



Estimating Discrete-Choice Models of Product Differentiation

Steven T. Berry

The RAND Journal of Economics, Vol. 25, No. 2 (Summer, 1994), 242-262.

Stable URL:

<http://links.jstor.org/sici?sici=0741-6261%28199422%2925%3A2%3C242%3AEDMOPD%3E2.0.CO%3B2-Z>

The RAND Journal of Economics is currently published by The RAND Corporation.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/rand.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Estimating discrete-choice models of product differentiation

Steven T. Berry*

This article considers the problem of “supply-and-demand” analysis on a cross section of oligopoly markets with differentiated products. The primary methodology is to assume that demand can be described by a discrete-choice model and that prices are endogenously determined by price-setting firms. In contrast to some previous empirical work, the techniques explicitly allow for the possibility that prices are correlated with unobserved demand factors in the cross section of markets. The article proposes estimation by “inverting” the market-share equation to find the implied mean levels of utility for each good. This method allows for estimation by traditional instrumental variables techniques.

1. Introduction

■ Traditional “supply-and-demand” analysis has long been a staple of empirical economics. This analysis attempts to uncover cost and demand information from market data under the assumption of a static, perfectly competitive equilibrium. In recent years, increasing attention has been paid to estimating demand and cost parameters under imperfect competition. Much, though not all, of this existing literature on estimation under imperfect competition is tied to homogeneous goods markets.

This article considers the problem of estimating supply-and-demand models in markets with product differentiation. In common with some previous articles, market demand is derived from a general class of discrete-choice models of consumer behavior. The utility of consumers depends on product characteristics and individual taste parameters; product-level market shares are then derived as the aggregate outcome of consumer decisions. Firms are modelled as price-setting oligopolists, and endogenous market outcomes are derived from an assumption of Nash equilibrium in prices.

The proposed estimation methods do not require the econometrician to observe all relevant product characteristics. The presence of unobserved product characteristics allows for a product-level source of sampling error. More importantly, it reintroduces the econometric problem of endogenous prices (or “simultaneity”) that is familiar from studies of homogeneous goods markets. In these studies, the “error” in the demand equation is usually

* Yale University.

I would like to thank Ariel Pakes, James Levinsohn, Alvin Klevorick, Richard Levin, Robert Porter, and Michael Whinston for helpful discussions and/or comments on earlier drafts. Referees and a Coeditor of this Journal made invaluable suggestions. Johan Gauderis provided excellent programming assistance. This work was partially supported by NSF grant no. SES-9122672 (received jointly with James Levinsohn and Ariel Pakes) and by an Olin Fellowship at the National Bureau of Economic Research. I am, of course, responsible for remaining errors.

given an explicit structural interpretation as representing unobserved (by the econometrician) demand factors. These demand factors are, by inspection of the supply curve, seen to be correlated with prices. It is well known that ignoring the correlation between price and the demand error frequently leads to findings of upward sloping demand curves and other anomalies.

As is illustrated below, similar problems arise in the study of differentiated products markets when some product characteristics are unobserved. Importantly, unobserved product characteristics are a feature in many markets that economists study. Characteristics such as style are inherently difficult to quantify but are frequent determinants of demand. In some markets, products may be physically similar but differ in consumers' perceptions about quality, durability, status, or service at point-of-sale. Also, in practice, the number of product characteristics that are important to consumers may be much larger than the number of observations available to the econometrician, making it impossible to estimate the separate effects of each characteristic.

The endogeneity of prices that follows from the presence of unobserved product characteristics is not just an econometric quibble. A later set of Monte Carlo results will demonstrate that (as in the homogeneous goods case) estimation methods that ignore the endogeneity of prices in the presence of unobserved product characteristics can be severely misleading. In a more concrete example, the importance of price endogeneity is illustrated by Trajtenberg's (1989) careful study of the medical CT scanner market. This study notes that, in some cases, prices appear to have a positive effect on demand. In Trajtenberg's model, this finding implies that an increase in price increases consumer benefits. Consistent with the arguments made here for the importance of unobserved product characteristics, Trajtenberg attributes this anomaly to the presence of unobserved product quality. Empirical results on the automobile industry, reported in Berry, Levinsohn, and Pakes (1993), are also consistent with the importance of accounting for unobserved product characteristics.

In the homogeneous goods case, demand parameters can be consistently estimated in the presence of unobserved demand factors via the use of traditional instrumental variables methods. However, in the context of discrete-choice models, both prices and unobserved product characteristics enter demand equations in a nonlinear fashion. This frustrates any straightforward application of instrumental variables methods.

This article introduces a method for avoiding the nonlinear instrumental variables problem. This method inverts the function defining market shares to uncover the mean utility levels of various products as specified by the primitives of the model. These mean utility levels can then be related to product characteristics and prices using instrumental variables techniques. The mean utility method is applicable to a wide class of discrete-choice models, does not rely on the existence of a unique equilibrium, and frequently involves a smaller computational burden as compared to previously used alternatives.

After a brief description of related models, I shall outline the basic framework of discrete-choice demand and oligopolistic pricing. The method of recovering mean utility levels is then introduced and discussed. To illustrate, I show how to implement the method in several special cases, including logit, nested logit, and the vertical differentiation model. Final sections of the article discuss some problems with and extensions of this approach and also provide some Monte Carlo evidence.

2. Previous empirical models of differentiated products oligopoly

- Markets with perfectly homogeneous goods are empirically rare, although not unknown. Accordingly, empirical models of differentiated products oligopoly have received some attention. Empirical studies of differentiated products oligopoly address topics as varied as the mode of market conduct (e.g., Bresnahan, 1987), the welfare effects of the

introduction of new products (Trajtenberg, 1989), or of deregulation (Morrison and Winston, 1986). The focus in this article will be on estimating structural demand and cost parameters without reference to specific applications; applications that are related to the concerns of this article include those of Berry (1990), who estimates separate cost and demand effects of airline hubbing, and Berry, Levinsohn, and Pakes (1993), whose article embodies and extends the ideas of this article in an empirical study of the automobile industry.¹

Perhaps the simplest approach for dealing with endogenous prices in a differentiated products industry is to posit simple aggregate (that is, market level) demand curves in which quantity demanded is decreasing in a firm's own price and increasing in the price of its rivals. Consider, for example, the constant elasticity framework:

$$\ln(q_j) = \alpha_j + \sum_k \eta_{jk} \ln(p_k) + \epsilon_j, \quad (1)$$

where η_{jk} is the elasticity of good j with respect to the price of good k . When ϵ and p are correlated, the demand system described by (1) can be easily estimated by traditional instrumental variables techniques. The well-known problem, however, is that a system of N goods gives N^2 elasticities to estimate, which is a very large number in many real-world applications. For example, in the automobile industry model of Bresnahan (1987), there are close to 100 distinct products, implying almost 10,000 separate elasticities.

It is possible to avoid the problem of "too many elasticities" by placing *a priori* restrictions on the pattern of cross-price elasticities. For example, a researcher could decide that many of the cross-price elasticities are equal to zero or that many sets of cross-price elasticities are equal to each other. This approach is obviously arbitrary, and in many markets, economic theory will provide little guidance on such restrictions.

It is desirable, therefore, to put some structure on the demand problem in order to reduce the number of demand parameters. This article will impose such structure by making assumptions on consumer utility. The utility of a given consumer is assumed to depend on the characteristics of the chosen product, on random consumer "tastes," and on a small set of parameters to be estimated. Market demand is then derived as the aggregation of individual consumer choices. Explicitly deriving aggregate demand from consumer choices has several advantages. This approach avoids the problems of (1) by deriving all the relevant demand elasticities from a much smaller number of utility parameters. Also, the resulting model can make predictions about the demand for new products and about the demand for dissimilar products found in different markets. Finally, such a model allows us to move easily between statements about aggregate demand and statements about consumer utility.

Discrete-choice models are a common, tractable, and parsimonious method for obtaining the desired structure on demand. This parsimony comes at some cost, as the models rule out the purchase of multiple items and do not easily incorporate dynamic aspects of demand. Furthermore, they typically place important parametric restrictions on the demand structure. Discrete-choice models of product demand have, of course, a long history in econometrics, most notably influenced by McFadden (e.g., McFadden (1974)). Recently, discrete-choice models have received increasing attention in the theoretical literature on differentiated products oligopoly, either as a means of justifying particular assumptions on aggregate demand or as an independent focus of analysis.²

¹ The model of this article is also related to the hedonic pricing models of Griliches (1971), Rosen (1974), and Epple (1987). Indeed, my model implicitly produces a hedonic equilibrium pricing function that depends on product characteristics. However, the focus in this article on structural estimation with price-setting firms and unobserved demand characteristics differs from the typical focus in the hedonic literature.

