

MARKET POWER AND JOINT DOMINANCE IN U.K. BREWING*

MARGARET E. SLADE†

Market power and joint dominance are examined in U.K. brewing. I assess *unilateral* and *coordinated* effects, where the latter is equated with joint dominance, and show how one can distinguish between the two econometrically. The application makes use of two demand equations: the nested logit of McFadden [1978a] and the distance-metric of Pinkse, Slade, and Brett [2002]. The two equations yield very different predictions concerning elasticities and markups. Nevertheless, although there is evidence of market power using either demand model, that power is due entirely to unilateral effects. In other words, neither model uncovers evidence of coordinated effects (tacit collusion).

I. INTRODUCTION

INDUSTRIAL ECONOMISTS ARE FREQUENTLY ASKED to assess the extent of market power that firms in an industry possess. For example, when two or more firms propose a merger, competition authorities must decide if the firms possess market power and if the merger will lead to an increase in that power. Furthermore, it is becoming increasingly common to supplement traditional competition policy analysis with econometric evaluations of industry performance. European Union (EU) and North American authorities, however, take different approaches to the analysis of competitive harm.

To illustrate, in North America, merger policy tends to be based on the notion of *unilateral effects*. In other words, authorities attempt to determine if firms in an industry have market power and how a merger will affect that power, assuming that the firms act in an uncoordinated fashion (see, e.g.,

* This research was supported by a grant from the Social Sciences and Humanities Research Council of Canada. I would like to thank Joris Pinkse for collaboration on related papers and for suggestions for this one. I would also like to thank the following people for helpful comments on earlier versions of the paper: Steven Berry, David Genesove, Lars Mathiesen, Mark McCabe, Marc Rysman, Frank Verboven, an anonymous referee, and participants at seminars at the University of British Columbia, the Helsinki School of Economics, the University of Chicago GSB, the London School of Economics, the London Business School, CREST, the EARIE meetings in Madrid, an NBER workshop on Industrial Organization, the International IO Conference in Boston, Frontier Economics, and the U.K. Office of Fair Trading.

† Author's affiliation: Department of Economics, University of Warwick, Coventry CV4 7AL, U.K.

e-mail: m.slade@warwick.ac.uk

Shapiro [1996]). In practice, this change is often evaluated as a move from one static Nash equilibrium to a second equilibrium with fewer players.¹

European authorities, in contrast, tend to base their policy on the notion of dominance. In other words, they seek to determine if a single firm or group of firms occupies a dominant position and if the merger will strengthen that position. Traditionally, single-firm dominance was emphasized. However, the notion of joint dominance has assumed increasing importance due to high-profile merger cases such as Nestle/Perrier, Gencor/Lonrho, and Airtours. Joint dominance is usually taken to mean tacit collusion or *coordinated effects*.² In contrast to the evaluation of unilateral effects, where econometric evidence has been introduced in court, the econometric evaluation of coordinated effects has received little attention.

In this paper, I propose an econometric technique that can be used to evaluate price/cost margins and to decompose those margins into unilateral and coordinated effects. The technique is based on the decomposition that is developed in Nevo [2001]. The organization of the paper is as follows.

The next section compares the E.U. and North American approaches to evaluating abuse-of-dominance and merger cases. In particular, it discusses how rules of thumb based on unilateral and coordinated effects can lead to different choices of antitrust cases to investigate. In addition, if a case is investigated, the two approaches can lead to different conclusions concerning competitive harm.

Section 3 describes the U.K. brewing industry, which is characterized by moderately high margins (approximately 30%), a relatively large number of producers (about 60), a much larger number of brands (many hundreds), and moderate to high horizontal concentration (Hirshman/Herfindahl index, approximately 1,800). Furthermore, in recent years, both the structure of the industry and consumers' demand for product characteristics have witnessed dramatic changes.

Section 4 discusses the demand side of the market. I use two classes of simple models in the analysis. Simple models are emphasized because time is an important factor in competition policy cases. Indeed, since substantial efficiencies can accompany changes in market structure and industry practices, authorities would like to decide cases in a matter of months, not years. In addition, if econometric evidence will be presented in court, it must be transparent, easy to interpret, and reproducible using standard software packages. The first and more familiar class of demand model, which includes the logit and nested logit of McFadden [1974 and 1978a], has been used

¹ See, e.g., Werden and Froeb [1994]; Hausman, Leonard and Zona [1994]; Nevo [2000]; Ivaldi and Verboven [2000]; Jayaratne and Shapiro [2000]; and Pinkse and Slade [2002].

² See, e.g., Lexicon [1999a], Kuhn [2000], Kuhn and Motta [2001], and Compte, Jenny and Rey [2002].

extensively by economists to evaluate mergers.³ With the logit class, own and cross-price elasticities depend on a brand's market and submarket shares. The second class, the distance-metric, is developed in Pinkse, Slade, and Brett [2002] and Pinkse and Slade [2002]. With a distance metric, own and cross-price elasticities depend on brand characteristics and a set of measures of the distance between those characteristics.

Section 5 describes how the estimated price/cost margins, which are summary statistics for the degree of market power that the firms possess, can be decomposed into unilateral and coordinated effects. Furthermore, since competition authorities have little control over product characteristics but can influence market structure, the unilateral effect is further decomposed into a portion that is due to differentiation or differences in product attributes and one that is due to concentration or multiproduct production. This decomposition is accomplished by considering pricing games that involve different ownership patterns.

Section 6 deals with estimation, section 7 describes the data, section 8 presents the empirical results, and section 9 concludes. To anticipate, I find that, although firms in the industry have substantial market power, that power is due entirely to unilateral effects. In particular, there is no evidence of a dominant group. In addition, I find that whereas brand characteristics appear to be important determinants of margins, firm characteristics do not.

II. E.U. AND NORTH AMERICAN APPROACHES

Competition authorities are responsible for policing violations of many sorts of antitrust laws. I limit attention, however, to abuse-of-dominance and merger cases. With those cases, the common goal of E.U. and North American authorities is to evaluate market power and how certain practices or acts contribute to that power. The approaches that they take to pursuing their common goal, however, differ.

With abuse-of-dominance cases, authorities on both sides of the Atlantic must determine whether a single firm or group of firms occupies a dominant position before they can determine whether that firm or group has abused its position. When a single firm is involved, although market share alone does not determine dominance, a firm's share must typically exceed 40% before it is considered dominant.⁴ With joint dominance, in contrast, often no single firm has such a large market share, but some group of firms has a joint share that is large (perhaps in excess of 60%). It is clear, however, that regardless of industry concentration, *some* group of firms will have a large share of *any* market. For this reason, it must be demonstrated that the group behaves in a

³ See, e.g., Werden and Froeb [1994], Werden [1999], Jayaratne and Shapiro [2000], and Ivaldi and Verboven [2000].

⁴ These numbers are not exact but are only indicative of common practice.

coordinated fashion to control the market. In other words, unilateral effects are not sufficient to establish joint dominance, and coordinated effects assume primary importance. It is therefore useful to have econometric techniques that can be used to distinguish between the two.

With merger cases, authorities must have some screening process to determine which cases will be investigated. Typically this involves assessing whether market power exists before determining whether the merger will increase that power. In performing that exercise, North American authorities have traditionally relied heavily on measures of industry concentration such as the Hirschman/Herfindahl index (HHI), whereas E.U. authorities, as well as many national competition bureaus within the E.U., have relied more heavily on the notion of dominance. Unfortunately, these two approaches can lead to different choices of mergers to investigate. For example, in the application below, based on the HHI and average price/cost margins, the U.K. brewing industry is fairly concentrated and firms in the industry have substantial market power. I find, however, that that power is due entirely to unilateral effects. This means that North American guidelines would dictate that mergers in the industry be closely scrutinized whereas European guidelines would not.

The factors that are considered in screening mergers are indicative of the different approaches. When unilateral effects are the focus, market shares and entry barriers assume primary importance. When coordinated effects are the focus, in contrast, a number of additional factors is considered, including product homogeneity, stable and symmetric market shares, stagnant and inelastic demand, similar costs and low levels of technical change. Those factors tend to discount mergers in consumer-product and high-tech industries.

If authorities decide to investigate a merger, they must perform a more formal assessment of whether that merger will increase market power. It is relatively straightforward to update concentration indices to reflect post-merger shares or to perform merger simulations based on unilateral effects.⁵ However, the predictions of game-theoretic models of how changes in market structure affect incentives to collude are very fragile. This means that, whereas the notion of unilateral effects is operational, the notion of joint dominance is not.

