$\begin{aligned} \beta_k &\sim N(0, 1) \quad \forall k = 1, 2 \\ \sum_{k=1}^2 (\alpha_{ik})^2 &= 1 \quad \forall i = 1, 2. \end{aligned}$$

In other words, we let each taste parameter β_k be distributed according to a standard normal, and we assume that variants of the product lie on the unit radius circle. Under these assumptions, the type vector θ is distributed according to a bivariate normal with

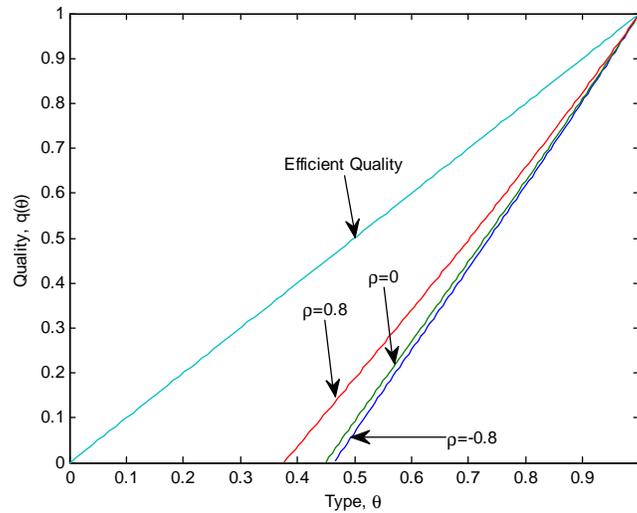
mean vector $m = (0, 0)$ and dispersion matrix $\Sigma = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$, where $\rho = \frac{\alpha_1 \cdot \alpha_2}{\|\alpha_1\| \|\alpha_2\|}$.

We now use the symmetric equilibrium conditions derived above to characterize the solution of this model for different values of ρ . In order to compute a numerical solution, we need to truncate the distribution of each type component at a common upper bound (θ_H). It may be easily shown that the truncation point only affects the terminal conditions and not the differential equations governing the solution. The probability density function is then given by

$$f(\theta_1, \theta_2) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{\theta_1^2 - 2\rho\theta_1\theta_2 + \theta_2^2}{2(1-\rho^2)}\right),$$

which means that the market share function may be expressed in closed form via the error function $\text{erf}(\theta_i)$. In Figure 3, we let costs be quadratic, and let the correlation parameter ρ take values in the set $\{-0.8, 0, 0.8\}$.

Figure 3: Duopoly Quality Provision - Bivariate Normal



The numerical results confirm the intuition that positive correlation across type components increases the efficiency of the quality supply schedules. Quite surprisingly, even values of the correlation coefficient as high as 0.8 do not dramatically improve the efficiency of the allocation. In the opposite case of negative correlation, low types' outside options are higher, and for some parameter values the equilibrium quality provision lies below the monopoly level.

7 Comparison with the Literature

The papers most closely related to the present work are those by Armstrong and Vickers (2001), Rochet and Stole (2002), and Yang and Ye (2008). As alluded to previously, in these papers buyers value quality uniformly, and their brand preferences are given by seller-specific additive utility shocks (see utility function formulation (2)). These shocks are assumed to be uncorrelated with the quality of the purchased product, and distributed independently from the tastes for quality in the population. Thus, these models separate the consumer’s “vertical” preferences over verifiable product quality from her “horizontal” brand preferences. However, as a result, the relative value of purchasing products of similar quality from different firms is independent of the quality of the chosen products. We now illustrate the different theoretical results of these papers and our model. We then turn to their empirical predictions and discuss the implications of taking each model to the data, in the context of the CPU market example.

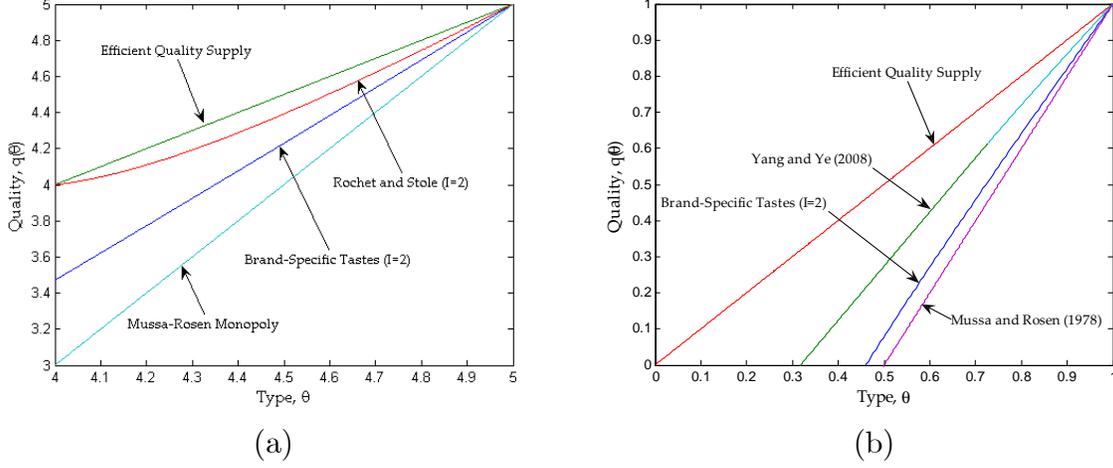
7.1 Model Predictions

The main apparent difference between our results and those of Rochet and Stole (2002) lies in the equilibrium quality schedules. When the market is covered, Armstrong and Vickers (2001) and Rochet and Stole (2002) predict the efficient quality levels are produced in equilibrium. When not all types (t_L, \cdot) participate, the Rochet and Stole (2002) allocation is characterized by efficiency at the top and at bottom. This is in contrast with our result in Proposition 2, which predicts quality distortions for all types, including θ_L . Figure 4, panel (a), considers an example where types θ_i (and vertical types t) are distributed uniformly on $[4, 5]$. It shows that our model predicts more severe quality distortions for all types.