² Theoretical works that apply discrete-choice models to the study of oligopoly product differentiation include Shaked and Sutton (1982), Sattinger (1984), Perloff and Salop (1985), Anderson, DePalma, and Thisse (1989), and Caplin and Nalebuff (1991).

In an empirical study that is related to the approach taken here, Bresnahan (1987) uses a discrete-choice model with vertically differentiated products to study the automobile market. In this model, consumers care about product quality, which is modelled as depending on observed product characteristics. Bresnahan's model has several features that I shall also employ below. Consumer utility will be modelled as depending on product characteristics, consumers will be allowed to purchase an "outside good," and explicit use can be made of the first-order conditions of price-setting firms. In a defect shared with Bresnahan (and with nearly all empirical studies of differentiated products), product characteristics will be treated as exogenous, although product prices are determined within the model. However, in contrast to Bresnahan's model, I shall consider the presence of unobserved product characteristics and shall discuss a much broader class of discrete-choice demand models. Both of these issues will suggest the use of estimation methods that are substantially different than those used by Bresnahan.

3. The model

■ The primitives of the model are the characteristics of products, consumer preferences, and the equilibrium notion. All characteristics and all decisions are assumed to be observable by all participants in the market. However, the econometrician does not observe all of the product characteristics and may not observe the decisions of individual consumers. The econometrician is assumed to observe the market outcomes of price and quantities sold by each firm.

For now, I shall assume that we observe a large number, R , of independent markets. There are N_r firms in market r , with each firm producing one product. For product j in market r , observed characteristics are denoted by the vector $z_{jr} \in \mathbb{R}^k$. (For simplicity, I often drop the market subscript r .) The elements of z_j include characteristics that affect demand (x_j) and marginal costs (w_j). The characteristics of all firms in the market are included in the vector $\mathbf{z} = (z_1, \dots, z_N)$. Similarly, $\mathbf{x} = (x_1, \dots, x_N)$ and $\mathbf{w} = (w_1, \dots, w_N)$.

The unobserved characteristics of product j are (ξ_j, ω_j) , where ξ_j is an unobserved demand characteristic and ω_j is an unobserved cost variable. The unobserved characteristics in a market are assumed to be mean independent of \mathbf{z} and independent across markets. Together, \mathbf{z} , $\boldsymbol{\xi}$, and $\boldsymbol{\omega}$ define the data that are causally "exogenous" to the firm's pricing decisions. Assuming that the unobservables are mean independent of the product characteristics amounts to treating the product characteristics as econometrically exogenous. While common, this assumption is unreasonable in many cases. The problems raised by models with endogenous product characteristics are discussed in Section 9.

□ **The discrete choice model.** The utility of consumer i for product j depends on the characteristics of the product and the consumer: $U(x_j, \xi_j, p_j, \nu_i, \theta_d)$, where x_j , ξ_j , p_j , and θ_d are observed product characteristics, unobserved (by the econometrician) product characteristics, and price and demand parameters, respectively. The term ν_i captures consumer-specific terms that are not observed by the econometrician. All the estimators discussed below require parametric assumptions on the consumer-specific variables; these assumptions are analogous to the choice of a functional form for a homogeneous goods demand equation. Different choices for the utility function and for the density of ν will have important implications for the resulting model.

I shall focus on a simple random coefficients specification for utility, which is quite simple, yet flexible enough to illustrate the main points of the article. In this specification, the utility of consumer i for product j is given by

$$u_{ij} = x_j \tilde{\beta}_i - \alpha p_j + \xi_j + \epsilon_{ij}, \quad (2)$$

where the (unobserved to the econometrician) consumer-specific taste parameters are $\tilde{\beta}_i$ and ϵ_{ij} . The parameter α is written as invariant across consumers, although this is not necessary. The term ξ_j might be thought of as the mean of consumers' valuations of an unobserved product characteristic such as product quality, while the ϵ_{ij} represents the distribution of consumer preferences about this mean.

For simplicity, I shall decompose consumer i 's taste parameter for characteristic k as

$$\tilde{\beta}_{ik} = \beta_k + \sigma_k \zeta_{ik}, \quad (3)$$

where β_k is the mean level of the taste parameter for product k and the mean-zero ζ_{ik} has, e.g., an identically and independently distributed standard normal distribution across individuals and characteristics. Combining (2) with (3), we can write

$$u_{ij} = x_j \beta + \xi_j - \alpha p_j + v_{ij},$$

with

$$v_{ij} = \left[\sum_k x_{jk} \sigma_k \zeta_{ik} \right] + \epsilon_{ij}. \quad (4)$$

The term v_{ij} is thus a mean-zero, heteroskedastic error that captures the effects of the random taste parameters. I denote the mean utility level of product j , which will play an important role below, as

$$\delta_j \equiv x_j \beta - \alpha p_j + \xi_j. \quad (5)$$

It is common in traditional logit and probit models to assume that the variation in consumer tastes enters only through the additive term ϵ_{ij} , which is assumed to be identically and independently distributed across consumers and choices. While parsimonious, it is rarely noticed that this assumption places very strong restrictions on the pattern of cross-price elasticities from the estimated model. I shall here summarize the more detailed discussion of this problem that is found in Berry, Levinsohn, and Pakes (1993).³ In the model with identically and independently distributed consumer tastes, only the mean utility levels, δ_j , differentiate the products. Therefore, all properties of market demand, including market shares and elasticities, are determined solely by the δ_j . In particular, cross-price elasticities can only depend on the value of δ_j , with no additional effect from individual product characteristics or prices. In the automobile market, for example, this property implies that any pair of cars (j, k) with the same pair of market shares (s_j, s_k) will have the same cross-price elasticity with any given third product. This property will hold regardless of whether both j and k are small inexpensive cars or one car is a subcompact and one is a luxury car. It is important to note that this property is a function of the identically and independently distributed additive error and not of any specific distributional assumption (such as logit) on the errors.

Models that have random coefficients, $\tilde{\beta}_i$, on the product characteristics avoid the problem of *a priori* unreasonable substitution effects. An increase in the price of product j affects only those consumers who currently purchase good j . In the random coefficients model, these consumers will typically have values for $\tilde{\beta}_i$ that differ from the mean. These selected consumers will therefore substitute toward a particular group of products; generally, these products will "resemble" product j . This same property will hold for many

³ See also the earlier, related discussions in Tversky (1972), Hausman and Wise (1978), and McFadden (1981).

specifications in which consumer and product characteristics are interacted, so that differences between consumers have a systematic effect on their preferences. For example, this property will hold in studies that use consumer data and interact observed characteristics of consumers with product characteristics.

Given the functional form assumptions, the discrete-choice market share function, s_j , is derived in the usual way. Each consumer purchases one unit of the good that gives the highest utility. That is, conditional on the characteristics $(\mathbf{x}, \boldsymbol{\xi})$ and prices \mathbf{p} , consumer i will purchase one unit of good j if and only if for all $k \geq 0$ and $k \neq j$, $U(x_j, \xi_j, p_j, \mathbf{v}_i, \theta_d) > U(x_k, \xi_k, p_k, \mathbf{v}_i, \theta_d)$. This implicitly defines the set of unobservable taste parameters, v_{ij} , that result in the purchase of good j . Define the set of consumer unobservables that lead to the consumption of good j as $A_j(\boldsymbol{\delta}) = \{\mathbf{v}_i \mid \delta_j + v_{ij} > \delta_k + v_{ik}, \forall k \neq j\}$. The market share of the j th firm is then the probability that \mathbf{v}_i falls into the region A_j . Given a distribution, $F(\cdot, \mathbf{x}, \sigma)$, for \mathbf{v} , with density $f(\cdot, \mathbf{x}, \sigma)$, this market share is

$$s_j(\boldsymbol{\delta}(\mathbf{x}, \mathbf{p}, \boldsymbol{\xi}), \mathbf{x}, \boldsymbol{\theta}) = \int_{A_j(\boldsymbol{\delta})} f(\mathbf{v}, \mathbf{x}, \sigma_v) d\mathbf{v}, \quad (6)$$

where the integral is over the set of consumer unobservables implicitly defined by A_j .

To complete the specification of the demand system, we should discuss the size of the market that allows us to move between market shares and observed quantities in the presence of an outside alternative.