In this paper, I develop techniques that can be used for premerger evaluation or for establishing joint dominance in abuse-of-dominance cases. The more difficult task of comparing post-merger predictions of industry

⁵ When products are differentiated, market shares are less informative predictors of market power. For this reason, it has become customary to supplement the process of defining a market and calculating shares within that market with merger simulations. Those simulations assess how margins vary as the number of players in a static game changes. In other words, they assess unilateral effects.

performance across approaches is left to the future. In particular, there is a need for more robust theoretical models of joint dominance before the predictions of those theories can be quantified.

III. THE U.K. BREWING INDUSTRY

The U.K. brewing industry is interesting for a number of reasons. In particular, it has recently undergone rapid change with respect to consumer tastes, product offerings and market structure. In addition, both its horizontal and vertical organization have been subjected to numerous reviews by several levels of government.⁶

Historically, the U.K. brewing industry developed in a very different fashion from those in, for example, the U.S., Canada, and France, which were dominated by a few large brewers that sold rather homogeneous national brands of lagers. Indeed, the U.K. industry, which was relatively unconcentrated, produced a large variety of ales, and regional variation in product offerings was substantial. Moreover, national advertising played a less important role. In the last decade, however, there has been a succession of mergers that have increased concentration in brewing and have caused the industry to move towards a more North American style. Nevertheless, U.K. brewing is still less concentrated than its counterparts in the U.S., Canada, and France, where beer tends to be mass produced. It is substantially more concentrated, however, than its counterpart in Germany, where specialty beers predominate.

Substantial changes in both consumption and production have occurred in the industry in the last few decades. To illustrate, beers can be divided into three broad categories: ales, stouts, and lagers. Although U.K. consumers traditionally preferred ales to lagers, the consumption of lager has increased at a rapid pace. Indeed, from less than 1% of the market in 1960, lager became the dominant drink in 1990, when it began to sell more than ale and stout combined.⁷

A second important aspect of beer consumption is the popularity of 'real' or cask-conditioned ale. Real products are alive and undergo a second fermentation in the cask, whereas keg products are sterilized. Although real products' share of the ale market has increased, as a percentage of the total beer market, which includes lager, they have lost ground.

A final trend in consumption is the rise in popularity of premium beers, which are defined as brands with alcohol contents in excess of 4.2%. Traditional ales are of lower strength than stouts and lagers. In addition, keg

⁶ There have been about 30 reviews since 1969.

⁷ Most U.K. lagers bear the names of familiar non-British beers such as Budweiser, Fosters, and Stella Artois. Almost all, however, are brewed under license in the U.K. and are therefore not considered to be imports.

products tend to contain less alcohol than real products. Many of the more recently introduced brands, however, particularly the lagers and hybrid ales,⁸ are premium beers with relatively high alcohol contents.

With respect to production, the number of brewers has declined steadily. Indeed, in 1900, there were nearly 1,500 brewery companies, but this number fell dramatically and is currently around sixty. In addition to incorporated brewers, however, there are approximately 200 microbreweries operating at very small scales. In spite of increases in industry concentration, most brewers are still small, and few produce products that account for more than 0.5% of local markets.

For many decades, there were six national brewers that jointly accounted for about 75% of sales. If a dominant group existed, therefore, the nationals were the natural members of that group. Furthermore, in the early 1990s, mergers reduced the number of nationals from six to four, and the increase in horizontal concentration that accompanied those mergers made the possibility of joint dominance more likely.

This snapshot of the U.K. beer industry shows significant changes in tastes and consumption habits as well as a decline in the number of firms that cater to those tastes. Nevertheless, there is still considerable variety in brand offerings and brand characteristics. Brewer market power could therefore result from fewness, differentiation, collusion, or from a combination of the three. To disentangle those effects, we turn to the econometric model.

IV. DEMAND MODELS

Firms can possess market power because they have few competitors and thus operate in concentrated markets. Even when there are many producers of similar items, however, they can possess market power if their products have unique features that cause rival products to be poor substitutes. To evaluate power in markets where products are differentiated, it is therefore important to have good estimates of substitutability.

When a product is homogeneous, a single price prevails in the market. There is therefore just one price elasticity of demand—the own-price elasticity—to estimate. When products are differentiated, in contrast, the number of brands can be very large, often several hundred, and the number of price elasticities is formidable. One must therefore place some structure on the estimation.

A number of demand specifications have been developed recently to deal with the problem of an abundance of elasticities. I use two relatively simple ones here: the nested multinomial logit (NML) and the distance metric

⁸ A hybrid is a keg ale that uses a nitrogen and carbon-dioxide mix in dispensing that causes it to be smoother and to more closely resemble a cask ale.

(DM).⁹ Since the first is well known and the second is developed elsewhere (Pinkse, Slade, and Brett, 2002 and Pinkse and Slade, 2002), I simply reproduce the estimating equations and the equations for the own and cross-price elasticities of demand.

IV(i). *The Nested Logit*

The NML demand equation is based on the random-utility model in which there are n brands of a differentiated product, $q = (q_1, \dots, q_n)^T$ as well as an outside good q_0 , and an individual consumes one unit of the product that yields the highest utility. The NML is distinguished from the ordinary logit by the fact that the n brands are partitioned into G groups, indexed by $g = 1, \dots, G$, and the outside good is group 0. The partition is chosen so that like products are in the same group. For example, when the differentiated product is beer, the groups might be lager, ale, and stout. The NML estimating equation is

$$(1) \quad \ln(s_i) - \ln(s_0) = \beta^T x_i - \alpha_i p_i + \sigma_g \ln(\bar{s}_{i/g}) + \zeta_i,$$

where s_i is product i 's volume share of the entire market, x_i is a vector of observed characteristics of that product, p_i is its price, $\bar{s}_{i/g}$ is its share of the group g to which it belongs, and ζ_i is an unobserved (by the econometrician) product characteristic. The parameter $\sigma_g (0 \leq \sigma_g \leq 1)$ measures within-group- g correlation of utility, and the ordinary logit is obtained by setting σ_g equal to zero for all g .¹⁰

Let η_{ij} denote the price elasticity of demand, $(\partial q_i / \partial p_j)(p_j / q_i)$. The NML elasticities are then

$$(2) \quad \eta_{ii} = \alpha_i p_i \left[s_i - \frac{1}{1 - \sigma_g} + \frac{\sigma_g}{1 - \sigma_g} \bar{s}_{i/g} \right],$$

$$\eta_{ij} = \begin{cases} \alpha_j p_j [s_j + \frac{\sigma_g}{1 - \sigma_g} \bar{s}_{j/g}] & \text{if } j \neq i \text{ and } j \in g \\ \alpha_j p_j s_j & \text{if } j \neq i \text{ and } j \notin g. \end{cases}$$

Equation (1) is more flexible than the standard NML. In particular, the coefficient of p_i , α_i , is allowed to depend on the characteristics of that product.¹¹ In other words, $\alpha_i = \alpha(x_i)$. Furthermore, within-group correlation of utilities, σ_g , is allowed to depend on the group.¹² Nevertheless, as

⁹ A third demand specification has been used to model substitution patterns in antitrust cases. This class involves multi-stage budgeting in which the final stage (brand choice) is modeled using a flexible functional form (e.g., an Almost-Ideal-Demand System). See, Hausman, Leonard and Zona [1994]. A flexible functional form, however, is empirically intractable when there is a large number of brands in any group, as is the case here.

¹⁰ With the ordinary logit, the error, ζ , is iid across groups, and substitution possibilities are completely symmetric (e.g., all products belong to the same group).

¹¹ A similar generalization is used by Davis [2002], who allows α to depend on the group, g .

¹² A similar generalization is used by Brenkers and Verboven [2002].

equation (2) shows, the cross-price elasticity between i and j is independent of i . This means that the off-diagonal elements in a column of the elasticity matrix take on at most two values, depending on whether the rival product is in the same or a different group.

IV(ii). *The Distance Metric*

Brands of a differentiated product can compete along many dimensions. For empirical tractability, however, one must limit attention to a small subset of those dimensions. Nevertheless, it is not desirable to exclude possibilities *a priori*. The distance-metric demand model allows the researcher to experiment with and determine the strength of competition along many dimensions. It can thus be used to construct an empirically tractable demand equation that relies on few *a priori* assumptions. In particular, virtually any hypothesis concerning the way in which products compete (any distance measure) can be assessed in the DM framework. However, only the most important measures are typically used in the final specification.

Unlike the NML, with the DM it is assumed that individuals have a systematic taste for diversity and thus might want to consume more than one brand. Furthermore, individuals are allowed to purchase variable amounts of each brand. Finally, all individuals consume the outside good.

The DM model is based on a normalized-quadratic, indirect-utility function (Berndt, Fuss, and Waverman 1977 and McFadden 1978b) in which the prices of the differentiated products as well as individual incomes have been divided (or normalized) by the price of the outside good. This utility function is in Gorman polar form and can therefore be aggregated to obtain brand-level demands. In particular, aggregation does not depend on the distribution of unobserved consumer heterogeneity or of income. However, simplicity in aggregation is not obtained costlessly, since one must assume that all consumers have the same marginal utility of income.