This difference is due to the equilibrium composition of market shares, and to their sensitivity to prices in the two models. To clarify this point, remember that in the Armstrong and Vickers (2001) and Rochet and Stole (2002) models, at a symmetric equilibrium with full market coverage, each firm serves a constant fraction of all vertical types t . In equilibrium, each firm effectively serves a population of consumers whose types are distributed according to $(1/I) f(t)$. Conversely, in the symmetric equilibrium of our model, types in the market segment of firm i are distributed according to $F^{I-1}(\theta_i) f(\theta_i)$, which indicates a relatively larger presence of high types.

To obtain a meaningful comparison between the different models, we appropriately modify the initial distributions of types. In particular, consider the model in Yang and Ye (2008), which extends the one in Rochet and Stole (2002) to the case of imperfect market coverage. Vertical types t in the model of Yang and Ye (2008) are uniformly distributed on $[0, 1]$. For

Figure 4: Equilibrium Quality Provision - Comparison



all types t , market shares are defined by $x^*(U_i(t), U_j(t))$, which is the solution to:

$$U_i(t) - x^* = U_j(t) - (1 - x^*).$$

At a symmetric equilibrium, over the full coverage region, market shares are equal to one-half. Furthermore, the sensitivity of market shares to utility provision is constant and equal to

$$\frac{\partial x^*(U_i(t), U_j(t))}{\partial U_i(t)} = \frac{1}{2} \quad \forall t.$$

In our model, let type components θ_i be distributed according to $F(\theta_i) = \sqrt{\theta_i}$ on $[0, 1]$. As in the model of Yang and Ye (2008), the equilibrium market share of each firm is given by $F(\theta_i) f(\theta_i) = 1/2$. However, in the present work, market shares are determined by the solution to $U_j(\theta_j^*) = U_i$. Therefore, their sensitivity to utility provision is given by

$$\frac{\partial \theta_j^*(U_i(\theta_i), U_j)}{\partial U_i} = \frac{1}{U_j'(\theta_j^*)} = \frac{1}{q_j(U_i(\theta_i), U_j)}.$$

Since quality is an increasing function of type, utility provision has a positive but declining effect on market shares. In other words, the value to firm i of providing additional utility to high-valuation types is much lower than in Yang and Ye (2008), because those types require larger discounts in order to switch brands. Figure 4, panel (b), compares our equilibrium allocation with the model in Yang and Ye (2008), holding constant the composition of the equilibrium market shares. Distortions are greater in the present work, despite the stochastically dominated initial distribution of types. We can then conclude that differences in the

sensitivity of market shares drive the main qualitative differences between the two classes of models.

7.2 Implications for Observable Variables

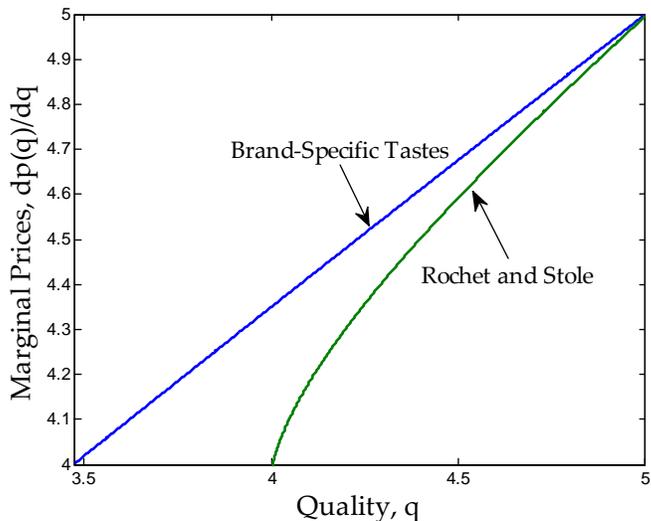
The most immediate implication of quality distortions in terms of observable variables concerns the variety of products offered by each firm. In particular, when the market is covered and there is no bunching, our model predicts a wider range of product qualities than the previous models. Quite simply, the upper bound on quality is identical ($\theta_H = c'(q_H)$), while the lowest offered quality is below the efficient level for the lowest type ($\theta_L > c'(q_L)$).

A more subtle implication of quality distortions is related to the actual nonlinear prices $p_i(q_i)$. Under our linear utility specification, the *marginal* price charged for each quality level is equal to the type of the consumer buying that product. This is most clearly seen from the first order condition for type θ 's choice from firm i 's menu,

$$\theta_i = p'_i(q_i). \tag{18}$$

More severe quality distortions imply that a given quality level is assigned to a higher-valuation buyer, and hence it is sold at a higher marginal price. To illustrate the difference with the Rochet and Stole (2002) model, Figure 5 shows the equilibrium marginal prices, when types θ_i (and vertical types t) are distributed uniformly on $[4, 5]$.

Figure 5: Equilibrium Marginal Prices - Comparison



7.3 Inference from Data

In the Rochet and Stole (2002) and in our model, the econometrician can use price and quality data, together with the first-order condition for the consumer’s problem (18) to make inference about the valuations of the consumers buying each product. For example, if the utility function is assumed to be known, the highest and lowest marginal prices identify θ_H and θ_L .

In addition, if data on production costs is available, the presence of quality distortions at the top and at the bottom of the menu can be tested directly. This is the case, for example, of McManus (2007), who studies specialty coffee chains, and finds no distortions at the top of the menu (large, sweet espresso drinks), and positive distortions in the size of all other coffee products. He interprets the diminishing distortions towards the bottom as evidence of more intense competition for these items. The results of McManus (2007) are in line with the findings of Proposition 2, as well as with the our model’s feature of more price-sensitive demand for buyers of low-quality items.

In the absence of information on cost, but with knowledge of each product’s market share, the existing models and the present work can yield very different estimates of the distribution of valuations and of demand elasticity. We illustrate these differences in the context of our motivating example of the CPU market. In this market, the relevant product characteristics are clock speed and cache memory. Table 1 summarizes these features for the main Intel and AMD processors for sale in 2007. Consistent with our approach in section 6.2, AMD processors generally run faster than their Intel counterparts, but have a smaller cache memory.