□ **Market size and the outside good.** The measure of consumers in a market is denoted M . This number is either observed as the population of a market or left as a parameter to be estimated. The observed output quantity of the firm is then⁴

$$q_j = M s_j(\mathbf{x}, \boldsymbol{\xi}, \mathbf{p}, \theta_d). \quad (7)$$

In addition to the competing products $j = 1, \dots, N$, I shall also assume the existence of an outside good, $j = 0$. Consumers may choose to purchase the outside good instead of one of the N "inside" products. The distinction is that the price of the outside good is not set in response to the prices of the inside goods. In the absence of an outside good, consumers are forced to choose from the inside good and demand depends only on differences in prices. Therefore, a general increase in prices will not decrease aggregate output; this is an unfortunate feature of some discrete models that have been applied to the empirical study of differentiated products markets (e.g., Morrison and Winston (1986)).

However, the presence of the outside good with market share s_0 means that observations on the output quantities of the N firms (q_1, \dots, q_N) are not sufficient to calculate the market shares of the $N + 1$ total alternatives. If the total market size M is directly observed, then s_j can be calculated easily as $s_j = q_j/M$. For example, Berry, Levinsohn, and Pakes (1993) set market size equal to the number of households in the economy. Otherwise, M will have to be estimated. When there is information on a number of markets, M can be parameterized as depending on market-level data (such as population) that vary across markets and that affect the aggregate level of output (e.g., Berry (1990)). Methods for estimating M will be application specific, and in the remainder of this article, I assume that M is observed.

⁴ If output quantity is formed as the sum of M distinct draws on consumer preferences, then (7) represents only the expected output quantity. Observed market shares would include a random error with a variance of $s_j(1 - s_j)/M$. I shall treat (7) as the observed output quantity, which is consistent with assuming a continuous measure of consumers, rather than M distinct consumers.

□ **The supply side.** The N firms in a market are assumed to be price setters, although alternate models of firm behavior are easy to incorporate. Total costs for firm j are given by the function $C_j(q_j, w_j, \omega_j, \gamma)$, and marginal costs are $c_j(q_j, w_j, \omega_j, \gamma)$, where γ is a vector of unknown parameters. Profits for firm j are then

$$\pi_j(\mathbf{p}, \mathbf{z}, \boldsymbol{\xi}, \omega_j, \theta) = p_j M_{\phi_j}(\mathbf{x}, \boldsymbol{\xi}, \mathbf{p}, \theta_d) - C_j(q_j, w_j, \omega_j, \gamma), \quad (8)$$

where $\theta = (\theta_d, \gamma)$.

Assuming the existence of a pure-strategy interior equilibrium, the price vector satisfies the usual first-order conditions⁵

$$[p_j - c_j(q_j, w_j, \omega_j, \gamma)][\partial \phi_j(\mathbf{x}, \boldsymbol{\xi}, \mathbf{p}, \theta_d) / \partial p_j] + \phi_j(\mathbf{x}, \boldsymbol{\xi}, \mathbf{p}, \theta_d) = 0$$

or

$$p_j = c_j + \phi_j / \left| \partial \phi_j / \partial p_j \right|. \quad (9)$$

If the N such first-order conditions define a unique equilibrium for all possible values of the observed and unobserved data and for all possible values of the parameters, then the first-order conditions implicitly define the “reduced-form” price, $\beta_j(\mathbf{z}, \boldsymbol{\xi}, \boldsymbol{\omega}, \theta)$, as a function of the exogenous data and the parameters. The equilibrium price, together with the demand function, then defines the reduced-form expression for equilibrium quantities: $q_j(\mathbf{z}, \boldsymbol{\xi}, \boldsymbol{\omega}, \theta) = M_{\phi_j}(\mathbf{x}, \boldsymbol{\xi}, \beta_j(\mathbf{z}, \boldsymbol{\xi}, \boldsymbol{\omega}, \theta), \theta_d)$. In a later section, I shall note that these reduced forms could be used as the basis of an estimation routine. The next section, however, introduces a simpler and more general method.

4. Estimating from the mean utility levels

■ The discrete-choice model of the last section is entirely traditional except for the unobserved product characteristic ξ_j . However, the presence of ξ_j raises a difficult econometric problem. Consider a demand equation that relates observed market shares, s_j , to the market shares that are predicted by the model, ϕ_j :

$$s_j = \phi_j(\mathbf{x}, \mathbf{p}, \boldsymbol{\xi}, \boldsymbol{\theta}). \quad (10)$$

The right-hand side of this equation contains both prices and product level demand errors. We expect the unobserved product characteristics to be correlated with prices; thus, the right-hand side prices are endogenous in the sense that they are correlated with the unobservables. Instrumental variables methods are a traditional solution to this problem of endogenous prices. However, the unobservables enter (10) in a nonlinear fashion, thus frustrating the application of traditional instrumental variables methods.

To solve this problem, I propose transforming the market shares so that the unobserved product characteristics appear as a linear term. Let us begin with the simple case in which the distribution of consumer unobservables is known, so that market shares depend only on mean utility levels:

$$s_j = \phi_j(\bar{\boldsymbol{\delta}})(j = 1, \dots, N). \quad (11)$$

At the true values of $\bar{\boldsymbol{\delta}}$ and of market shares, \mathbf{s} , these equations must hold exactly. (The

⁵ I shall consider only pure strategy equilibria in this article; mixed strategy equilibria would complicate the analysis considerably. Caplin and Nalebuff (1991) provide a useful discussion of the existence of equilibrium in this class of models.

distinction between the observed market shares s_j and the market share function $s_j(\delta)$ is important here.) The point is that the mean utility levels δ_j contain the aggregate error ξ_j ; therefore, conditional on the true values of δ (and given a density, f) the model should fit the data exactly.

The exact fit of the model conditional on the mean utility levels δ can be exploited in an estimation procedure. If the vector-valued equation $\mathbf{s} = s(\delta)$ can be inverted to produce the vector $\delta = s^{-1}(\mathbf{s})$, then the observed market shares (together with the distributional assumption on ν) uniquely determine the means of consumer utility for each good.

Under weak regularity conditions on the density of consumer unobservables, the existence of a unique $\delta^*(\mathbf{s})$ that satisfies $\mathbf{s} = s(\delta^*(\mathbf{s}))$ is established in the Appendix. There, I show that (conditional on setting the mean utility of the outside good, δ_o , equal to zero) the market share function is one-to-one. I also establish that, for every possibly observed vector of market shares, \mathbf{s} , there is a vector of utility means $\delta \in \mathbb{R}^{N+1}$ that will create that observed vector by the relation $\mathbf{s} = s(\delta)$. Thus, every vector of observed market shares can be explained by one and only one vector of utility means. For any density $f(\cdot, x)$, we can therefore calculate the vector δ from observations on the market shares alone.

The unique, calculated vector $\delta(\mathbf{s})$ can then be used in a simple estimation procedure. When the density of ν is known exactly, so that the market share function depends on no unknown parameters other than the vector δ , then the calculated mean utility levels can be treated as a known, nonlinear transformation of the market shares, \mathbf{s} . From (5), for the true values of (β, α) ,

$$\delta_j(\mathbf{s}) = x_j\beta - \alpha p_j + \xi_j. \quad (12)$$

We can treat (12) as an estimation equation and use standard instrumental variables techniques to estimate the unknown parameters. That is, we can run an appropriate instrumental variables regression of $\delta_j(\mathbf{s})$ on (x_j, p_j) to estimate (β, α) , treating ξ_j as an unobserved error term. The fact that $\delta_j(\mathbf{s})$ is a transformation of the original data on market shares is not important; except for the computational problem of inverting the market share function, this is little different than similar estimation procedures that take some other transformation of the observed data (e.g., logarithms) as a dependent variable.

The correlation of p_j with ξ_j suggests the use of instruments for prices. Cost variables that are excluded from x_j , such as input prices that vary across firms, are traditional instruments in homogeneous goods markets and they continue to be appropriate in the present context. Interestingly, in product differentiation models with exogenous characteristics, the characteristics of other firms ($x_k, k \neq j$) are also appropriate instruments. These are appropriate because they are excluded from the utility function (u_{ij} does not depend on x_k) and they are correlated with prices via the markups in the first-order conditions. This is a specific example of the general proposition that, in imperfectly competitive markets, demand-side instruments can be variables that affect markups as well as variables that affect marginal cost. Thus, it may be possible to obtain consistent estimates of demand parameters even in the absence of excluded "cost-side" variables. This idea was developed in detail by Berry, Levinsohn, and Pakes (1993).

The idea of estimating the demand parameters β and α by an instrumental variables regression of δ_j on characteristics and prices is similar to the homogeneous goods regression of output quantities on demand factors and prices. In the homogeneous goods case, demand parameters can be estimated with cost-side instruments under the relatively weak assumption that the demand error is uncorrelated with the instruments. Similarly, estimates of the demand parameters (β, α) can be obtained in the present case by inverting the market share function without the need for assumptions on either the parametric distribution of the unobservables, ξ , or on the actual process that generates prices.