Since the indirect-utility function is quadratic, the market-level demand equations are linear in normalized prices and income,¹³ and brand sales can be written as

$$(3) \quad q_i = a_i + \sum_{j=1}^n b_{ij}p_j - \gamma_i y,$$

where $B = [b_{ij}]$ is an arbitrary $n \times n$ symmetric, negative-semidefinite matrix, and prices, $p = (p_1, \dots, p_n)^T$, and aggregate income, y , have been divided by p_0 .

¹³ By Roy's identity, demands must be divided by the marginal utility of income, which is a price index here that differs over time but not by consumer. In a short time series, this number can be set equal to 1.

Equation (3) clearly has more parameters than can be estimated using a single cross section or short panel. It is therefore assumed that a_i and b_{ii} , $i = 1, \dots, n$, are functions of the characteristics of brand i , $a_i = a(x_i)$ and $b_{ii} = b(x_i)$. For example, when the product is beer, the characteristics might be the brand's alcohol content, product type (e.g., lager, ale, or stout), and brewer identity. Furthermore, the off-diagonal elements of B are assumed to be functions of a vector of measures of the distance between brands in some set of metrics, $b_{ij} = g(d_{ij})$.¹⁴ For example, when the product is beer, the measures of distance, or its inverse closeness, might be alcoholic-content proximity and dummy variables that indicate whether the brands belong to the same product type and whether they are brewed by the same firm.

As with the NML, since the intercepts depend on product characteristics, the equation is transformed from one in which consumers demand brands into one in which they demand the characteristics that are embodied in those brands. Furthermore, own-price elasticities also depend on the characteristics. In contrast to the NML, the off-diagonal elements b_{ij} that determine substitutability between brands depend on distance measures. This allows one to test hypotheses such as, for example, 'brands that have similar alcohol contents are closer substitutes'.

Let Z be the matrix of observed brand and market variables with typical row $z_i = (x_i^T, y^T)$.¹⁵ If in addition, there are unobserved brand and market variables ξ with typical element ξ_i , (3) can be written in matrix notation as

$$(4) \quad q = Z\beta + Bp + \xi,$$

where β is a vector of parameters that must be estimated. The random variable ξ can be heteroskedastic and spatially correlated. However, as with the NML, ξ is assumed to be mean independent of the observed characteristics, $E[\xi_i|X] = 0$.

The own and cross-price elasticities that are implied by equation (4) are

$$(5) \quad \eta_{ii} = \frac{p_i b(x_i)}{q_i} \quad \text{and} \quad \eta_{ij} = \frac{p_j g(d_{ij})}{q_i}.$$

A feature that distinguishes the DM from the NML is that, with the former, cross-price elasticities depend on attributes of both brands— i and j —whereas with the latter, they depend only on the characteristics of j . This means that cross-price elasticities can be modeled more flexibly. Indeed, by choosing appropriate distance measures, one can obtain models in which

¹⁴ One can restrict the indirect-utility functions, the individual demand equations, or the aggregate demand equation (as is done here). The three possibilities are equivalent. In particular, the restrictions on aggregate demand do not imply that individuals have similar preferences.

¹⁵ A market is a geographic region, time period pair. It is assumed that brands that are sold in different markets are not substitutes.

substitution patterns depend on *a priori* product groupings, as with the nested logit. There are, however, many other possibilities. Furthermore, hybrid models that include more than one distance measure are possible.

V. EVALUATING UNILATERAL AND COORDINATED EFFECTS

The term market power usually denotes the ability of firms to charge prices in excess of marginal costs. The most common measure of market power is the Lerner index, or price/cost margin, $L_i = (p_i - c_i)/p_i$.

If one has exogenous estimates of marginal costs, one can calculate n price/cost margins L_i , one for each brand, and it is possible to decompose those margins into unilateral and coordinated effects. Furthermore, since market structure can to some extent be controlled by competition authorities, whereas the degree of product differentiation cannot, it is useful to further decompose the unilateral effect. There are then three components: one that is due to differentiation, one that is due to market structure, and a third that is due to collusion.¹⁶ Finally, the sum of the first two is the unilateral effect, whereas the third, if positive, is the coordinated effect.

This procedure involves solving first-order conditions to obtain equilibrium prices of different games and calculating the associated margins. To illustrate, suppose that there are K sellers of the differentiated product and that player k , $k = 1, \dots, K$, controls a set of prices p_i with $i \in J_k$, where $\mathcal{J} = [J_1, J_2, \dots, J_K]$ is a partition of the integers $1, \dots, n$. Let P_{J_k} be the set of prices that k controls. Assume also that sellers of the differentiated product play a game, whereas the outside good is competitively supplied. For given \mathcal{J} and prices p_j with $j \notin J_k$, player k chooses P_{J_k} to

$$(6) \quad \max_{P_{J_k}} \pi_k = \sum_{i \in J_k} (p_i - c_i)q_i - F_k,$$

where c_i is the constant marginal cost of producing brand i and F_k is the fixed cost for firm k .¹⁷

The i th first-order condition is

$$(7) \quad q_i + \sum_{j \in J_{k_i}} (p_j - c_j) \frac{\partial q_j}{\partial p_i} = 0,$$

where J_{k_i} is the element of the partition to which p_i belongs.

¹⁶This decomposition is due to Nevo [2001].

¹⁷With this specification, economies of scope enter only through the firms' fixed costs.

Equation (7) nests the following games:

- i) *Bertrand behavior with single-product firms*: $K = n$.
- ii) *Bertrand behavior with multiproduct firms*: $1 < K < n$.
- iii) *Joint-profit-maximizing behavior*: $K = 1$.

Given a demand equation, a partition \mathcal{J} , and a set of marginal costs, one can solve the first-order conditions (7) for equilibrium prices and margins, $\tilde{p}_{\mathcal{J}_i}$ and $\tilde{L}_{\mathcal{J}_i} = (\tilde{p}_{\mathcal{J}_i} - c_i) / \tilde{p}_{\mathcal{J}_i}$, $i = 1, \dots, n$, of the corresponding Bertrand game. Moreover, with the DM demand equation, this calculation normally involves only matrix inversion.

The first step in the decomposition is to evaluate the market power that results from differentiation alone. One does this by solving game i). With this game, each element of the partition \mathcal{J} is a singleton, and there are n Bertrand players or decision makers, one for each brand. The margins that correspond to the equilibrium prices of this game express the market power that is due to differentiation. The implicit comparison is with marginal-cost pricing or $L = 0$.

The second step is to evaluate the market power that results from concentration, or equivalently, fewness or multibrand ownership. To do this, one solves game ii), where the partition \mathcal{J} with $1 < K < n$ corresponds to the observed brand-ownership pattern (the *status quo* partition). The margins that are associated with this game express the market power that is due to a combination of differentiation and fewness. Furthermore, differences in the margins that are associated with the two games measure the additional power that is due to fewness (i.e., to the fact that there are K rather than n firms).

The third step is to estimate the coordinated effect. If one interprets tacit collusion in the game-theoretic sense—as obtaining an outcome that is preferred by the players to the Nash equilibrium of the one-shot game—then the residual market power that has not been explained is the coordinated effect. Let the vector of Lerner indices evaluated at the exogenous cost estimates and observed prices be L_o , where o stands for observed. Then differences between L_o and the margins of the second game, if positive, can be attributed to collusion. One cannot distinguish, however, between tacit and overt collusion. Furthermore, if collusion is believed to be tacit, one cannot determine the sort of dynamic game that underlies that collusion, at least not using the methods that are described here.

This procedure can be used to test for the existence of and to estimate the average magnitude of coordinated effects, but it does not indicate which firms price less competitively or which brands are less competitively priced. It is possible to investigate that issue by examining differences in observed and predicted margins, which I call *excess margins*. Let the vector $\theta_{\mathcal{J}} = L_o - \tilde{L}_{\mathcal{J}}$, denote those differences. One can think of the θ s as

parameters that measure the extent of the deviation from the null hypothesis of static Nash-equilibrium behavior for a given game.

Since there are n first-order conditions (7) and n parameters, $\theta_{\mathcal{J}_i}$, one can solve the first-order conditions to obtain a vector of implicit excess margins, $\hat{\theta}_{\mathcal{J}}$. Alternatively, one can assume that excess margins depend on brand, firm, and regional characteristics, $\theta_{\mathcal{J}_i} = f_{\mathcal{J}}(z_i)$, and estimate the first-order conditions. Let the econometrically estimated excess margins be $\hat{\theta}_{\mathcal{J}}$. $\hat{\theta}_{\mathcal{J}}$ and $\tilde{\theta}_{\mathcal{J}}$ can be used to test hypotheses concerning any static game (e.g., $E(\theta_{\mathcal{J}}) = 0$) as well as to assess systematic deviations from the Nash equilibrium of that game.