Table 1: CPU Comparison - 2007

Intel				AMD			
Model	Clock speed	L2 cache	Price	Model	Clock speed	L2 cache	Price
Core 2 Duo E6300	1.83GHz	2MB	\$183	Athlon 64 X2 4400+	2.3GHz	1MB	\$170
Core 2 Duo E6400	2.13GHz	2MB	\$224	Athlon 64 X2 5000+	2.6GHz	1MB	\$222
Core 2 Duo E6600	2.4GHz	4MB	\$316	Athlon 64 X2 5600+	2.8GHz	2MB	\$326
Core 2 Duo E6700	2.66GHz	4MB	\$530	Athlon 64 X2 6000+	3.0GHz	2MB	\$459
Core 2 Extreme X6800	2.93GHz	4MB	\$999	Athlon 64 FX-70	2.6GHz	4MB	\$599
Core 2 Quad Q6600	2.4GHz	8MB	\$851	Athlon 64 FX-72	2.8GHz	4MB	\$799
Core 2 Extreme QX6700	2.66GHz	8MB	\$999	Athlon 64 FX-74	3.0GHz	4MB	\$999

Source: <http://techreport.com/reviews/2007q1/cpus>

A competitive nonlinear pricing model may be employed to analyze the two firms' choices of product characteristics and prices simultaneously. In particular, the present model would construct a separate quality measure for each brand, by considering different linear combinations of clock speed and cache memory. Conversely, the existing models are based on a one dimensional vertical type. Thus, they would define quality as a linear combination of the two characteristics which is *common to both brands*.⁹

In both models, the observed market shares provide information on the underlying distribution of consumers' valuations for each product. The main difference in this respect is that the Rochet and Stole (2002) model would attribute each purchase of a high quality product to a "high t " type, and explain the choice of brand entirely through consumer-firm fixed effects (equivalently, the logit errors). Conversely, our model would recognize that consumers sort based on their idiosyncratic needs, and hence that the observed market shares correspond to a selected sample of consumers.¹⁰ In other words, some high-quality purchases should be attributed to consumers with a particularly high taste for one brand, but not for the other. Therefore, our model would yield much lower estimates of the number of consumers with uniformly high valuations. This is a general feature of empirical applications of discrete choice models, but in particular of Song (2007), who considers a one-dimensional pure characteristics approach, and of Hendel (1999), who matches buyer characteristics with product characteristics in a multinomial logit model.

Finally, in order to explain the observed marginal prices, the Rochet and Stole (2002) model would estimate a single transportation cost parameter for all vertical types t . Instead in the present model, the degree of horizontal differentiation is implicit in the distribution of brand preferences. As shown in section 6.2, this distribution summarizes the buyers' tastes for characteristics and the properties of each brand's products. Compared to the Rochet and Stole (2002) results, the implied sensitivity of demand will be much lower at the top of the type distribution. For example, if a consumer of the top Intel product (QX6700) were to switch to AMD, she would most likely choose the FX-72 or FX-74, which have considerably different characteristics (most notably, half as much cache memory). Even though these products might have very similar quality levels q_{INTEL} and q_{AMD} , this buyer requires a very large discount in order to consider switching brands. Conversely, market share sensitivity is higher towards the bottom, where absolute differences in product characteristics are smaller.

⁹This might already be a problem, as there exists no linear combination of clock speed and L2 cache that yields increasing price functions for both brands in Table 1.

¹⁰Loosely speaking, many different applications benefit from a large cache memory, while a single, simpler application benefits from a higher clock speed.

8 Conclusion

This paper develops a competitive nonlinear pricing model, where the buyers' valuation of quality depends on the product's brand. In particular, buyers of high quality items are willing to pay larger brand premia, leading to a lower price elasticity for high quality items. In the symmetric equilibrium, brand-specific tastes for quality restore the quality distortions that may be absent from earlier random participation models, and rule out the possibility of marginal cost pricing.

However, our model does not nest those of Armstrong and Vickers (2001), Rochet and Stole (2002) and Yang and Ye (2008), and is to be considered a complement rather than substitute to their approach. Another contribution of this paper is to build a tighter connection between the empirical and theoretical literature on differentiated products oligopoly. In particular, it relates the pure characteristics demand model of Berry and Pakes (2007) to a competitive nonlinear pricing problem. Under particular functional forms, the resulting screening model with correlated types may be solved for a symmetric equilibrium. This allows us to trace out the implications of consumers' preferences over multidimensional product characteristics, when firms simultaneously choose prices and qualities.

A natural extension of our model considers a dynamic game in which product characteristics are determined in a first stage and competitive price discrimination (e.g. through two-part tariffs) takes place in the second stage. Another, more ambitious extension consists of integrating proportional and additive brand preference components in the consumer's choice problem. At present, these extensions constitute the object of further research, but both represent important steps towards adopting the present model for empirical work.

Appendix

Proof of Proposition 1: Suppose all firms $j \neq i$ are offering an identical menu of contracts $(q_j(\theta_j), U_j(\theta_j))$. The Hamiltonian for firm i may be written as

$$H(q_i, U_i, \theta_i, \lambda_i) = (\theta_i q_i - c(q_i) - U_i) F^{I-1}(\theta_j^*(U_i(\theta_i), U_j)) f(\theta_i) + \lambda_i q_i. \quad (19)$$

The first order conditions for (19) with respect to quality and utility provision are given by:

$$(\theta_i - c'(q_i)) F^{I-1}(\theta_j^*(U_i(\theta_i), U_j)) f(\theta_i) + \lambda_i = 0 \quad (20)$$

$$-\frac{\theta_i q_i - c(q_i) - U_i}{q_j(U_i(\theta_i), U_j)} (I-1) F^{I-2}(\theta_j^*(U_i(\theta_i), U_j)) f^2(\theta_i) \quad (21)$$

$$+ F^{I-1}(\theta_j^*(U_i(\theta_i), U_j)) f(\theta_i) = \lambda_i'(\theta_i) \quad (22)$$

$$\lambda_i(\theta_H) = 0.$$

At a symmetric equilibrium, $q_j(\theta_j) = q_i(\theta_i)$ and $\theta_j^* = \theta_i$. Differentiating (20) with respect to θ_i , solving for $\lambda_i'(\theta_i)$, and substituting into (21) delivers the expressions for first order conditions (9)-(11). \square