In the differentiated products model, there is one demand equation, of the form in (12), for each good in each market. If we have access to a large sample of independent markets, then we can obtain consistent estimates of demand parameters by treating each market as a separate observation. This approach allows for arbitrary correlation between demand unobservables within markets. If, instead, we assume that the ξ_j are independent across firms, then the demand parameters β and α might be estimated from a dataset containing cross-sectional information on a large number of firms within a single market. This contrasts with the homogeneous goods case, in which a single market and single time period imply a single observation on demand.⁶

Next I present two simple special cases, in which it is quite easy to solve for mean utility levels as a function of observed market shares. Logit, the first example, is the best-known special case of the model in Section 3. The second example, the vertical differentiation model, is a simple variant of the random coefficients model which has been prominent in the empirical literature (e.g., Bresnahan (1987)).

Example: the logit model. Suppose we begin with the utility function in (2) and make the familiar assumption that $\tilde{\beta}_i = \beta$ (no random coefficients) and that ϵ_{ij} is identically and independently distributed across products and consumers with the "extreme value" distribution function $\exp(-\exp(-\epsilon))$. The market share of product j is then given by the well-known logit formula

$$s_j(\mathbf{\delta}) = e^{\delta_j} / \left(\sum_{k=0}^N e^{\delta_k} \right). \quad (13)$$

With the mean utility of the outside good normalized to zero,

$$\ln(s_j) - \ln(s_o) = \delta_j \equiv x_j\beta - \alpha p_j + \xi_j, \quad (14)$$

so δ_j is uniquely identified directly from a simple algebraic calculation involving market shares. Thus, the logit case suggests a simple instrumental variables regression of differences in log market shares on (x_j, p_j) . This case is unusual as, in many cases, $\mathbf{\delta}$ must be solved for numerically. Unfortunately, as noted, the logit model produces unreasonable substitution patterns.

Example: the vertical differentiation model. In the vertical model of product differentiation (see Shaked and Sutton (1982) and Bresnahan (1987)) consumers agree about the quality of each good but disagree about the value of quality. Consider the utility function

$$u_{ij} = \psi_j v_i - p_j, \quad (15)$$

with ψ_j representing the quality of product j and v_i a scalar random variable representing the value that consumer i places on quality. Assume that quality depends linearly on the observed and unobserved characteristics of product j :

$$\psi_j = x_j\beta + \xi_j. \quad (16)$$

Also, order the products in increasing quality, $\psi_o < \psi_1 < \dots < \psi_N$, and denote the cumulative distribution of v as $F(v)$, with density $f(v)$.

⁶ However, in the case of a single market, prices are correlated across firms (via the first-order conditions) even if the ξ_j are independent across firms. This raises important econometric issues of dependence that I shall not address here.

If ν_i is assumed to have mean one (which is just a normalization on the units of quality), then mean utility is still $\delta_j = x_j\beta - \alpha p_j + \xi_j$, with $\alpha \equiv 1$. In this model, ν_i is effectively a single random coefficient that interacts with observed and unobserved characteristics. That is, we can write utility as a slight variant of (4):

$$u_{ij} = \left[\sum_k x_{jk}\beta_k \nu_i \right] - p_j + \nu_i \xi_j.$$

In some respects, the vertical model is the opposite extreme of the logit model. In the logit model, there are as many consumer characteristics, ϵ_{ij} , as there are products. In the vertical model, there is only one consumer characteristic: ν_i , the “taste” for quality. In the logit model, all products are strict substitutes for one another, while in the vertical model, only products that are adjacent in the quality dimension are substitutes.

It is well known that market shares in the vertical model are defined by the cutoff points:

$$\Delta_j = (p_j - p_{j-1})/(\psi_j - \psi_{j-1}), \quad 1 \leq j \leq N. \quad (17)$$

For consistency of notation, also define cutoffs $\Delta_0 = -\infty$ and $\Delta_{N+1} = \infty$. In equilibrium, Δ_j is increasing in j and consumer i purchases good j if and only if $\Delta_{j+1} > \nu_i > \Delta_j$, giving market shares of

$$s_j = F(\Delta_{j+1}) - F(\Delta_j). \quad (18)$$

We can use this market share equation to solve recursively for the cutoff points and then, from prices together with the definition of the cutoffs, can solve for the implied quality levels. Solving for Δ_j in the market share equation gives the recursive relationship

$$\Delta_j = F^{-1}(F(\Delta_{j+1}) - s_j), \quad (19)$$

with initial value $\Delta_N = F^{-1}(1 - s_N)$ (3). Given values for the price and quality of the outside good, p_o and ψ_o , the remaining values for quality can then be recursively determined from

$$\psi_j = \psi_{j-1} + (p_j - p_{j-1})/\Delta_j. \quad (20)$$

Since $\delta_j = \psi_j - p_j$, solving for the quality levels in (20) is equivalent to solving for mean utility levels. Note that (16) and (20) do not separately identify ψ_o from the mean of $x_j\beta$, so we can normalize ψ_o to zero. In a more complicated framework, the price of the outside good could be estimated, but for simplicity, I assume $p_o \equiv 0$.

If we then maintain the assumption that the ξ 's have mean zero conditional on the x 's, β can be estimated by regressing the calculated ψ_j 's on the product characteristics. If the distribution of the taste for quality depends on any unknown parameters to be estimated, then the estimation procedure must be modified, as discussed in the next section.

As noted, both the logit and the vertical differentiation model place very strong restrictions on the pattern of estimated cross-price elasticities. The following section will discuss the extension to richer models.

5. Estimating density parameters

■ In the immediately preceding discussion, I have assumed that the density of the vector of consumer characteristics is known to the researcher. While this assumption is imposed by much of the existing literature, one might prefer to assume a parametric family of

densities. In this section, I shall assume that the density of v_{ij} , $f(\cdot, \mathbf{x}, \sigma)$, depends on a vector of unknown parameters, σ , which is to be estimated.

In many cases, we may have no particular interest in σ itself but are instead concerned that a narrow distributional assumption on tastes will yield unreasonable estimates of economically interesting values, such as cross-price elasticities. For example, as in the random coefficients model, one may wish to specify a set of parameters that allow us to estimate the relationship between product characteristics and substitution patterns.

Once the distribution of consumer characteristics is parameterized to depend on density parameters σ , the market share function and the implied mean utility levels will also vary with σ . The mean utility levels are implicitly defined from the vector of equation $\mathbf{s} = \mathcal{A}(\boldsymbol{\delta}, \sigma)$. Inverting, the demand equation is

$$\delta_j(\mathbf{s}, \sigma) = x_j\beta + \alpha p_j + \xi_j. \quad (21)$$

We can continue to use an instrumental variables technique to estimate (σ, β, α) . However, note that the parameters σ will now frequently enter the estimating equation in a nonlinear fashion, so nonlinear least-squares (or generalized method of moments (Hansen, (1982)) techniques may be necessary to estimate the model parameters. Furthermore, the presence of σ increases the number of parameters to be estimated and so increases the number of required instruments.

To illustrate, I next consider two models, the nested logit and the full random coefficients model, in which the distribution of consumer tastes depends on unknown parameters to be estimated. Each of these models involves interactions between consumer and product characteristics, where the interactions are modelled as depending on a small number of parameters.

Example: nested logit. In contrast to the simple logit model, the nested logit model or “tree extreme value” model (McFadden, 1978; and Cardell, 1991) preserves the assumption that consumer tastes have an extreme value distribution but allows consumer tastes to be correlated (in a restricted fashion) across products j . This allows for more reasonable substitution patterns as compared to the simple logit model. In this section, I shall briefly review the nested logit model and show how to analytically invert the market share function.

I follow Cardell’s (1991) exposition of the nested logit, which has the advantage of using an explicit factor structure that is similar to the random coefficients model. First group the products into $G + 1$ exhaustive and mutually exclusive sets, $g = 0, 1, \dots, G$. Denote the set of products in group g as \mathcal{J}_g . The outside good, $j = 0$, is assumed to be the only member of group 0. For product $j \in \mathcal{J}_g$, assume that the utility of consumer i is

$$u_{ij} = \delta_j + \zeta_{ig} + (1 - \sigma)\epsilon_{ij}, \quad (22)$$

where, once again, $\delta_j = x_j\beta - \alpha p_j + \xi_j$ and ϵ_{ij} is an identically and independently distributed extreme value. For consumer i , the variable ζ is common to all products in group g and has a distribution function that depends on σ , with $0 \leq \sigma < 1$. Cardell shows that the distribution of ζ is the unique distribution with the property that, if ϵ is an extreme value random variable, then $[\zeta + (1 - \sigma)\epsilon]$ is also an extreme value random variable. As the parameter σ approaches one, the within group correlation of utility levels goes to one, and as σ approaches zero, the within group correlation goes to zero.