VI. ESTIMATION

VI(i). *Demand*

The demand equations (1) and (3) contain endogenous right-hand side variables (own and rival prices and group shares) and are therefore estimated by instrumental-variables (IV) techniques. Estimation of the nested logit is entirely straightforward.

The DM equation (3) can be estimated by either parametric or semiparametric IV methods. With the parametric estimator used here, $g(\cdot)$ is a parametric function of the distance measures d_{ij} .

The issue of identification of the DM equation is complicated by the fact that the Z variables can enter both the linear part of the model, $Z\beta$, and the g function. In particular, it is not immediately obvious that g is identified, even by functional form. However, if the discrete distance measures (such as product groupings) are used in g , but no corresponding product dummies are included in Z , which is the case with the results reported later, g can be identified. In general, this procedure will not work well if price distributions and/or locations in taste space do not vary much across categories. Fortunately, with the application, there is substantial variation in both across product types.

One must also choose a stochastic specification for ξ . The covariance-matrix estimator that is used, which is nonparametric, is similar to the one that is proposed in Newey and West [1987] in a time-series context. In particular, observations that are 'close' to one another are assumed to have nonzero covariances, where closeness is measured by one or more of the distance measures (see the appendix).

VI(ii). *Excess Margins*

The first-order conditions (7) can be solved for a vector of excess margins, θ , that can be modeled as functions of the brand characteristics. In the absence of information on functional form, a simple linear

relationship is used,

$$(8) \quad \theta_i = \lambda^T z_i + \phi_i,$$

where z_i is a vector of observed firm, brand, and regional characteristics, and ϕ_i is an unobserved variable that affects markups. With the DM specification, the equation that is estimated is

$$(9) \quad Y \equiv L_o - E^{-1}e = Z\lambda + \phi,$$

where e is an n vector of ones, $e = (1, 1, \dots, 1)^T$, the $n \times n$ matrix, E , is defined by

$$E_{ij} = \begin{cases} -\epsilon_{ji} \frac{R_i}{R_j} & \text{if } j \in J_k \\ 0 & \text{if } j \notin J_k, \end{cases}$$

and $R_i = p_i q_i$. In other words, E is a weighted elasticity matrix where the weights are ratios of revenue shares.

I use a two-step procedure to estimate the first-order conditions. In the first step, the parameters of the demand equation are estimated as in the previous subsection. In the second step, the estimated demand parameters and the postulated excess-margin function is substituted into (9), and the remaining parameters are estimated. The only complication is that the standard errors of the second-stage parameters must be adjusted to reflect the fact that the demand equation was itself estimated. The method that is used to do this, which is described in the appendix, is based on suggestions of Newey [1984] and Murphy and Topel [1985]. An advantage of a two-step procedure is that misspecification of the first-order condition does not contaminate the demand estimates, in which one typically has more confidence.

VI(iii). *Choosing Instruments and Testing Their Validity*

An important issue is the choice of instruments. In particular, one needs instruments that vary by brand and market. The exogenous demand and cost variables, Z and c , are obvious candidates, and some of them (e.g., coverage) vary by brand and market. A number of other possible instruments have been discussed in the differentiated-products literature. For example, Hausman, Leonard, and Zona [1994], assume that systematic cost factors are common across regions and that short-run shocks to demand are not correlated with those factors. This allows them to use prices in one city as instruments for prices in another. Berry, Levinsohn, and Pakes [1995], in contrast, suggest that one can use rival-product characteristics as instruments.

My identifying assumptions involve a combination of the two suggestions. The first set of instruments exploits the panel nature of the data.

Specifically, I assume that prices in region one are valid instruments for prices in region two and vice versa. The brands in my sample are not brewed locally and thus have a common cost component. Moreover, it is likely that the error term in the demand equation principally reflects local promotional activity that is apt to be uncorrelated across regions.

Price in the other region, p_{-r} , can enter the instrument set directly. Moreover, it can be used to construct additional instruments. This is done by premultiplying the price vector by weighting matrices W , where each W is an element of the distance vector, d . To illustrate, suppose that W^1 is the same-product-type matrix. The instrument $W^1 p_{-r}$ has as i th element the average in the other region of the prices of other brands that are of the same type as i .¹⁸

The second set of instruments exploits the conditional independence assumption; when $E[\xi_i|X] = 0$, rival characteristics can be used to form instruments by premultiplying vectors of characteristics by weighting matrices.¹⁹ To illustrate, suppose that x^1 is the vector of alcohol contents of the brands (a column of the matrix X) and that W^1 is the same-product-type matrix. Then the instrument $W^1 x^1$ has as i th element the average alcohol content of rival brands that are of the same type as i .

Unfortunately, there are circumstances under which price instruments, p_{-r} will not be valid. In particular, anything that is unobserved and is common to both demand equations is problematic. For example, national advertising can cause demand shocks to be correlated across regions. Fortunately, national advertising creates less of a problem here than with, for example, U.S. beer, which is much more heavily promoted.²⁰ There are also circumstances under which instruments formed from rival characteristics will not be valid. This will be the case, for example, if rival characteristics should enter the demand equation directly. A formal test of exogeneity, one that is valid in the presence of heteroskedasticity and spatial correlation of an unknown form, is derived in Pinkse and Slade [2002] and is used here. Intuitively, the test involves assessing correlation between instruments and residuals, taking into account the fact that the residuals are not errors but are estimates of errors.

¹⁸The weighting matrices are normalized so that the rows sum to one.

¹⁹The use of rival characteristics here is somewhat different from their use in much of the differentiated-products literature (e.g., Berry, Levinsohn, and Pakes [1995]). There, rival characteristics are used as instruments for own price. Here, they are principally used as instruments for rival prices.

²⁰Figures taken from the MMC cost study indicate that advertising and marketing expenditures are less than one percent of sales.

VII. THE DATA

VII(i). *Demand Data*

The data are a panel of brands of draft beer sold in different markets, where a market is a time-period/regional pair.²¹ The panel also includes two types of establishments. Brands that are sold in different markets are assumed not to compete, whereas brands that are sold in the same market but in different types of establishments are assumed to compete.

Most of the demand data were collected by StatsMR, a subsidiary of A.C. Nielsen Company. An observation is a brand of beer sold in a type of establishment, a region of the country and a time period. Brands are included in the sample if they accounted for at least one half of one percent of one of the markets. There are 63 brands. Two types of establishments are considered: multiples and independents; two regions of the country: London and Anglia; and two bimonthly time periods: Aug./Sept. and Oct./Nov., 1995. There are therefore potentially 504 observations. Not all variables were available for all observations, however. When data for an observation were incomplete, the corresponding observation was also dropped in the other region.²² This procedure reduced the sample to 444 observations.

Establishments are divided into two types. Multiples are public houses that either belong to an organization (a brewer or a chain) that operates 50 or more public houses or to estates with less than 50 houses that are operated by a brewer. Most of these houses operate under exclusive-purchasing agreements (ties) that limit sales to the brands of their affiliated brewer. Independents, in contrast, can be public houses, clubs, or bars in hotels, theaters, cinemas, or restaurants.

For each observation, there is a price, sales volume, and coverage. Each is an average for a particular brand, type of establishment, region, and time period. Price, which is measured in pence per pint, is denoted PRICE. Volume, which is measured in 100 barrels, is denoted VOL, and coverage, which is the percentage of outlets that stocked the brand, is denoted COV.

VOL is the dependent variable in the distance-metric demand equation. With the nested logit specifications, in contrast, the dependent variable is LSHARE—the natural logarithm of the brand's overall market share—where the market includes the outside good. The outside good here consists of all other alcoholic beverages. Beer's share of the alcoholic beverage market averages 55%.²³

²¹ Draft sales in the U.K. in 1995 were just under 70% of total sales.

²² Observations were dropped in both regions because prices in one region are used as instruments for prices in the other.

²³ When the NML is estimated, the log of the share of the outside good is moved to the right-hand side of the equation, where it is captured by market fixed effects.

In addition, there are data that vary by brand but not by region, establishment type, or time period. These variables are product type, brewer identity, and alcohol content.

Brands are classified into four product types: lagers, stouts, keg ales, and real ales. Unfortunately, three brands—Tetley, Boddingtons, and John Smiths—have both real and keg-delivered variants. Since it is not possible to obtain separate data on the two variants, the classification used by StatsMR was adopted. Dummy variables that distinguish the four product types are denoted $PROD_i$, $i = 1, \dots, 4$. The product types also form the basis of the groups for the NML specifications, and those specifications include an explanatory variable LGS , the natural logarithm of the brand's share of the group to which it belongs.