Proof of Proposition 2: (1.) Consider the necessary conditions (9)-(11) and drop the argument (θ_i) . Rearranging condition (9) one obtains:

$$U(\theta_i) = c'(q(\theta_i)) q(\theta_i) - c(q(\theta_i)) - \frac{2 - c''(q(\theta_i)) q'(\theta_i) F(\theta_i)}{I-1} \frac{F(\theta_i)}{f(\theta_i)} q(\theta_i) \quad (23)$$

$$- \frac{\theta_i - c'(q(\theta_i)) F(\theta_i) f'(\theta_i)}{I-1} \frac{F(\theta_i) f'(\theta_i)}{f^2(\theta_i)} q(\theta_i),$$

which holds for all $\theta_i > \theta_L$. Since both F and f are assumed to be continuously differentiable, and $f > 0$ for all θ_i , by the continuity of U , we have

$$\lim_{\theta_i \rightarrow \theta_L} \frac{F(\theta_i)}{f(\theta_i)} = \lim_{\theta_i \rightarrow \theta_L} \frac{F(\theta_i) f'(\theta_i)}{f^2(\theta_i)} = 0.$$

Therefore, taking the limit of the right hand side of expression (23) as $\theta_i \rightarrow \theta_L$, we can conclude that:

$$U(\theta_L) = c'(q(\theta_L)) q(\theta_L) - c(q(\theta_L)),$$

which is strictly positive for all $q > 0$ since $c(q)$ is strictly convex.

(2.) No firm offers qualities in excess of the socially efficient level, as it could offer the same utility levels at a lower cost by reducing quality. Therefore, we seek to rule out the following allocation:

$$\begin{aligned} c'(q(\theta_i)) &< \theta_i \quad \forall \theta_i \in (\theta_L, \theta_H), \\ c'(q(\theta_L)) &= \theta_L, \\ c'(q(\theta_H)) &= \theta_H. \end{aligned}$$

Define the profit margin on type θ_i as

$$\pi(\theta_i) \triangleq \theta_i q(\theta_i) - c(q(\theta_i)) - U(\theta_i).$$

Suppose towards a contradiction that $c'(q(\theta_L)) = \theta_L$. Condition (12) immediately implies $\pi(\theta_L) = 0$. Differentiating and using the incentive compatibility constraint (4), we obtain

$$\pi'(\theta_L) = (\theta_L - c'(q(\theta_L)))q'(\theta_L) + q(\theta_L) - U'(\theta_L) = 0. \quad (24)$$

At a symmetric equilibrium, first order condition (21) may be written as

$$F^{I-1}(\theta_i) f(\theta_i) - \frac{\pi(\theta_i)}{q(\theta_i)} (I-1) F^{I-2}(\theta_i) f^2(\theta_i) = \lambda'(\theta_i). \quad (25)$$

Both $q(\theta_i)$ and $f(\theta_i)$ are strictly positive. Therefore, as $\theta_i \rightarrow \theta_L$, condition (24) implies that the second term in (25) goes to zero at rate $(d\theta_i)^I$. Conversely, the first term goes to zero at rate $(d\theta_i)^{I-1}$. Consequently, there exists an $\varepsilon > 0$ such that

$$\int_{\theta_L}^{\theta_L + \varepsilon} \lambda'(\theta_i) d\theta_i = \lambda(\theta_L + \varepsilon) - \lambda(\theta_L) = \lambda(\theta_L + \varepsilon) > 0,$$

since first order condition (20) implies $\lambda(\theta_L) = 0$. But then first order condition (20) also implies the firm is offering quality $c'(q(\theta_L + \varepsilon)) > \theta_L + \varepsilon$, which is in excess of the efficient level and clearly sub-optimal. \square

We now use restriction (12) from Proposition 2 to implement two algorithms to compute the symmetric equilibrium.

Computing the solution under full market coverage: First, choose an initial value q_0 for $q(\theta_L)$ on $[\theta_L, \theta_H]$. Then solve the system using $c'(q(\theta_H)) = \theta_H$ and $q(\theta_L) = q_0$ as boundary conditions, and verify whether condition (12) holds. If it does, then the equilibrium has the desired properties; otherwise, adjust the initial value for q_0 and go back to the first

step. This procedure must be repeated until the right-hand side of (12) converges to the corresponding $U(\theta_L; q_0)$. The limit q_0 , along with the associated schedules $q(\theta_i)$ and $U(\theta_i)$, result in the computed equilibrium. To verify that the second order conditions are satisfied, use first order condition (20) to compute the equilibrium value of $\lambda(\theta_i)$. Then check that the computed $q(\theta_i)$ and $U(\theta_i)$ maximize the modified Lagrangean of Seierstad and Sydsaeter (1977):

$$\mathcal{L}(q_i, U_i, \theta_i) = (\theta_i q_i - c(q_i) - U_i) F^{I-1}(\theta_j^*(U_i(\theta_i), U_j)) f(\theta_i) + \lambda(\theta_i) q_i + \lambda'(\theta_i) U_i. \quad (26)$$

Computing the solution under partial market coverage: First, choose an initial value for θ_0 on $[\theta_L, \theta_H]$. Then solve the system using $c'(q(\theta_H)) = \theta_H$ and $U(\theta_0) = 0$ as boundary conditions and verify whether the condition $q(\theta_0) = 0$ holds or not. If it does, then the equilibrium has the desired properties; otherwise, adjust the initial value for θ_0 and go back to the first step. This procedure must be repeated until $q(\theta_0)$ converges to zero. The limit θ_0 , along with the associated schedules $q(\theta_i)$ and $U(\theta_i)$, result in the computed equilibrium. To verify that the second order conditions are satisfied, use first order condition (20) to compute the equilibrium value of $\lambda(\theta_i)$. Then check that the computed $q(\theta_i)$ and $U(\theta_i)$ maximize the modified Lagrangean given in (26).

The following lemmas are instrumental to the proof of Proposition 3.