We can interpret (22) as a random coefficients model involving random coefficients ζ_{ig} only on group-specific dummy variables. That is, if d_{jg} is a dummy variable equal to one if $j \in \mathcal{J}_g$ and equal to zero otherwise, we can rewrite (22) as

$$u_{ij} = \delta_j + \sum_g [d_{jg}\zeta_{ig}] + (1 - \sigma)\epsilon_{ij},$$

which is similar to (4). Thus, the nested logit model allows us to model correlation between groups of similar products in a simple way. However, unlike the more general random coefficients model, the nested logit allows correlation patterns to depend only on groupings of products that are determined prior to estimation and not on the values of continuous variables.

If product j is in group g , the well-known formula for the market share of product j as a fraction of the total group share is

$$\bar{d}_{j/g}(\delta, \sigma) = [e^{\delta_j/(1-\sigma)}]/D_g, \tag{23}$$

where the denominator of this expression for a product in group g is

$$D_g \equiv \sum_{j \in \mathcal{J}_g} e^{\delta_j/(1-\sigma)}.$$

Similarly, the probability of choosing one of the group g products (the group share) is

$$\bar{d}_g(\delta, \sigma) = \frac{D_g^{(1-\sigma)}}{\left[\sum_g D_g^{(1-\sigma)} \right]}, \tag{24}$$

giving a market share of

$$d_j(\delta, \sigma) = \bar{d}_{j/g}(\delta, \sigma) \bar{d}_g(\delta, \sigma) = \frac{e^{\delta_j/(1-\sigma)}}{D_g^\sigma \left[\sum_g D_g^{(1-\sigma)} \right]}. \tag{25}$$

With the outside good as the only member of group zero and with $\delta_o \equiv 0$, $D_o = 1$ and so

$$d_o(\delta, \sigma) = 1 / \left[\sum_g D_g^{(1-\sigma)} \right].$$

Having set out the basic model, we can now derive a simple analytic expression for mean utility levels. Taking logs of market shares,

$$\ln(s_j) - \ln(s_o) = \delta_j/(1 - \sigma) - \sigma \ln(D_g). \tag{26}$$

This expression depends on the unknown value of D_g . Taking the log of the group share in (24), $\ln(D_g) = [\ln(\bar{s}_g) - \ln(s_o)]/(1 - \sigma)$, where the observed group share is denoted \bar{s}_g . Substituting this into (26) and combining terms gives the analytic expression for $d_j^{-1}(s, \sigma)$:⁷

$$\delta_j(s, \sigma) = \ln(s_j) - \sigma \ln(\bar{s}_{j/g}) - \ln(s_o). \tag{27}$$

This is the same as the logit equation (14), except for the additional term $\sigma \ln(\bar{s}_{j/g})$.

Setting $\delta_j = x_j\beta - \alpha p_j + \xi_j$ and substituting in from (27) for δ_j gives

$$\ln(s_j) - \ln(s_o) = x_j\beta - \alpha p_j + \sigma \ln(\bar{s}_{j/g}) + \xi_j, \tag{28}$$

so that estimates of β , α , and σ can be obtained from a linear instrumental variables regression of differences in log market shares on product characteristics, prices, and the log of the within group share. This last term is endogenous, suggesting the need for additional

⁷ This formula can be easily (but somewhat tediously) extended to the case of multiple levels of nests and different correlation parameters, σ_g , for different groups.

exogenous variables that are correlated with the within group share. These variables might include the characteristics of other firms in the group.

Note that the nested logit model is an example with nontrivial interactions between product and consumer characteristics that, on the demand side, still allows for linear estimation techniques. Because the nested logit only allows for simple patterns of correlation between products, I shall briefly return to the full random coefficients model, which allows for more complicated patterns.

Example: the full random coefficients model. In this model, the market share equation is now difficult to calculate, but the general discussion of solving for the vector δ does not substantially change. Each set of values for the σ_k , the standard deviations of the random coefficients in (4), will imply a different relationship between the observed market shares s and the utility means δ . Typically, one will have to solve for the δ 's numerically.

There remains the problem that, for a large number of products, and for arbitrary assumptions on consumer tastes, the integral defining market share in the random coefficients model may be difficult to calculate. In the context of market level data, this is effectively an aggregation problem. Pakes (1986) suggests the use of simulation methods to solve such aggregation problems, and an extension of this technique is employed in Berry, Levinsohn, and Pakes (1993). This last article also shows how to incorporate information on the empirical distribution of consumer characteristics (such as the actual distribution of income) into the random coefficients framework. Interested readers are referred to that article for details on using simulation to calculate market shares and to solve for δ_j .

Obviously, there is a tradeoff between the larger (but still feasible) computational burden of the Berry, Levinsohn, and Pakes algorithm versus the simple, but still quite restrictive, nested logit. The nested logit may be preferred when a heavy penalty is placed on computational complexity, or when a researcher wants to model substitution effects as depending only on predetermined classes of products. The random coefficients model will be preferred when a premium is placed on estimating richer patterns of demand.

6. The pricing equation and the supply side

■ Sections 4 and 5 used no information from the price-setting process. If we are willing to assume that observed prices are the result of an interior, pure strategy Nash equilibrium in prices, then we can also make use of the information contained in the first-order equations for equilibrium prices in (9). Note that, under the assumptions of Section 3, $\partial s_j / \partial p_j = -\alpha \partial s_j / \partial \delta_j$, so the first-order conditions can be rewritten as depending on $\partial s_j / \partial \delta_j$. Then, given the vector of utility levels as derived from the inverse market share function, the term $\partial s_j / \partial \delta_j$ can be obtained by simple analytic or numeric differentiation of the market share function evaluated at the appropriate value of δ . Thus, given the distribution of consumer tastes, both δ_j and $\partial s_j / \partial \delta_j$ can be treated as known transformations of the data.

The discussion of cost-side estimation is eased if we assume that marginal cost is linear in the unobservable cost term ω_j . If we make the simple assumption $c_j = \bar{c}(q_j, w_j, \gamma) + \omega_j$, then the first-order condition implies that⁸

$$p_j = \bar{c}(q_j, w_j, \gamma) + \frac{1}{\alpha} [s_j / (\partial s_j / \partial \delta_j)] + \omega_j. \quad (29)$$

Equation (29) can now be treated as an estimation equation in much the same manner as (21). The observable right-hand-side variables of (29) are the terms defining the mean marginal cost function, $\bar{c}(q_j, w_j, \gamma)$, and the markup term, $s_j / [\partial s_j / \partial \delta_j]$. The parameters to

⁸ While not necessary, the assumption that marginal cost is linear in the unobservable is useful because it produces a linear error in (29).

be estimated are the cost function parameters, γ , and the marginal disutility of price increase, α . Note that the right-hand side of (29) includes variables that are econometrically endogenous in the sense that they are correlated with ω_j ; these endogenous variables are outputs, q_j , market share, s_j , and the market share derivatives, $\partial s_j / \partial \delta_j$. Therefore, appropriate instruments must once again be found.

Excluded demand-side parameters (elements of x_j that are not included in w_j) are, as usual, available as cost-side instruments. However, it may be unreasonable to assume the existence of x variables that are valued by consumers but do not affect marginal costs. Once again, the characteristics of other firms are also available as instruments. In equilibrium, these characteristics will be correlated with own-firm output and therefore correlated with market shares and with $\partial s_j / \partial \delta_j$. Also, in a cross section or time series of markets with differing populations, population is a potential instrument for output quantities, q_j .

Finally, it is obviously possible to jointly estimate the demand and supply equations, (21) and (29). Joint estimation would take into account the cross-equation restrictions on parameters: α and the substitution parameter, σ , affect both demand and supply. Once again, an example of joint estimation is found in Berry, Levinsohn, and Pakes (1993).

Examples of first-order conditions. In this subsection, I shall discuss the supply equations that are implied by the simple special cases discussed above. To simplify the examples, I make the assumption that marginal cost is constant in output and linear in product characteristics,

$$c_j = w_j \gamma + \omega_j. \tag{30}$$

It is easy to derive the first-order conditions for the logit model. Because in this model $\partial s_j / \partial \delta_j = s_j(1 - s_j)$, the first-order condition is

$$p_j = \frac{1}{\alpha(1 - s_j)} + c_j,$$

which, given (30), implies a supply equation for the logit model of

$$p_j = w_j \gamma + \frac{1}{\alpha(1 - s_j)} + \omega_j, \tag{31}$$

where the parameters to be estimated are γ and $(1/\alpha)$. The logit joint estimation problem is then defined by this equation together with the logit demand equation from (14).