There are ten brewers in the sample: the four nationals (Bass, Carlsberg-Tetley, Scottish Courage, and Whitbread), two brewers without tied estate²⁴ (Guinness and Anheuser Busch) and four regional brewers (Charles Wells, Greene King, Ruddles, and Youngs). Brewers are distinguished by dummy variables, $BREW_i$, $i = 1, \dots, 10$.

Each brand has an alcohol content that is measured in percentage. This continuous variable is denoted ALC . Moreover, brands whose alcohol contents are greater than 4.2% are called premium, whereas those with lower alcohol contents are called regular beers. A dichotomous alcohol content variable, $PREM$, that equals one for premium brands and zero otherwise, was therefore created.

Dummy variables that distinguish establishment types ($MULT = 1$ for multiples), regions of the country ($LONDON = 1$ for London), national brewers ($NAT = 1$ for nationals) and time periods ($PER1 = 1$ for the first period), were also created.

Finally, a variable, NCB , was created as follows. First, each brand was assigned a spatial market where brand i 's market consists of the set of consumers whose most preferred brand is closer to i in taste space than to any other brand. Euclidean distance in alcohol/coverage space was used in this calculation. Specifically, i 's market consists of all points in alcohol/coverage space that are closer to i 's location in that space than to any other brand's location. NCB_i is then the number of brands that share a market boundary (in the above sense) with i , where boundaries consist of indifferent consumers (i.e., loci of points that are equidistant from the two brands).²⁵

A number of interaction variables are also used. Interactions with price are denoted $PRVVV$, where VVV is a characteristic. To illustrate, $PRALC_i$ denotes price times alcohol content, $PRICE_i \times ALC_i$.

²⁴ Brewers without tied estate are not vertically integrated into retailing.

²⁵ The details of this construction can be found in Pinkse and Slade [2002].

TABLE I
SUMMARY STATISTICS BY PRODUCT TYPE^a LONDON AND ANGLIA DRAFT BEER
BRANDS IN SAMPLE

THREE MAJOR GROUPS				
Variable	Units	Lager	Stout	Ale
Average Price	Pence per pint	175.3	184.0	154.6
Total Volume	100 barrels	8732	1494	4451
Average Volume	100 barrels	47.5	67.9	18.7
Market Share	%	59	10	31
Average Coverage	%	10.1	31.3	6.3
Alcohol Content	%	4.3	4.1	3.9
Number of Brands		25	4	34

ALES			
Variable	Units	Cask Conditioned ('Real')	Keg
Average Price	Pence per pint	158.3	148.2
Total Volume	100 barrels	3092	1359
Average Volume	100 barrels	20.3	15.8
Market Share	%	21.5	9.5
Average Coverage	%	7.0	5.2
Alcohol Content	%	4.1	3.7
Number of Brands		21	13

^aAverages taken over brands, regions, and time periods.

Table I shows summary statistics by product type. The top half contains statistics for the three major product groups: lagers, stouts, and ales, whereas the bottom gives statistics for the two types of ales. In these tables, total volume is the sum of sales for that product type, whereas average volume is average sales per establishment. The table shows that stouts are the most expensive, lagers have the highest alcohol contents, and stouts have the highest coverage rates. The last statistic, however, is somewhat misleading, since it is due to the fact that Guinness is an outlier that is carried by most establishments. Finally, cask-conditioned ales have higher prices and sell larger volumes than keg ales. However, the volume statistics must also be viewed with caution, since some of the most popular brands have keg variants.

VII(ii). *The Metrics*

Using the same data, Pinkse and Slade [2002] experimented with a number of metrics or measures of similarity of beer brands. These include several discrete measures: same product type, same brewer and various measures of being nearest neighbors or sharing a market boundary in product-characteristic space. Two continuous measures of closeness, one in alcohol content and the other in coverage space, were also used.

They found that one metric stands out in the sense that it has the greatest explanatory power, both by itself and in equations that include several measures. That metric, *WPROD*, is the same-product-type measure that is set equal to one if both brands are, for example, real ales and zero otherwise, and then normalized so that the entries in a row sum to one. A second measure, a similar-alcohol-content measure, is also included in their final specification. That metric is calculated as $WALC_{ij} = 1 / (1 + 2|ALC_i - ALC_j|)$. I use the same metrics here.

To create average rival prices, the vector, *PRICE*, is premultiplied by each distance matrix, *W*, and the product is denoted *RPW*. For example, *RPPROD* is *WPROD* × *PRICE*, which has as *i*th element the average of the prices of the other brands that are of the same type as *i*.

VII(iii). *Cost Data*

The Monopolies and Mergers Commission performed a detailed study of brewing and wholesaling costs by brand and company. In addition, they assessed retailing costs in managed public houses.²⁶ A summary of the results of that study is published in MMC [1989]. Although the assessment of costs was conducted on a brand and company basis, only aggregate costs by product type are publicly available.

Brewing and wholesaling costs include material, delivery, excise, and advertising expenses per unit sold. Retailing costs include labor and wastage. Finally, combined costs include VAT. Table II summarizes those costs by product type. Two changes to the MMC figures were made. First, their figures include overhead, which is excluded here because it is a fixed cost. Second, their figures do not include advertising and marketing costs. Nevertheless, several of the companies report advertising expenditures per unit sold, and the numbers in the table are averages of those figures.

The last row of the table contains the updated cost figures in 1995 pence per pint. Updating was performed to reflect inflation. To do this, the closest available price index for each category of expense was collected and expenditures in each category were multiplied by the ratio of the appropriate price index in 1995 to the corresponding index in 1985.

If average variable costs in brewing are constant, these numbers are marginal costs, but if average variable costs vary with output, they either over or underestimate marginal costs. However, it is difficult to predict the direction of the bias. Indeed, due to the presence of fixed costs, average variable costs can increase with output even when there are increasing returns.

²⁶ Managed public houses are owned and operated by a brewer.

TABLE II
BREWER COSTS AND MARGINS^a

	Lager	Stout	Real Ale	Keg Ale
<i>Brewing and Wholesaling</i>				
Duty	16.4	0.0	17.0	16.9
Materials	2.3	0.0	2.7	2.5
Other	5.0	0.0	3.9	5.3
Bought-in-Beer	1.5	39.2	—	—
Delivery	5.6	0.0	4.2	4.4
Advertising and Marketing ^{b)}	0.9	0.0	0.8	0.8
B&W Cost	31.7	39.2	28.6	29.9
Transfer Price	45.4	54.2	41.6	41.1
B&W Profit	13.7	15.0	13.0	11.2
<i>Retailing</i>				
Transfer Price	45.4	54.2	41.6	41.1
Wastage	1.1	1.4	1.0	1.0
Labor	33.4	35.0	34.0	32.6
Retail Cost	79.9	90.6	76.6	74.7
Takings	94.1	104.7	82.4	81.1
Retail Profit	14.2	14.1	5.8	6.4
<i>Combined</i>				
VAT	12.3	13.7	10.8	10.6
Combined Cost	78.5	89.3	74.4	74.1
Combined Profit	15.6	15.4	8.0	7.0
<i>Updated Costs</i>				
Brewing, Wholesaling, and Retailing	132	147	125	124

Source: MMC (1989).

^{a)}Excludes overhead.

^{b)}1% of takings.

VIII. EMPIRICAL RESULTS

VIII(i). Demand

Nested Logit Demand. A number of specifications of the nested-logit demand equation (2) are shown in Table III. The first is obtained by setting $\sigma = 0$ and $\alpha_i = \alpha$, which yields the standard logit (numbers 1 and 2). The second has $\sigma > 0$ and $\alpha_i = \alpha$, which is the standard nested logit (numbers 3–5). The third is also nested but allows α to vary by brand, $\alpha_i = \alpha(x_i)$ (number 6). The fourth allows σ to vary by group (number 7), and the last allows both α and σ to vary (number 8). Furthermore, some equations are estimated by ordinary least squares (OLS) whereas others are estimated by instrumental variables (IV) with PRICE and LGSH endogenous. The method of estimation is indicated in the last column of the table.²⁷

All specifications contain alcohol content (ALC), the log of coverage (LCOV), time-period, regional, product type and brewer fixed effects. Finally, with the equations that are nested, brands are partitioned into four groups according to product type: lager, stout, keg ale and real ale.

²⁷ When σ varies by group, product-type dummies are interacted with the other instruments to create instruments that vary by group.