Lemma 1 *Under the quadratic costs assumption (13), if $q'(\theta_i) \leq (>) 1$ for all $\theta'_i \leq \theta_i$, then $q^2(\theta_i)/2 \leq (>) U(\theta_i)$.*

Proof of Lemma 1: We know that $U'(\theta_i) = q(\theta_i)$ by incentive compatibility and that $U(\theta_L) = q^2(\theta_L)/2$ from Proposition 2. Now rewrite $U(\theta_i)$ and $q(\theta_i)$ as

$$\begin{aligned} U(\theta_i) &= \int_{\theta_L - q(\theta_L)}^{\theta_L} [x - (\theta_L - q(\theta_L))] dx + \int_{\theta_L}^{\theta_i} q(x) dx \text{ and} \\ q(\theta_i) &= \int_{\theta_i - q(\theta_i)}^{\theta_i} [x - (\theta_i - q(\theta_i))] dx. \end{aligned}$$

Since $q'(\theta_i) \leq 1$ for all $\theta'_i \leq \theta_i$, the quantity $x - (\theta_i - q(\theta_i))$ lies below both $q(x)$ and $x - (\theta_L - q(\theta_L))$ for all $x \in [\theta_i - q(\theta_i), \theta_i]$. Therefore, $U(\theta_i) \geq q^2(\theta_i)/2$. The converse holds when $q'(\theta_i) > 1$. \square

Lemma 2 Under the uniform distribution and the quadratic cost assumptions (13)-(14), $U(\theta_i) \leq q^2(\theta_i)/2$ for all θ_i .

Proof of Lemma 2: We know from Lemma 1 that if there is a type θ_i such that $U(\theta_i) > q^2(\theta_i)/2$, then there must also exist a type $\theta'_i \leq \theta_i$ for which $q'(\theta'_i) < 1$. Furthermore, notice that notice that, under assumptions (13)-(14), equation (23) may be re-written as follows:

$$q'(\theta_i) = 2 + \left(\frac{U(\theta_i)}{q(\theta_i)} - \frac{q(\theta_i)}{2} \right) \frac{I-1}{\theta_i - \theta_L}. \quad (27)$$

In particular, this implies the following:

$$q'(\theta_i) \leq 1 \Leftrightarrow \frac{q(\theta_i)}{2} - \frac{U(\theta_i)}{q(\theta_i)} \geq \frac{\theta_i - \theta_L}{I-1}. \quad (28)$$

Two cases are possible.

(a) If $U(\theta''_i) > q^2(\theta''_i)/2$, then direct substitution into equation (27) and (28) yields a contradiction.

(b) If $U(\theta''_i) \leq q^2(\theta''_i)/2$, then consider the difference between the two functions $U(\theta_i)$ and $q^2(\theta_i)/2$:

$$\frac{d}{d\theta_i} \left(\frac{1}{2} q^2(\theta_i) - U(\theta_i) \right) = q(\theta_i) (q'(\theta_i) - 1). \quad (29)$$

The two conditions $q^2(\theta'_i)/2 \geq U(\theta'_i)$ and $q^2(\theta_i)/2 < U(\theta_i)$, together with (29) imply that there must exist a type $\theta''_i \in (\theta'_i, \theta_i)$ for which $q'(\theta''_i) < 1$ and $U(\theta''_i) > q^2(\theta''_i)/2$. One may then repeat the steps from part (a) and derive a contradiction. \square

Proof of Proposition 3: (1.) Assume that types are uniformly distributed and let the number of firms I increase. We want to show that (a) market coverage, (b) every type's utility, and (c) quality provision all (weakly) increase for all θ_i . Denote by $(q_I(\theta_i), U_I(\theta_i), \lambda_I(\theta_i))$ and $(q_{I+1}(\theta_i), U_{I+1}(\theta_i), \lambda_{I+1}(\theta_i))$ the equilibria with I and $I+1$ firms respectively. Consider the first order condition for quality provision (20) in both cases.

$$(\theta_i - q_I(\theta_i)) F^{I-1}(\theta_i) f(\theta_i) + \lambda_I(\theta_i) = 0 \quad (30)$$

$$(\theta_i - q_{I+1}(\theta_i)) F^I(\theta_i) f(\theta_i) + \lambda_{I+1}(\theta_i) = 0. \quad (31)$$

Multiplying the left hand side of (30) by $F(\theta_i)$ one obtains:

$$(\theta_i - q_I(\theta_i)) F^I(\theta_i) f(\theta_i) + F(\theta_i) \lambda_I(\theta_i) = 0.$$

Therefore, $\lambda_{I+1}(\theta_i)$ and $F(\theta_i) \lambda_I(\theta_i)$ provide comparable measures of quality distortions

$(\theta_i - q_{I+1}(\theta_i))$ and $(\theta_i - q_I(\theta_i))$. Now consider the first order conditions for utility provision in the two cases, and omit the argument (θ_i) for ease of notation:

$$\lambda'_I = F^{I-1}f - \frac{S_I - U_I}{q_I} (I-1) F^{I-2} f^2 \quad (32)$$

$$\lambda'_{I+1} = F^I f - \frac{S_{I+1} - U_{I+1}}{q_{I+1}} I F^{I-1} f^2. \quad (33)$$

(a) Market coverage is higher under $I + 1$ firms. Given the common terminal condition $q(\theta_H) = \theta_H$, this is equivalent to showing that quality provision $q_{I+1}(\theta_i)$ cannot cross $q_I(\theta_i)$ from below for the first time.

Suppose $q_{I+1}(\theta_i)$ crossed $q_I(\theta_i)$ from below at the point $\theta_i = \hat{\theta}_i$ and that $q_{I+1}(\theta_i) \leq q_I(\theta_i)$ for all $\theta_i \leq \hat{\theta}_i$. Then we would have $q_{I+1}(\hat{\theta}_i) = q_I(\hat{\theta}_i)$ and $q'_{I+1}(\hat{\theta}_i) > q'_I(\hat{\theta}_i)$. Differentiating (30) and (31), and evaluating them at $\hat{\theta}_i$, one obtains:

$$\begin{aligned} \lambda_{I+1}(\hat{\theta}_i) &= F(\hat{\theta}_i)\lambda_I(\hat{\theta}_i) \\ \lambda'_{I+1}(\hat{\theta}_i) &> \frac{d}{d\theta_i}(F(\hat{\theta}_i)\lambda_I(\hat{\theta}_i)). \end{aligned}$$

However, from conditions (32) and (33), one can write

$$\begin{aligned} \lambda'_{I+1} &= F^I f - \frac{S_{I+1} - U_{I+1}}{q_{I+1}} I F^{I-1} f^2 \\ &= F\lambda'_I - \frac{S_{I+1} - U_{I+1}}{q_{I+1}} I F^{I-1} f^2 + \frac{S_I - U_I}{q_I} (I-1) F^{I-1} f^2 \\ &= \frac{d}{d\theta_i}(F\lambda_I) - f\lambda_I - \frac{S_{I+1} - U_{I+1}}{q_{I+1}} I F^{I-1} f^2 + \frac{S_I - U_I}{q_I} (I-1) F^{I-1} f^2. \quad (34) \end{aligned}$$