In the vertical model, which departs slightly from the model of Section 3, the first-order conditions can be derived from $\partial s_j / \partial p_j = -[f(\Delta_{j+1})/(\psi_{j+1} - \psi_j) + f(\Delta_j)/(\psi_j - \psi_{j-1})]$. Defining price minus markup as

$$y_j \equiv p_j - s_j / \left| \partial s_j / \partial p_j \right| = p_j - s_j / [f(\Delta_{j+1})/(\psi_{j+1} - \psi_j) + f(\Delta_j)/(\psi_j - \psi_{j-1})], \tag{32}$$

estimates of γ can be found from an OLS regression of y_j on w_j . Further instruments are necessary in this model only when the density, f , is assumed to depend on unknown parameters. In this case, nonlinear instrumental variable methods may be necessary.

In the nested logit model, differentiating the market share equation (25) gives

$$\partial s_j / \partial \delta_j = \frac{1}{(1 - \sigma)} s_j [1 - \sigma \bar{s}_{j/g} - (1 - \sigma)s_j].$$

The implied pricing equation is

$$p_j = w_j \gamma + \left[\frac{(1 - \sigma)}{\alpha} \right] / [1 - \sigma \bar{s}_{j/g} - (1 - \sigma)s_j] + \omega_j. \tag{33}$$

If $\sigma = 0$, it is only the product share, s_j , and not the within group share, $\bar{s}_{j/g}$, that affects

the markup. Conversely, as σ approaches one, it is only $\bar{s}_{j/g}$ that matters. Thus, the relationship in the data between prices, product shares, and group shares will help to identify the substitution parameter, σ .

Thus, in the three previously discussed cases with analytic solutions for δ_j —logit, nested logit, and vertical differentiation—there are also analytic solutions for the supply-side first-order conditions. In the full random coefficients model, however, the term $\partial a_j / \partial \delta_j$ must often be obtained by numeric differentiation of the market share function.

7. Alternative methods of estimation

■ While the method of inverting for mean utility levels is very easy for some of the outlined special cases, in other cases, the procedure may appear to be overly burdensome. In this section I shall briefly discuss two methods that may appear to be obvious, simpler solutions to this problem but which are not. I shall also compare the mean utility method to the reduced form method of estimating differentiated products models. The reduced form method, which has been used in the past, imposes a very severe computational burden and also requires more restrictive assumptions.

Perhaps the most obvious econometric approaches for dealing with the unobserved characteristic are either to estimate ξ_j as a “fixed effect” or to “integrate out” over some assumed exogenous distribution for the unobserved heterogeneity ξ . Regarding the latter suggestion, the price-setting model is inconsistent with any assumption of an exogenous distribution for ξ , conditional on \mathbf{x} and \mathbf{p} . This follows from the first-order conditions for optimal prices, which imply that different values for ξ result in different levels of prices. Thus, integrating (6) over the distribution of ξ while holding prices fixed will not give the average level of market share that would be observed as ξ varies.

Neither is it possible to separately estimate values of ξ_j together with estimates of the coefficients on x_j and p_j . Remember that mean utility is given by $\delta_j = x_j\beta + \xi_j - \alpha p_j$. Obviously, combinations of values for $(\xi_1, \dots, \xi_N, \beta, \alpha)$ that give the same values of δ_j must also yield the same predictions for consumer behavior. Therefore, the vector ξ is not identified separately from the coefficients on firm-specific characteristics and prices. This result is familiar from any analysis of grouped data: it is not possible to estimate an individual group mean together with coefficients on variables that do not vary within the group.

Another approach to estimation, which has been used by Bresnahan (1987) and Berry (1990), requires solving for the reduced form of the model. Suppose we have established the existence of a unique equilibrium and are willing to assume the existence of a family of probability measures, $\Phi(\cdot/\mathbf{z}, \theta_u)$, for the random variables (ξ, ω) . A nonlinear least-squares (or method of moments) estimator can then be based on the difference between the observed price and the mean of the reduced-form price.

To obtain this estimator, write the expected values (conditional on product characteristics) of equilibrium prices and quantities as

$$\bar{p}(\mathbf{z}, \theta, \theta_u) = \int p(\mathbf{z}, \xi, \omega, \theta) \Phi(d(\xi, \omega)/\mathbf{z}, \theta_u)$$

and

$$\bar{q}(\mathbf{z}, \theta, \theta_u) = \int q(\mathbf{z}, \xi, \omega, \theta) \Phi(d(\xi, \omega)/\mathbf{z}, \theta_u), \quad (34)$$

where, once again, $\bar{p}(\mathbf{z}, \boldsymbol{\xi}, \boldsymbol{\omega}, \theta)$ is the reduced form function defining price and $\bar{q}(\mathbf{z}, \boldsymbol{\xi}, \boldsymbol{\omega}, \theta)$ is the reduced form for quantities. We can use these equations to rewrite the model as

$$\mathbf{p} = \bar{p}(\mathbf{z}, \theta^*, \theta_a^*) + \mathbf{v}$$

and

$$\mathbf{q} = \bar{q}(\mathbf{z}, \theta^*, \theta_a^*) + \mathbf{e}, \tag{35}$$

where the “prediction errors” (\mathbf{e}, \mathbf{v}) are, by construction, mean zero conditional on the observed firm characteristics \mathbf{z} . Thus, (35) can be used as the basis for a traditional non-linear least-squares estimator of the model parameters.

The reduced form method is linear in observed prices and quantities, which allows us to easily incorporate measurement error in prices and quantities. Indeed, Bresnahan (1987) models measurement error, instead of unobserved product characteristics, and adopts a reduced form approach to estimation.

However, there are several problems with the reduced form approach. The expected values in (34) are defined as integrals over implicitly defined functions. These are typically very difficult to calculate, especially because a nonlinear estimation routine will need to evaluate this function at many possible parameter values. Also, the integrals defining these expected values implicitly depend on the existence of a unique equilibrium for all observed values of \mathbf{x} and for almost all values of the unobservables ($\boldsymbol{\xi}, \boldsymbol{\omega}$). Thus, existence of a unique equilibrium at a particular set of values for ($\boldsymbol{\xi}, \boldsymbol{\omega}$) is not sufficient. As noted by Caplin and Nalebuff (1991), it is very difficult to establish uniqueness in this class of models. Interestingly, Caplin and Nalebuff establish uniqueness of equilibrium in the special cases of the logit and vertical differentiation model. However, we have seen that the mean utility method is easy in these cases, especially when compared to solving the integrals in (34).

8. Monte Carlo experiments

■ In the introduction, I noted that Trajtenberg (1989) provided an example of empirical work in which unobserved characteristics appear to have a dramatic effect on some real-world parameter estimates. In order to provide a simple example of how the methods of this article can correct for such a bias, Table 1 supplies Monte Carlo results for estimation

TABLE 1 Monte Carlo Parameter Estimates 100 Random Samples of 500 Duopoly Markets Logit Utility

Parameter	True Value	$(\sigma_{\xi d} = 1)$		$(\sigma_{\xi d} = 3)$	
		(1) OLS	(2) IV	(3) OLS	(4) IV
β_o	5	3.46 (.158)	4.98 (.226)	0.378 (.415)	4.89 (.738)
β_x	2	1.41 (.058)	1.99 (.091)	.325 (.127)	1.95 (.272)
α	1	.726 (.029)	.995 (.039)	.181 (.076)	.979 (.128)

Notes: The values given in the table are empirical means and (standard errors).

The utility function is $u_{ij} = \beta_o + \beta_x x_j + \sigma_{\xi d} \xi_j - \alpha p_j + \epsilon_{ij}$.

Marginal cost is $c_j = e^{\gamma_o + \gamma_x x_j + \sigma_{\xi d} \xi_j + \gamma_{\omega} \omega_j + \sigma_{\omega} \omega_j}$.

of the logit model with an unobserved characteristic. The data for these experiments were created as follows. Each simulated sample consists of 500 duopoly markets (a reasonable number of markets for the airline examples of Morrison and Winston (1986) and Berry (1990)).

With a slight abuse of the notation for ξ and ω , the utility of each consumer in each market is given by $u_{ij} = \beta_o + \beta_x x_j + \sigma_d \xi_j - \alpha p_j + \epsilon_{ij}$, with ϵ_{ij} being the appropriate logit error. The utility of the outside good is given by $u_{io} = \epsilon_{io}$, where the ϵ_{io} has the same distribution as the other ϵ_{ij} . Marginal cost is constrained to be positive and is given by $c_j = e^{\gamma_o + \gamma_x x_j + \sigma_c \xi_j + \gamma_w w_j + \sigma_w \omega_j}$.