TABLE III
NESTED LOGIT DEMAND EQUATIONS, DEPENDENT VARIABLE: LSHARE

α and σ Constant								
#	PRICE ($-\alpha$)	Slope ($-\alpha_i$) ^{a)}	LGSH (σ)	LGSH ₁ (σ_1)	LGSH ₂ (σ_2)	LGSH ₃ (σ_3)	LGSH ₄ (σ_4)	
1	-0.0040 (-0.7)						Logit OLS	
2	-0.0059 (-1.1)						Logit IV	
3	-0.0022 (-0.8)		0.822 (29.8)				NML OLS	
4	-0.0026 (-1.3)		0.881 (11.6)				NML IV	
5 ^{b)}	0.0051 (1.4)		0.729 (8.6)				NML IV	
<i>α Variable (Price interacted with characteristics)</i>								
6	-0.0029 (-1.2)	-0.0024 (-1.0)	0.776 (14.5)				NML IV	
<i>σ Variable</i>								
7	-0.0050 (-1.7)			0.762 (12.4)	0.668 (9.6)	0.944 (12.9)	0.739 (9.3)	NML IV
<i>α and σ Variable</i>								
8	0.0039 (0.6)	0.0034 (1.1)		0.594 (4.7)	0.590 (4.0)	0.667 (6.8)	0.599 (6.7)	NML IV

$\ln(s_0)$ in market/time-period fixed effects.

Other explanatory variables: ALC, LCOV, PER1, LONDON, product and brewer fixed effects.

Four groups, lager, stout, keg ale, and real ale.

t-statistics in parentheses.

^{a)}Evaluated at the mean of the data.

^{b)}Includes PREM.

A number of empirical regularities appear in the table. First, the coefficient of PRICE is never significant at conventional levels and is sometimes positive. Second, when the same specification is estimated by OLS and IV, the slope is more negative with the IV estimation, as one would expect. Third, comparing ordinary and nested logits, the slope is always more negative with the former. Nevertheless, due to the presence of the group-share variable, LGSH, demand is more elastic with the latter (see below). Finally, with the nested specifications, the coefficients of LGSH are positive, less than one, highly significant, and assume magnitudes that imply high within-group correlation of tastes.

In the second half of Table III, $-\alpha$ is the constant coefficient of price, and $-\alpha_i = -\alpha(x_i)$ is the slope of the demand equation evaluated at the mean of the product characteristics.²⁸ This section of the table shows that when prices are interacted with characteristics, slopes are neither more negative nor statistically more significant than when they are not. When σ varies by

²⁸The characteristics that are included in this specification are the same as with the DM specifications that are presented in Table IV.

group, in contrast, the slope of the demand equation is steeper and marginally significant. Finally, when both α and σ vary, the slope is positive.

When calculating elasticities and markups, to give the NML the benefit of the doubt, the specification with the steepest slope is used (number 7).

Brand own-price elasticities are calculated holding the prices of all other brands constant. First, consider the ordinary logit elasticities. The mean and median own-price elasticities are -0.98 and -0.97 respectively, which are not reasonable values. Indeed, demand for individual brands should be highly elastic, since there are many close substitutes for a given brand. Turning to the NML, the mean own-price elasticity is -3.6 , which is also the median elasticity. Demand is therefore substantially more elastic with the nested logit than with the logit. Compared to estimates reported in Hausman, Leonard, and Zona [1994], however, where own-price elasticities for U.S. brands average -5.0 , the NML own-price elasticities still seem low.

Brand cross-price elasticities are calculated allowing the price of a single rival brand to increase, holding own price and the prices of all other rivals constant. With the logit, those elasticities vary only by brand. At the mean of the data, this elasticity is 0.009 , which is also low. The logit-demand specification is therefore not very satisfactory for this application.

NML cross-price elasticities take on two values per brand, one for brands in the same group and one for brands in different groups. At the mean of the data, those elasticities are 0.195 for the former and 0.0074 for the latter, an indication that most substitution is within groups, as the estimates of σ already suggested. Finally, there is substantial variation in partial cross-price elasticities across groups. Those differences, however, are driven almost entirely by differences in same-group shares (i.e., by differences in the number of brands in each group).

The logit elasticities are not significant at conventional levels. The NML own-price elasticities, as well as cross-price elasticities for products in the same group are significant at 1%. However, significance is due to the presence of the parameter σ_g and not due to α_i .

Distance-Metric Demand. Table IV summarizes the estimated distance-metric-demand equations. The first two specifications in this table, however, are included only for comparison with the logit and NML. Recall that the coefficients of price in those equations were sometimes positive and never significant at conventional levels. To demonstrate that this finding is not simply due to functional form, linear equations are shown in which prices are not interacted with characteristics and distance-weighted rival prices are not included. As with the logit and NML, the slopes of those equations are not consistently negative or significant at conventional levels.

The third equation in table IV is the DM specification. This equation is divided into three sections: the intercepts, $A_i = \beta^T z_i$, and the own-price terms, b_{ii} , are functions of the characteristics. The characteristics in b_{ii} however, have been interacted with price to allow the own-price elasticities

TABLE IV
DISTANCE-METRIC DEMAND EQUATIONS, DEPENDENT VARIABLE: VOL

Estimation Technique	1 IV	2 IV	3 IV	4 OLS
<i>Own Price (b_{ii})</i>				
PRICE	0.348 (1.3)	-0.811 (-1.2)	-1.125 (-2.9)	-0.871 (-2.6)
PRCOVR			0.165 (7.8)	0.153 (7.4)
PRPREM			-0.030 (-0.1)	-0.025 (-0.7)
PRNCB			-0.117 (-2.7)	-0.106 (-2.5)
<i>Rival Price (b_{ij})</i>				
RPPROD			0.712 (2.6)	0.747 (2.9)
RPALC			0.215 (1.6)	0.172 (1.0)
<i>Intercept (A_i)</i>				
LCOV	30.64 (11.9)	32.27 (11.4)	60.29 (11.7)	56.81 (13.6)
ALC	9.145 (1.4)	6.660 (0.5)	8.801 (0.7)	8.36 (0.7)
MULT	-25.93 (-4.4)	-10.47 (-1.1)	-10.97 (-1.9)	-16.03 (-3.1)
PER1	2.229 (0.5)	-0.221 (-0.1)	3.806 (0.8)	3.886 (0.8)
LONDON	30.22 (6.1)	36.60 (6.2)	31.49 (6.4)	31.13 (6.4)
Product Fixed Effects	no	yes	no	no

t-statistics in parentheses.

Standard errors corrected for heteroskedasticity and spatial correlation of an unknown form.

to vary with those characteristics. The rival-price terms b_{ij} , $j \neq i$, in contrast, depend on the distance measures.

In theory, all characteristics that are included in z_i could enter both A_i and b_{ii} . In practice, however, each characteristic is highly correlated with the interaction of that characteristic with price. For this reason, the variables that appear in A_i and those that appear in b_{ii} are never the same. An attempt was made to allocate the variables in a sensible fashion. Nevertheless, the allocation is somewhat arbitrary.²⁹ In addition, since coverage was found to be an important determinant of both brand-market size and own-price elasticity, coverage is included in both parts of the equation. To avoid collinearity, different functional forms are used in the two parts, with $LCOV = \log(COV)$ and $COVR = 1/COV$.

First, consider the own-price effect, b_{ii} , in the third specification. In contrast to the earlier findings, this slope is both negative and significant. Moreover, this is true not only of the coefficient of price, but also of most of the interaction terms. In particular, premium and popular brands have

²⁹ Other specifications were estimated, and the principal conclusions were not affected.

steeper (i.e., more negative) slopes (recall that COVR is an inverse measure of coverage), and when a brand has a large number of neighbors, its sales are more price sensitive. Allowing the slope to vary with the characteristics is therefore important here.

The second part of the equation, which assesses the determinants of brand substitutability, shows that the coefficient of the same-product-type variable, RPPROD, is both positive and significant at 1%. This implies that competition is stronger among brands that are in the same group. The coefficient of the similar-alcohol-content variable, RPALC, is positive and significant at 10%. The DM demand equation is thus similar to a nested logit, where the nests are product types. In addition to the product groupings, however, beers with similar alcohol contents tend to compete regardless of type, but the strength of that rivalry is less pronounced.

Finally, consider the intercepts, A_i . In all specifications, high coverage is associated with high sales. In addition, sales are higher in independent establishments and in London. Furthermore, a high alcohol content has a positive but weak effect on sales.

For comparison purposes, the last column of Table IV contains OLS estimates of the DM demand equation. The table shows that the OLS estimates of the coefficients of the endogenous variables tend to be smaller in magnitude than the IV estimates but are similar in significance.

Prior to using the DM demand equation, its identification and regularity were assessed. First, I determined that the instruments do indeed explain the endogenous variables. Second, I used the test of correlation between the residuals and various groups of instruments that is discussed in Pinkse and Slade [2002]. This process uncovered no evidence of endogeneity. Finally, the regularity conditions that are implied by economic theory were assessed. Symmetry, which implies that the system of demand equations is integrable, was imposed on the estimation *a priori*. With respect to curvature, all of the eigenvalues of the estimated matrix \hat{B} , which is the Hessian of the indirect-utility function, are nonnegative. This must be the case if \hat{B} is negative semidefinite and shows a close adherence to quasi-convexity of the indirect-utility function.