Now, since $q_{I+1}(\theta_i) \leq q_I(\theta_i) \forall \theta_i \leq \hat{\theta}_i$, then $U_{I+1}(\hat{\theta}_i) \leq U_I(\hat{\theta}_i)$. Hence at $\theta_i = \hat{\theta}_i$, $(S_I - U_I)/q_I \leq (S_{I+1} - U_{I+1})/q_{I+1}$. Then, from (34), one obtains

$$\begin{aligned} \lambda'_{I+1} - \frac{d}{d\theta_i}(F\lambda_I) &= -f\lambda_I - F^{I-1}f^2 \left(\frac{S_{I+1} - U_{I+1}}{q_{I+1}} I - \frac{S_I - U_I}{q_I} (I-1) \right) \\ &< -f\lambda_I - F^{I-1}f^2 \frac{S_I - U_I}{q_I}. \end{aligned}$$

Then, using Lemma 2 and (30), the following inequality can be established,

$$\begin{aligned}
\lambda'_{I+1} - \frac{d}{d\theta_i}(F\lambda_I) &< -f\lambda_I - F^{I-1}f^2 \frac{S_I - U_I}{q_I} \\
&= (\hat{\theta}_i - q_I)F^{I-1}f^2 - F^{I-1}f^2 \frac{S_I - U_I}{q_I} \\
&= F^{I-1}f^2 \left(-\frac{q_I}{2} + \frac{U_I}{q_I} \right) \leq 0,
\end{aligned}$$

which contradicts $\lambda'_{I+1}(\hat{\theta}_i) > d(F(\hat{\theta}_i)\lambda_I(\hat{\theta}_i))/d\theta_i$. Hence, the possibility of $q_{I+1}(\theta_i) \leq q_I(\theta_i)$ for all $\theta_i \leq \hat{\theta}_i$ is ruled out. Incidentally, this also proves that quality provision with $I + 1$ firms cannot be everywhere lower than under only I firms (let $\hat{\theta}_i = \theta_H$).

(b) Every type's utility increases: $\forall \theta_i, U_{I+1}(\theta_i) \geq U_I(\theta_i)$. Note that if there exists a type θ_i such that $q_{I+1}(\theta_i) < q_I(\theta_i)$, then, given the common terminal condition, the function q_{I+1} must cross q_I from below at least once (potentially at θ_H). Therefore, there must also exist a type $\hat{\theta}_i$ for which

$$\lambda_{I+1}(\hat{\theta}_i) = F(\hat{\theta}_i)\lambda_I(\hat{\theta}_i) \text{ and} \quad (35)$$

$$\lambda'_{I+1}(\hat{\theta}_i) > \frac{d}{d\theta_i}(F(\hat{\theta}_i)\lambda_I(\hat{\theta}_i)) \quad (36)$$

Furthermore, if at $\theta_i = \hat{\theta}_i$ it is the case that $(S_I - U_I)/q_I < (S_{I+1} - U_{I+1})/q_{I+1}$ then the previous argument can be replicated to show that $\forall \theta_i, q_{I+1}(\theta_i) \geq q_I(\theta_i)$.

Suppose however, that for all $\hat{\theta}_i$ such that (35) and (36) hold, it is the case that $(S_I - U_I)/q_I \geq (S_{I+1} - U_{I+1})/q_{I+1}$. This inequality may be expressed as

$$\begin{aligned}
\frac{S_I - U_I}{q_I} &\geq \frac{S_{I+1} - U_{I+1}}{q_{I+1}} \\
\Leftrightarrow \frac{q_{I+1}}{2} + \frac{U_{I+1}}{q_{I+1}} &\geq \frac{q_I}{2} + \frac{U_I}{q_I}.
\end{aligned} \quad (37)$$

However, inequality (37) must also hold for points $\hat{\theta}_i$ where $q_{I+1}(\hat{\theta}_i) = q_I(\hat{\theta}_i)$. Therefore,

$$U_{I+1} \geq U_I, \quad \forall \hat{\theta}_i : q_{I+1}(\hat{\theta}_i) = q_I(\hat{\theta}_i).$$

This means that $U_{I+1} \geq U_I$ at all points where q_{I+1} crosses q_I from below. These points represent the local minima of the function $g(\theta_i) = U_{I+1}(\theta_i) - U_I(\theta_i)$. Since $g(\hat{\theta}_i) \geq 0$ it can be concluded that:

$$U_{I+1}(\theta_i) \geq U_I(\theta_i) \text{ for all } \theta_i.$$

(c) Quality provision increases everywhere. We have established that $U_{I+1}(\theta_i) \geq U_I(\theta_i)$ for all θ_i and that $q_{I+1}(\theta_i)$ must cross $q_I(\theta_i)$ from above for the first time. Therefore, if there exists a point for which $q_{I+1}(\theta_i) < q_I(\theta_i)$, using the fact the common terminal conditions, it follows that there must exist a point $\tilde{\theta}_i$ for which

$$\begin{aligned} q_{I+1}(\tilde{\theta}_i) &< q_I(\tilde{\theta}_i) \\ q'_{I+1}(\tilde{\theta}_i) &= q'_I(\tilde{\theta}_i). \end{aligned}$$

Using the expression for the derivative of the quality provision schedule from condition (23)

$$\begin{aligned} q'_I(\tilde{\theta}_i) &= 2 + \left(\frac{U_I(\tilde{\theta}_i)}{q_I(\tilde{\theta}_i)} - \frac{q_I(\tilde{\theta}_i)}{2} \right) \frac{f(\tilde{\theta}_i)(I-1)}{F(\tilde{\theta}_i)} \\ &= 2 + \left(\frac{U_{I+1}(\tilde{\theta}_i)}{q_{I+1}(\tilde{\theta}_i)} - \frac{q_{I+1}(\tilde{\theta}_i)}{2} \right) \frac{f(\tilde{\theta}_i)I}{F(\tilde{\theta}_i)} \\ &= q'_{I+1}(\tilde{\theta}_i). \end{aligned}$$