The exogenous data x_j , ξ_j , w_j , and ω_j are all created as independent standard normal random variables. The term ξ is here a product characteristic that affects both demand and cost, while ω is some variable (such as an input price) that affects only costs. Note that β_o , β_x , α , γ_o , γ_x , and γ_w are parameters to be estimated, whereas the parameters σ_d , σ_c , and σ_w help to describe the effect of the unobservables ξ and ω . Values for the parameters were chosen by ad hoc experimentation to yield a moderate variance in market shares and prices across markets, without driving market shares of the duopolists toward zero in too many markets. The chosen values are $\beta_o = 5$, $\beta_x = 2$, $\alpha = 1$, $\gamma_o = 1$, $\gamma_x = .5$, and $\gamma_w = \sigma_w = \sigma_c = .25$. Columns 1 and 2 of Table 1 present Monte Carlo estimation results from samples of markets, with the standard deviation of the unobserved characteristics in the utility function set to $\sigma_d = 1$. Since the coefficient on x_j is set to 2, the total variance in the implied mean, δ_j , is 5, 80% of which is accounted for by the observed term x_j . In contrast, columns 2–4 present results for samples of markets with $\sigma_d = 3$, so that the variance of δ_j is 13, almost 70% of which is accounted for by the unobserved term $\sigma_d \xi_j$.

For each market, I first calculate the equilibrium values of prices and market shares. I assume that the hypothetical econometrician observes these data (including the market share of the outside good, s_o) along with x and w . The terms ξ and ω are, as usual, unobserved by the econometrician. As in (14), the mean utility level of good j can be found as $\delta_j = \ln(s_j) - \ln(s_o)$.

Two estimation methods for the demand parameters are presented. In the first method, δ_j is regressed on x_j and p_j without regard to the endogeneity of prices. These results, in columns 1 and 3, are comparable to those differentiated product studies that do not consider unobserved characteristics. In the second method, the observed cost factors, w_j , and the demand characteristic of the rival firm are used as instruments for price. These results are in columns 2 and 4. We see that, even when the observed characteristic accounts for 80% of the variance in mean utility levels, the coefficient on price is systematically underestimated by OLS. Simple calculations show that, for many samples, the OLS estimate of α implies that firms are pricing on an inelastic portion of their demand curves, thus falsely appearing to reject relevant economic theory. In column 3, where the observed characteristic accounts for only 25% of the variation in mean utility levels, the OLS estimates sometimes indicate that consumers prefer to pay higher prices (i.e., $-\hat{\alpha} > 0$). The instrumental variable method, in contrast, provides reasonable estimates of the coefficients, thus correcting for the bias in the OLS estimates.

9. Extensions

■ This article leaves many estimation issues yet to be explored. Much of this exploration will be most fruitful in the context of a particular industry study. Issues that might be examined include questions of how to estimate market size, M , when this is not directly observed and how to make optimal use of potential instruments, such as the characteristics of other firms. The first question is taken up in Berry (1990) and Greenstein (1992), while approximations to the optimal instruments are developed in Berry, Levinsohn, and Pakes (1993). In the remainder of this section, I shall briefly discuss some additional extensions.

□ **Consumer data.** Researchers increasingly have access to data on individual consumer decisions. In this case, we might parameterize utility as $u_{ij} = \delta_j + \mathbf{y}_i\beta_j + \epsilon_{ij}$, where \mathbf{y}_i is a vector of observed consumer characteristics. Individual consumer data could be used to estimate the product-specific means δ_j . Call these estimates $\hat{\delta}_j$. The estimated $\hat{\delta}_j$ could then be treated in much the same way as the δ_j derived from aggregate data on market shares (although additional complications now arise from the estimation error in $\hat{\delta}_j$). That is, the $\hat{\delta}_j$ could be “regressed” (using instrumental variables techniques to account for the endogeneity of prices) on product characteristics and prices. This procedure is analogous to techniques that are familiar from the empirical literature on linear grouped data models, in which estimates of individual group means might be explained via a regression on group characteristics. The nonlinear nature of the discrete-choice market-share function prevents us from using the obvious alternative in linear models, which is to include the group-specific data directly in a linear regression equation.

□ **Different specifications for utility.** It would be useful to extend the methodology above to incorporate yet more general models of consumer utility and firm behavior. As long as they incorporate unobserved product characteristics, such models will likely continue to face a nonlinear instrumental variables problem of the sort discussed above. The method of this article suggests solving backward from observed data to uncover the product specific unobservables, ξ and ω ; this method may also be useful in more general specifications. To extend the methodology on the demand side, it is necessary to prove a result similar to that found in the Appendix: namely, that given parameters, each vector of observed data can be explained by only one vector of product-specific unobservable demand characteristics. Similarly, on the pricing side, it will be necessary to provide an analog to (29), which demonstrates that the data and parameters together uniquely determine the cost-side unobservable.

□ **Measurement error.** Measurement error in observed prices, characteristics, or quantities may also create difficulties for the estimation procedure outlined above. Because prices are already treated as endogenous variables, measurement error in prices may not be a serious problem (as long as the only effect of price is as a linear term in δ_j .) However, measurement error in output quantities presents a more serious problem. The nonlinear inversion of market shares to uncover δ_j may be quite sensitive to measurement error in observed market shares. As noted, the reduced-form method is not sensitive to measurement error in the left-hand-side variables, price and quantity, and thus, this method may be preferable (when feasible) in the presence of mismeasured quantities.

□ **Endogenous product characteristics.** The estimation techniques of this article rely on the traditional assumption that the unobserved product-level errors are uncorrelated with observed product characteristics. Given that firms choose the characteristics of their products, this assumption may be unreasonable. However, a solution to the problem of “endogenous x ’s” requires a reasonable model of the dynamic process that generates product characteristics. This project goes well beyond the static framework of the current article. Similarly, there are other circumstances that call the static demand model into question, such as the modelling of consumer dynamics and durable goods.

In one useful advance, Pakes and McGuire (1991) show that simple discrete-choice product differentiation models have useful properties when employed as models of single-period profits in dynamic models of equilibrium firm behavior, such as that of Ericson and Pakes (1989). Combining the endogenous pricing models of this article with dynamic models of investment in product quality would allow for the endogeneity of both market characteristics and market prices. I shall leave this as a topic for future research.

10. Conclusions

■ In this article, I have considered methods for estimating product differentiation models in the presence of unobserved product characteristics. While homogeneous goods models are almost never estimated while ignoring the correlation of prices and demand errors, it has been commonplace to ignore this correlation in more complex studies of differentiated products markets with discrete-choice demand models. I suggest “inverting” the discrete-choice market-share function to find implied levels of mean utility. These mean utility levels can then be treated in much the same fashion as observed output quantities in the homogeneous goods model. For some leading special cases, it is quite easy to invert the market-share function. More complicated models impose a greater computational burden, but this burden may still be less than what is required by alternative estimation methods, such as solving for the reduced form.

The worth of the methods suggested here must ultimately be established in empirical applications. Some early success can be reported in this regard. Berry, Levinsohn, and Pakes (1993) extend the methods of this article in several directions in order to estimate the parameters of an equilibrium model of differentiated products supply and demand in the automobile industry. Consistent with the Monte Carlo results reported here, they show that allowing for unobserved product characteristics, which are correlated with prices, improves estimates of own-price elasticities. They also extend the random coefficients framework of this article and obtain plausible estimates of product-level cross-price elasticities. Greenstein (1992) has also reported plausible demand estimates in a study of the computer industry that employs a vertical differentiation model with unobserved product characteristics. Thus, while much work remains to be done, there are potentially useful empirical applications for the methods presented here.

I should emphasize in closing that the techniques of this article rely on a number of restrictive assumptions. These include assumptions that demand is well approximated by a static discrete-choice model and that the distribution of consumer tastes is known up to a parameter vector. More importantly, and more difficult to solve, I assume that product characteristics are econometrically exogenous. A solution to this last problem awaits further progress on dynamic models of firm behavior.

Appendix

■ **The inverse of the market-share equation.** As in Section 3, consider the utility function $u_{ij} = \delta_j + v_{ij}$. Holding $\delta_o = 0$, I shall prove that a unique vector $\delta = \delta^{-1}(s)$ exists. Assume that the market share function, $d(\delta)$, has the following properties, which are sufficient, but not necessary, for the results which follow: d is everywhere differentiable with respect to δ , and its derivatives obey the following strict inequalities— $\partial d_j / \partial \delta_j > 0$ and $\partial d_j / \partial \delta_k < 0$, $k \neq j$. A sufficient (but not necessary) condition for these properties is that for all possible values of \mathbf{x} , the density of consumer characteristics, $f(\mathbf{v}, \mathbf{x})$, is strictly positive and continuous for all $\mathbf{v} \in \mathbb{R}^{N+1}$. Also note that for any finite values of $\{\delta_k, k \neq j\}$, d_j approaches arbitrarily close to zero as δ_j goes to $-\infty$, while d_j approaches arbitrarily close to one as δ_j goes to ∞ .