Turning to the elasticities, with the DM specification, brand own-price elasticities vary with the characteristics of each brand. The mean own-price elasticity, however, is -4.6 . Demand is therefore considerably more elastic than with either the logit or the NML specifications. Furthermore, it is similar to, but slightly smaller in magnitude than, the Hausman, Leonard, and Zona [1994] average of -5.0 . The median own-price elasticity is -4.1 , which reflects an asymmetric distribution with a few large values in the upper tails.

All DM own-price elasticities are significant at 1%. Cross-price elasticities for brands of the same type (e.g., two lagers) are also significant at 1%. When brands are of different types (e.g., a lager and a stout), however, their cross-price elasticities are not significant at 5% but are at 10%.

TABLE V
SUMMARY OF ELASTICITY ESTIMATES, AVERAGES ACROSS BRANDS

Own and Cross-Price Elasticities				
Demand Model	Own-Price Elasticity		Cross-Price Elasticity	
Logit	- 0.98		0.0091	
Nested Logit	- 3.6		0.0543	
Distance Metric	- 4.6		0.0632	
AIDS Hausman, Leonard, and Zona (1995)	- 5.0		0.12	

Own-Price Elasticities by Product Type				
Demand Model	Lager	Stout	Keg Ale	Real Ale
Nested Logit	- 3.6	- 2.3	- 12.1	- 2.9
Distance Metric	- 4.5	- 2.5	- 8.8	- 4.1

Finally, Table V compares average own and cross-price elasticities across models. It shows that as one moves from the logit to the nested-logit to the DM specification, the magnitudes of the elasticities increase. For comparison purposes, the table also contains the average elasticities for US brands of beer that were estimated by Hausman, Leonard, and Zona [1995], which are somewhat larger than the DM estimates.

The second half of Table V shows average own-price elasticities by product type. With both NML and DM specifications, keg ales have the most elastic demand, followed by lagers, and then by real ales. The demand for stouts is the least elastic.

VIII(ii). *Decomposition of Market Power*

Corresponding to any demand equation and partition \mathcal{J} that determines brand ownership, there is a set of static Nash-equilibrium prices and margins. Table VI summarizes the equilibrium prices and margins that are associated with various demand equations and games, where margins are calculated using predicted prices. Each of the predictions can be compared to the observed prices and margins that are summarized in the last row of the table.

For the nested logit, only *status-quo* prices are computed, where the *status-quo* game corresponds to the actual brand-ownership pattern. The table shows that the mean *status-quo* NML price is 211 pence per pint, which can be compared to the observed mean of 168. NML *status-quo* prices are thus on average 31% higher than observed prices, which means that, with the NML, observed behavior appears to be substantially more rivalrous than Bertrand. Furthermore, NML *status-quo* margins at the mean of the data are 45%, which can be compared to the observed margins, that average 30%.

TABLE VI
PREDICTED EQUILIBRIUM PRICES AND MARGINS

Demand Equation	Equilibrium	Mean Price	Standard Deviation	% Difference ^{a)}	Margins (L)
Nested Logit Distance Metric	Status Quo	211.0	38.1	31.4	45.1
	Marginal-Cost Pricing	129.1	5.2	- 23.1	0.0
	Single-Product Firms	159.4	19.8	- 5.1	23.5
	Status Quo	168.4	29.5	0.4	30.4
Observed Prices & Margins		167.8	20.2		29.9

^{a)}(Predicted-Observed)/Predicted.

One must conclude that either this market is very competitive or that the NML underestimates price sensitivity in the beer data. Although it is possible that the market is very competitive, the second alternative seems more plausible. The end result is that the estimated NML elasticities are relatively small in magnitude and insignificantly different from zero. If those estimates were taken seriously, Bertrand decision makers would choose prices that are substantially higher than the ones that are observed.

Table VI also shows three hypothetical equilibria that were calculated using the distance-metric demand equation: marginal-cost pricing, Bertrand pricing with single-product firms, and Bertrand pricing with multiproduct firms (the *status-quo* game).³⁰ The first results in prices that are on average 40 pence per pint lower than observed prices and in margins that are everywhere zero. Single-product prices, in contrast, which average 159 pence per pint, are only 9 pence lower than observed prices. This means that differentiation by itself endows the firms in this market with substantial market power and results in margins of about 23%. Finally, *status-quo* prices and margins are extremely close to observed prices and margins.

Using the DM demand equation, one can decompose the observed margins of 30% into three factors. The first – the differentiation effect – is due to the fact that brands of beer are not identical and consumers differ in their tastes for beer characteristics. This effect accounts for about three quarters of the total margin. The second – the concentration effect – is due to the fact that there are 10 rather than 63 brewers in the sample. This effect accounts for the remaining quarter, which means that there is nothing left over to be explained by tacit or overt collusion. In other words, whereas substantial market power is uncovered, all of it is due to unilateral effects, and no evidence of coordinated effects can be found.

Although this conclusion might have been unanticipated, it is similar to results reported in Nevo [2001] for the U.S. breakfast cereal industry, an

³⁰ Joint-profit-maximizing prices and margins are not shown. Indeed, since industry demand is estimated to be inelastic, the monopoly-markup model does not perform well.

industry where margins are even higher than in U.K. brewing. Furthermore, estimated margins for beer and cereals are substantially higher than average markups in other processed food industries, which are less than 20%. High markups are therefore unapt to be simply due to the presence of fixed costs in a zero-profit, free-entry equilibrium.

VIII(iii). *Further Analysis of Coordinated Effects*

The decomposition is not a statistical test. One can use the implicit excess-margins, $\hat{\theta}$, to test for coordinated effects more formally. A discussion of the tests used is contained in the appendix. First, consider the implicit estimates obtained from the NML demand equation, $\hat{\theta}_{NML}$. Both the mean and the median of those parameters are -0.38 , and one can reject the hypothesis that the estimates are on average zero. This means that, as before, NML behavior is estimated to be significantly more rivalrous than Bertrand. Next, consider the implicit DM excess margins, $\hat{\theta}_{DM}$. The mean of those parameters is -0.009 , the median is -0.024 , and the range is -0.30 to 0.56 . Furthermore, the t-statistic for the hypothesis that $1/n \sum_i \theta_{iDM} = 0$ is -0.66 , which means that, on average, predicted and observed margins are equal and Bertrand behavior cannot be rejected.

It appears that although the choice between NML and DM specifications strongly influences the conclusions that can be drawn concerning firm behavior, with neither specification is there any evidence of collusion or coordinated effects. This finding, however, does not rule out the possibility that some firms (i.e., the dominant group) behave in a collusive fashion while others behave more competitively. One can use an econometrically estimated excess-margin function to assess this possibility.

Given the insignificance of the logit and NML elasticities, joint estimation of a first-order condition with one of those equations does not seem worthwhile. In particular, since the first-stage estimates are imprecise, it is unlikely that the second-stage estimates, which build on the first, would be more accurate. The DM elasticities, however, are precisely measured, and a DM first-order condition is estimated.

Table VII shows three specifications for a DM excess-margin function. The first contains only the binary variable NAT that equals one if the firm that produces the brand is a national brewer – a member of the hypothesized dominant group. The second two, which differ from one another according to the measure of brand strength that is used, also contain brand characteristics. The table shows that the coefficients of NAT are never significant at conventional levels, which implies that brands that are brewed by the nationals are not more collusively priced and that those firms do not form a dominant group. On the other hand, the second and third specifications show that more popular and higher-strength brands, as well as those that are sold in multiple establishments, are less competitively

TABLE VII
EXCESS MARGINS, ACTUAL - STATUS QUO
2-STEP ESTIMATES USING THE DM DEMAND EQUATION

Equation	1 ^{a)}	2	3
NAT	0.005 (0.2)	-0.002 (-0.1)	-0.003 (-0.1)
LCOV		0.038 (3.7)	0.037 (3.4)
PREM		0.062 (3.8)	
ALC			0.047 (3.6)
MULT		0.052 (2.9)	0.052 (2.9)
PROD ₁ (lager)		-0.142 (-4.1)	-0.315 (-4.6)
PROD ₂ (stout)		-0.028 (-0.6)	-0.218 (-2.3)
PROD ₃ (keg ale)		-0.113 (-2.8)	-0.270 (-3.7)
PROD ₄ (real ale)		-0.204 (-6.9)	-0.374 (-6.1)

^{a)}Includes a constant.

t-statistics in parentheses.

Standard errors corrected for heteroskedasticity and spatial correlation of an unknown form.

priced.³¹ Finally, there appear to be few systematic differences across product types.