However, note that $U_{I+1}(\tilde{\theta}_i) \geq U_I(\tilde{\theta}_i)$ and $q_{I+1}(\tilde{\theta}_i) < q_I(\tilde{\theta}_i)$ imply

$$\left(\frac{U_{I+1}(\tilde{\theta}_i)}{q_{I+1}(\tilde{\theta}_i)} - \frac{q_{I+1}(\tilde{\theta}_i)}{2} \right) \frac{f(\tilde{\theta}_i)I}{F(\tilde{\theta}_i)} > \left(\frac{U_I(\tilde{\theta}_i)}{q_I(\tilde{\theta}_i)} - \frac{q_I(\tilde{\theta}_i)}{2} \right) \frac{f(\tilde{\theta}_i)(I-1)}{F(\tilde{\theta}_i)}.$$

Therefore, it is impossible that the two quality schedules increase at the same rate.

(2.) Because consumers have single-unit demand, under anonymous pricing, the multiproduct monopolist will sell product i to all consumers θ for which $\theta_i = \max_j \{\theta_j\}$. Conditional on selling product i , the monopolist will offer the Mussa and Rosen (1978) monopoly quality schedule for the distribution of the highest order statistic $F^I(\theta_i)$.

Denote the Mussa and Rosen (1978) quality provision under distribution $F(\theta_i)$ by

$$q_F(\theta_i) \triangleq \theta_i - \frac{1 - F(\theta_i)}{f(\theta_i)}.$$

Now consider two different distributions $F(\theta_i)$ and $G(\theta_i)$ with the associated quality functions $q_F(\theta_i)$ and $q_G(\theta_i)$. Following (for example) Krishna (2002), one can show that if $F(\theta_i)$ dominates $G(\theta_i)$ in terms of the likelihood ratio (*i.e.* if $f(\theta_i)/g(\theta_i)$ is nondecreasing), then $q_F(\theta_i) \leq q_G(\theta_i)$ for all θ_i . It can be easily shown that the distribution $F^{I+1}(\theta_i)$ dominates

$F^I(\theta_i)$ in terms of the likelihood ratio. The ratio of the densities is given by

$$\frac{dF^{I+1}(\theta_i)/d\theta_i}{dF^I(\theta_i)/d\theta_i} = \frac{(I+1)F^I(\theta_i)f(\theta_i)}{IF^{I-1}(\theta_i)f(\theta_i)} = \frac{(I+1)F(\theta_i)}{I},$$

which is clearly increasing in θ_i . Therefore, quality provision is decreasing in I for all θ_i . It immediately follows that market coverage (by each product) is (weakly) decreasing in I . Finally, since in the monopoly problem $U'(\theta_i) = q(\theta_i)$ and $U(\theta_L) = 0$ for all I , it follows that information rents $U(\theta_i)$ are decreasing in I for all θ_i . \square

Proof of Proposition 4: The Hamiltonian for each firm i 's problem may be written as

$$H(q, U, \lambda) = \int_{\theta_L}^{\theta_j^*(U_i(\theta_i), U_j)} (S(q_i(\theta_i), \theta_i) - U_i(\theta_i)) f(\theta_i, \theta_{-i}) d\theta_{-i} + \lambda q_i. \quad (38)$$

The necessary conditions for a symmetric equilibrium (15)-(17) are then obtained by differentiating (38) with respect to q_i and U_i , and by imposing the transversality condition $\lambda(\theta_H) = 0$. \square

Example: Analytical solution in the linear-quadratic model Let $I = 2$ and assume the support of the type distribution satisfies $\theta_H = 5\theta_L/2$. In this case, it is immediate to verify that (9)-(11) admit a quadratic solution, which is given by

$$U(\theta_i) = \frac{1}{6}(\theta_i - \theta_L)(5\theta_i - 2\theta_H), \quad (39)$$

and hence by

$$q(\theta_i) = \frac{5}{3}\theta_i - \frac{2}{3}\theta_H, \quad (40)$$

from which we can immediately verify that $q(\theta_L) = 0$. From first order condition (9), we can solve for the associated costate variable,

$$\tilde{\lambda}(\theta_i) = -\frac{2(\theta_H - \theta_i)(\theta_i - \theta_L)}{3(\theta_H - \theta_L)^2}.$$

In order to check that the second order conditions are satisfied, suppose all firms $j \neq i$ offer the rent function $U(\theta_j)$ given in (39), and consider the Hamiltonian:

$$H(\theta_i, q, U, \lambda) = (\theta_i q - q^2/2 - U) \frac{\theta_j^*(U) - \theta_L}{(\theta_H - \theta_L)^2} + \lambda q.$$

where the threshold type θ_j^* is the solution to $U = U(\theta_j)$, and therefore,

$$\theta_j^*(U) = \theta_L + \sqrt{\frac{6}{5}U}.$$

Now consider the maximized Hamiltonian

$$H^*(\theta_i, U, \tilde{\lambda}) = (\theta_i q^*(U) - q^*(U)^2/2 - U) \frac{\theta_j^*(U) - \theta_L}{(\theta_H - \theta_L)^2} + \tilde{\lambda}(\theta_i) q^*(U), \quad (41)$$

where

$$\begin{aligned} q^*(U) &= \theta_i + \frac{\tilde{\lambda}(\theta_i)}{\theta_j^*(U) - \theta_L} (\theta_H - \theta_L)^2 \\ &= \theta_i - \frac{10(\theta_H - \theta_i)(\theta_i - \theta_L)}{3\sqrt{30U}}. \end{aligned}$$

We can plug $q^*(U)$, $\tilde{\lambda}(\theta_i)$ and $\theta_j^*(U)$ into (41), and verify that the resulting expression is strictly concave in U for all $\theta_i \in [\theta_L, \theta_H]$. Therefore, we can apply the Arrow sufficient condition (see Seierstad and Sydsaeter (1987), Theorem 3.17), and conclude that (39)-(40) are an equilibrium.