I begin by defining the element-by-element inverse, $r_j(\delta, s_j)$. This function is defined as the value for the mean utility of the j th product such that the predicted value d_j exactly equals the observed value s_j . That is, r_j is implicitly defined as

$$s_j = d_j(\delta_1, \delta_2, \dots, r_j(\delta, s_j), \dots, \delta_N). \quad (\text{A1})$$

By the assumptions on the market-share function, this element-by-element inverse exists and is continuous and differentiable. Note that r_j is strictly increasing in δ_k and does not depend on δ_j . Also define the vector valued function $\mathbf{r} = (r_1, \dots, r_j)$.

The element-by-element inverse allows us to transform the problem of solving for the vector inverse into a fixed-point problem, for a vector δ satisfies $d(\delta) = \mathbf{s}$ if and only if $\delta = \mathbf{r}(\delta, \mathbf{s})$. The method of proof is to use a slight variant of Brouwer’s fixed-point theorem to prove existence of a fixed point of the element-by-element inverse. It is then necessary to show that there cannot be two such fixed points.

To establish existence, first hold $\delta_o = 0$ and note that $r_j(\delta, s_j)$ has a lower bound. This lower bound is $r_j(\delta', s_j)$, with δ' set equal to any vector in \mathbb{R}^{N+1} such that $\delta_k = -\infty$ for $k \neq (j, 0)$. Define $\underline{\delta}$ as the smallest

value across products of these lower bounds. There is no upper bound for r_j , but the following lemma allows one to establish existence in the absence of an upper bound.

Lemma. There is a value $\bar{\delta}$, with the property that if one element of $\mathbf{\delta}$, say δ_j , is greater than $\bar{\delta}$, then there is a product index k such that $r_k(\mathbf{\delta}, \mathbf{s}) < \delta_k$.

Proof. To construct $\bar{\delta}$, again set $\delta_k = -\infty, \forall k \neq (j, 0)$. Then define $\bar{\delta}_j$ as the value of δ_j that sets the market share function for the outside good, s_o , equal to the observed share s_o . Define $\bar{\delta}$ as any value greater than the maximum of the $\bar{\delta}_j$. Now, if for the vector $\mathbf{\delta}$ there is an element j such that $\delta_j > \bar{\delta}$, then $s_j(\mathbf{\delta}) < s_o$, which implies $\sum_{j=1}^N s_j(\mathbf{\delta}) > \sum_{j=1}^N s_j$, so there is at least one element k with $s_k(\mathbf{\delta}) > s_k$. For this k , $r_k(\mathbf{\delta}, \mathbf{s}) < \delta_k$.

Now define a new function which is a truncated version of r_j : $\bar{r}_j(\mathbf{\delta}, \mathbf{s}) = \min \{r_j(\mathbf{\delta}, \mathbf{s}), \bar{\delta}\}$. Clearly, $\bar{\mathbf{r}}(\mathbf{\delta}, \mathbf{s})$ is a continuous function which maps $[\underline{\delta}, \bar{\delta}]^N$ into itself, so, by Brouwer's fixed-point theorem, $\bar{\mathbf{r}}(\mathbf{\delta}, \mathbf{s})$ has a fixed point, $\mathbf{\delta}^*$. By the definition of $\underline{\delta}$ and $\bar{\delta}$, $\mathbf{\delta}^*$ cannot have a value at the upper bound, so $\mathbf{\delta}^*$ is in the interior of $[\underline{\delta}, \bar{\delta}]^N$. This implies that $\mathbf{\delta}^*$ is also a fixed point of the unrestricted function $\mathbf{r}(\mathbf{\delta}, \mathbf{s})$, which establishes existence.

A well-known sufficient condition for uniqueness is $\sum_k |\partial r_j / \partial \delta_k| < 1$. By the implicit function theorem, $\partial r_j / \partial \delta_k = -[\partial s_j / \partial \delta_k] / [\partial s_j / \partial \delta_j]$. From this, $\sum_k |\partial r_j / \partial \delta_k| < 1$ if and only if a dominant diagonal condition holds:

$$\sum_{k \neq (j,0)}^N \left| \frac{\partial s_j}{\partial \delta_k} \right| < \frac{\partial s_j}{\partial \delta_j}. \quad (\text{A2})$$

To establish this condition, note that increasing all the mean utility levels (including δ_o) by the same amount will not change any market share. Then, (A2) follows from

$$\sum_{k=0}^N \frac{\partial s_j}{\partial \delta_k} = 0 \Rightarrow \sum_{k=1}^N \frac{\partial s_j}{\partial \delta_k} = -\frac{\partial s_j}{\partial \delta_o} > 0.$$

Q.E.D.

References

- ANDERSON, S.P., DEPALMA, A., AND THISSE, J.-F. "Demand for Differentiated Products, Discrete Choice Models, and the Characteristics Approach." *Review of Economic Studies*, Vol. 56 (1989), pp. 21–35.
- BERRY, S.T. "Airport Presence as Product Differentiation." *American Economic Review*, Vol. 80 (1990), pp. 394–399.
- , LEVINSOHN, J., AND PAKES, A. "Automobile Prices in Market Equilibrium." NBER Working Paper no. 4264, 1993.
- BRESNAHAN, T.F. "Competition and Collusion in the American Auto Industry: The 1955 Price War." *Journal of Industrial Economics*, Vol. 35 (1987), pp. 457–482.
- CAPLIN, A. and NALEBUFF, B. "Aggregation and Imperfect Competition: On the Existence of Equilibrium." *Econometrica*, Vol. 59 (1991), pp. 25–60.
- CARDELL, N.S. "Variance Components Structures for the Extreme Value and Logistic Distributions." Mimeo, Washington State University, 1991.
- EPPLE, D. "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products." *Journal of Political Economy*, Vol. 95 (1987), pp. 59–80.
- ERICSON, R. AND PAKES, A. "An Alternative Theory of Firm and Industry Dynamics." Columbia University Discussion Paper no. 445, 1989.
- GREENSTEIN, S. "From Super-minis to Super-computers: Estimating Surplus in the Computing Market." Mimeo, University of Illinois, 1992.
- GRILICHES, Z., ed. *Price Indexes and Quality Change*. Cambridge, Mass.: Harvard University Press, 1971.
- HANSEN, L.P. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*, Vol. 50 (1982), pp. 1029–1054.
- HAUSMAN, J.A. AND WISE, D.A. "A Conditional Probit Model for Qualitative Choice: Discrete Decisions Recognizing Interdependence and Heterogeneous Preferences." *Econometrica*, Vol. 46 (1978), pp. 403–426.
- McFADDEN, D. "Conditional Logit Analysis of Qualitative Choice Behavior." In P. Zarembka, ed., *Frontiers in Econometrics*. New York: Academic Press, 1974.
- . "Modelling the Choice of Residential Location." In A. Karlqvist, et al., eds., *Spatial Interaction Theory and Planning Models*. Amsterdam: North-Holland, 1978.
- . "Econometric Models of Probabilistic Choice." In C.F. Manski and D. McFadden, eds., *Structural Analysis of Discrete Data with Econometric Applications*. Cambridge, Mass.: MIT Press, 1981.
- MORRISON, S.A. AND WINSTON, C. *The Economic Effects of Deregulation*. Washington, D.C.: The Brookings Institution, 1986.
- PAKES, A. "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks." *Econometrica*, Vol. 54 (1986), pp. 755–784.

- and MCGUIRE, P. "Computation of Markov Perfect Equilibria: A Differentiated Products Version of the Ericson-Pakes Model." Mimeo, Yale University, 1991.
- PERLOFF, J.M. AND SALOP, S.C. "Equilibrium with Product Differentiation." *Review of Economic Studies*, Vol. 52 (1985), pp. 107–120.
- ROSEN, S. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy*, Vol. 82 (1974), pp. 34–55.
- SATTINGER, M. "Value of an Additional Firm in Monopolistic Competition." *Review of Economic Studies*, Vol. 51 (1984), pp. 321–332.
- SHAKED, A. AND SUTTON, J. "Relaxing Price Competition Through Product Differentiation." *Review of Economic Studies*, Vol. 49 (1982), pp. 3–13.
- TRAJTENBERG, M. "The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners." *Journal of Political Economy*, Vol. 97 (1989), pp. 444–479.
- TVERSKY, A. "Elimination-by-Aspects: A Theory of Choice." *Psychological Review*, Vol. 79 (1972), pp. 281–299.