Example: FMG Copula A tractable functional form to introduce correlation is given by the Farlie-Gumbel-Morgenstern (FMG) family of copulas (see Nelsen (2006) for a detailed description and for the properties of this family). Given identical marginal distribution functions $F(\theta_i)$, and a parameter $\gamma \in [-1, 1]$, define the joint *cdf* and *pdf* (for the case of two firms) as

$$\begin{aligned} H(\theta_1, \theta_2) &= F(\theta_1) F(\theta_2) (1 + \gamma(1 - F(\theta_1))(1 - F(\theta_2))) \\ h(\theta_1, \theta_2) &= f(\theta_1) f(\theta_2) (1 + \gamma(1 - 2F(\theta_1))(1 - 2F(\theta_2))). \end{aligned}$$

The equilibrium market share function may be written as follows

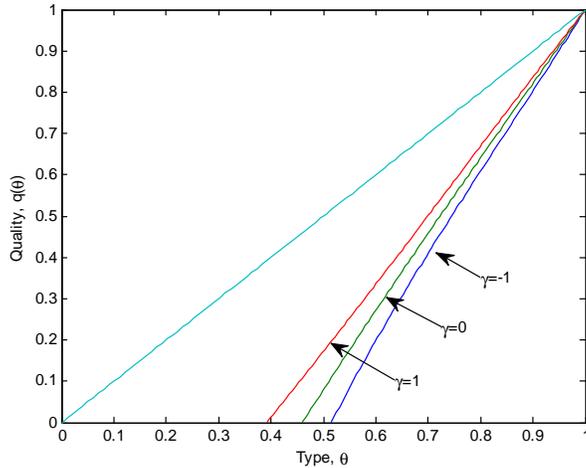
$$G(\theta_i) = \int_{\theta_L}^{\theta_i} h(\theta_i, t) dt = (1 + \gamma(1 - F(\theta_i))(1 - 2F(\theta_i))) F(\theta_i) f(\theta_i)$$

and

$$h(\theta_i, \theta_i) = f^2(\theta_i) (1 + \gamma(1 - 2F(\theta_i))^2).$$

These equations are equivalent to (20)-(22) when $\gamma = 0$. In Figure 6, let $F(\theta_i)$ be the uniform distribution on $[0, 1]$, let costs be quadratic, and draw the symmetric equilibrium quantity provision schedules for different values of the correlation parameter γ .

Figure 6: Duopoly Quality Provision - FMG Copula



References

- ANDERSON, S., A. DE PALMA, AND J. THISSE (1992): *Discrete Choice Theory of Product Differentiation*. MIT Press.
- ARMSTRONG, M., AND J. VICKERS (2001): “Competitive Price Discrimination,” *RAND Journal of Economics*, 32(4), 579–605.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890.
- BERRY, S., AND A. PAKES (2007): “The Pure Characteristics Demand Model,” *International Economic Review*, 48(4), 1193–1225.
- BUSSE, M., AND M. RYSMAN (2005): “Competition and price discrimination in yellow pages advertising,” *Rand Journal of Economics*, 36(2), 378–390.
- CHAMPSAUR, P., AND J. ROCHET (1989): “Multiproduct Duopolists,” *Econometrica*, 57(3), 533–557.
- CHEN, Y., AND M. RIORDAN (2008): “Price-increasing competition,” *RAND Journal of Economics*, 39(4), 1042–1058.
- ELLISON, G. (2005): “A Model of Add-On Pricing,” *Quarterly Journal of Economics*, 120(2), 585–637.
- HARTMANN, W., AND V. VIARD (2006): “Quantity-based price discrimination using frequency reward programs,” Discussion paper, Stanford Graduate School of Business.
- HENDEL, I. (1999): “Estimating multiple-discrete choice models: An application to computerization returns,” *Review of Economic Studies*, 66(2), 423–446.
- KRISHNA, V. (2002): *Auction Theory*. Academic Press.
- LANCASTER, K. (1971): *Consumer demand: a new approach*. Columbia University Press New York.
- LEDERMAN, M. (2007): “Do enhancements to loyalty programs affect demand? The impact of international frequent flyer partnerships on domestic airline demand,” *RAND Journal of Economics*, 38(4), 1134–1158.
- MARTIMORT, D., AND L. STOLE (2002): “The Revelation and Delegation Principles in Common Agency Games,” *Econometrica*, 70(4), 1659–1673.

- MCMANUS, B. (2007): “Nonlinear pricing in an oligopoly market: The case of specialty coffee,” *RAND Journal of Economics*, 38(2), 512.
- MIRAVETE, E. (2009): “Competing with menus of tariff options,” *Journal of the European Economic Association*, 7(1), 188–205.
- MIRAVETE, E., AND L. RÖLLER (2004): “Competitive non-linear pricing in duopoly equilibrium: the early US cellular telephone industry,” Discussion paper, University of Texas, Austin.
- MUSSA, M., AND S. ROSEN (1978): “Monopoly and Product Quality,” *Journal of Economic Theory*, 18(2), 301–317.
- NELSEN, R. (2006): *An Introduction to Copulas*. Springer Verlag.
- PERLOFF, J., AND S. SALOP (1985): “Equilibrium with product differentiation,” *Review of Economic Studies*, 52(1), 107–120.
- ROCHET, J., AND L. STOLE (2002): “Nonlinear Pricing with Random Participation,” *Review of Economic Studies*, 69(1), 277–311.
- SEADE, J. (1977): “On the shape of optimal tax schedules,” *Journal of Public Economics*, 7(2), 203–235.
- SEIERSTAD, A., AND K. SYDSAETER (1977): “Sufficient conditions in optimal control theory,” *International Economic Review*, 18(2), 367–391.
- (1987): *Optimal Control with Economic Applications*. North-Holland, Amsterdam.
- SONG, M. (2007): “Measuring Consumer Welfare in the CPU Market: An Application of the Pure Characteristics Demand Model,” *RAND Journal of Economics*, 38(2), 429–446.
- SPULBER, D. (1989): “Product Variety and Competitive Discounts,” *Journal of Economic Theory*, 48(2), 510–525.
- STOLE, L. (1995): “Nonlinear Pricing and Oligopoly,” *Journal of Economics and Management Strategy*, 4(4), 529–562.
- (2007): “Price Discrimination and Competition,” *Handbook of Industrial Organization*, Volume 3.
- YANG, H., AND L. YE (2008): “Nonlinear pricing, market coverage, and competition,” *Theoretical Economics*, 3(1), 123–153.