IX. CONCLUDING REMARKS

No evidence of coordinated effects in U.K. brewing has been found. In particular, the national brewers do not appear to form a dominant group. Nevertheless, the industry is relatively concentrated and the firms possess substantial market power. The rules of thumb that are used by North American competition authorities would therefore dictate close scrutiny of mergers in the industry. In contrast, unless the merger resulted in a firm with more than 40% of the market, the rules that are used by E.U. authorities would dictate a more lenient policy. However, using the criterion of unilateral effects, Pinkse and Slade [2002] conclude that a proposed merger between two U.K. brewing firms (Bass and Carlsberg/Tetley) that would have resulted in a merged firm with less than 40% of the market would have caused substantial price increases.

Differences in policies across jurisdictions are somewhat disconcerting, at least when they lead systematically to different decisions. With respect to

³¹ The third regularity is perhaps due to the fact that vertical relationships between brewer and retailer are more complex when public houses are owned by brewers or retail chains (see Slade [1998]).

mergers, the guidelines that have been adopted by the European Commission (E.C.) seem to systematically discount mergers in consumer product industries. For a while, it seemed that merger policies on the two sides of the Atlantic were converging. Indeed, in the recent Airtours decision,³² the E.C. stated that joint dominance is ‘not just about tacit collusion.’ Instead, it is ‘sufficient for oligopolists to act – independently – in ways which reduce competition.’³³ Nevertheless, the European Court reversed the Commission’s decision, and E.C. regulations still do not allow mergers to be prohibited unless they create or reinforce dominance.

Although brewer characteristics appear not to influence systematic departures from static Nash-equilibrium pricing, brand characteristic do. In particular, brands that have larger market shares, higher alcohol contents, or are sold in multiple establishments are found to be less competitively priced. Contrary to the U.K. Monopolies and Mergers Commission [1989] claim, however, pricing differences across product types appear to be small. In particular, I find no evidence that lagers are more collusively priced.

REFERENCES

- Berndt, E.R.; Fuss, M.A. and Waverman, L., 1977, ‘Dynamic Models of the Industrial Demand for Energy,’ Electric Power Institute, EA-580, Palo Alto, CA., U.S.A.
- Berry, S.; Levinsohn, J. and Pakes, A., 1995, ‘Automobile Prices in Market Equilibrium,’ *Econometrica*, 63, pp. 841–890.
- Brenkers, R. and Verboven, F., 2002, ‘Liberalizing a Distribution System: the European Car Market,’ mimeo.
- Compte, O.; Jenny, F. and Rey, P., 2002, ‘Capacity Constraints, Mergers, and Collusion,’ *European Economic Review*, 46, pp. 1–29.
- Davis, P., 2002, ‘Estimating Multi-Way Error Components Models with Unbalanced Data Structures,’ *Journal of Econometrics*, 106, pp. 67–95.
- Doob, J. L., 1953, *Stochastic Processes* (Wiley, New York).
- Hausman, J.; Leonard, G. and Zona, D., 1994, ‘Competitive Analysis with Differentiated Products,’ *Annales d’Econometrie et de Statistique*, 34, pp. 159–180.
- Ivaldi, M. and Verboven, F., 2000, ‘Quantifying the Effects from Horizontal Mergers: The European Heavy Trucks Market,’ University of Toulouse, mimeo.
- Jayaratne, J. and Shapiro, C., 2000, ‘Simulating Partial Asset Divestitures to “Fix” Mergers,’ *International Journal of the Economics of Business*, 7, pp. 179–200.
- Kuhn, K.-U., 2000, ‘An Economist’s Guide through the Joint Dominance Jungle,’ mimeo.
- Kuhn, K.-U. and Motta, M., 2001, ‘The Economics of Joint Dominance,’ mimeo.
- Lexicon, Ltd., 1999a, ‘Joint Dominance,’ Competition Memo, London.
- Lexicon, Ltd., 1999b, ‘The Airtours Case,’ Competition Memo, London.
- McFadden, D., 1974, ‘Conditional Logit Analysis of Qualitative Choice Behavior,’ in Zarembka, P. (ed.) *Frontiers in Econometrics* (Academic: New York).

³² 22.09.99, IV-M.1524.

³³ This statement is quoted in Lexicon [1999b]. For further analysis of the Airtours case, see Motta [1999].

- McFadden, D., 1978a, 'Modeling the Choice of Residential Location,' in Karlqvist, A. *et al.* (eds.) *Spatial Interaction Theory and Planning Models* (Amsterdam: North Holland).
- McFadden, D., 1978b, 'The General Linear Profit Function,' in Fuss, M. and McFadden, D. (eds.) *Production Economics: A Dual Approach to Theory and Applications*, Vol. 1, (Amsterdam: North Holland).
- Motta, M., 1999, 'EC Merger Policy and the Airtours Case,' European University Institute mimeo.
- Murphy, K.M. and Topel, R.H., 1985, 'Estimation and Inference in Two-Step Models,' *Journal of Business and Economic Statistics*, 3, pp. 370–379.
- Nevo, A., 2000, 'Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry,' *RAND Journal of Economics*, 31, pp. 395–421.
- Nevo, A., 2001, 'Measuring Market Power in the Ready-to-Eat Cereal Industry,' *Econometrica*, 69, pp. 307–342.
- Newey, W.K., 1984, 'A Method of Moments Interpretation of Sequential Estimators,' *Economics Letters*, 14, pp. 201–206.
- Newey, W.K. and West, K.D., 1987, 'A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimator,' *Econometrica*, 55, pp. 703–708.
- Pinkse, J., Slade, M.E. and Brett, C., 2002, 'Spatial Price Competition: A Semiparametric Approach,' *Econometrica*, 70, pp. 1111–1155.
- Pinkse, J. and Slade, M.E., 2002, 'Mergers, Brand Competition, and the Price of a Pint,' *European Economic Review*, forthcoming.
- Shapiro, C., 1996, 'Mergers with Differentiated Products,' *Antitrust*, spring, pp. 23–30.
- Slade, M.E., 1998, 'Beer and the Tie: Did Divestiture of Brewer-Owned Public Houses Lead to Higher Beer Prices?,' *Economic Journal*, 108, pp. 565–602.
- U.K. Monopolies and Mergers Commission, 1989, *The Supply of Beer* (Her Majesty's Stationery Office: London).
- Werden, G.J. and Froeb, L.M., 1994, 'The Effects of Mergers in Differentiated Products Industries: Logit Demand and Merger Policy,' *Journal of Law, Economics, and Organization*, 10, pp. 407–426.
- Werden, G.J., 1999, 'Expert Report in *United States v. Interstate Bakeries Corp. and Continental Baking Co.*, U.S. Department of Justice mimeo.

APPENDIX

ESTIMATION AND TESTING

*The Two-Step Estimation**Stage 1: Estimation of the Demand Equation*

One can write the demand equation as

$$f_1(X_1, \beta) = v,$$

where X_1 is an $n \times k_1$ matrix of endogenous and exogenous variables, β is a p_1 vector of parameters, and v is an n vector of errors. Let S be an $n \times m_1$ matrix of instruments with $m_1 > p_1$.

The IV estimator of β is the minimum over β of

$$v^T S(S^T \Omega_1 S)^{-1} S^T v,$$

where Ω_1 is a matrix that corrects for heteroskedasticity and spatial correlation of an unknown form. Specifically, Ω_1 has i, i element \hat{v}_i^2 and i, j element $\iota_{ij}\hat{v}_i\hat{v}_j$, where \hat{v} is the vector of two-stage least-squares residuals, and ι_{ij} equals one if j is one of i 's J closest neighbors and vice versa, one half if either i is one of j 's or j is one of i 's J closest neighbors (but not both), and zero otherwise. Closeness between i and j is measured here by $\text{WPROD}_{ij} \times \text{WALC}_{ij}$, where WPROD and WALC are the metrics that appear in the demand equation.

This yields $\tilde{\beta}$ and $\tilde{\Sigma}_\beta$, the IV estimates of β and $\text{Var}(\beta)$, where

$$\tilde{\Sigma}_\beta = [H_{1\beta}^T S (S^T \Omega_1 S)^{-1} S^T H_{1\beta}]^{-1},$$

and $H_{1\beta}$ is the $n \times p_1$ matrix $\partial f_1 / \partial \beta$ evaluated at $\tilde{\beta}$.

Stage 2: Estimation of the First-Order Condition

One can write the first-order condition as

$$Y(\beta) = X_2 \lambda + u \quad \text{or} \quad f_2(Y, X_2, \beta, \lambda) = u,$$

where Y is an endogenous variable, X_2 is an $n \times p_2$ matrix of exogenous variables, λ is a p_2 vector of parameters, and u is an n vector of errors. Since this equation is exactly identified, the IV estimates, $\tilde{\lambda}$, can be obtained by simply solving the moment conditions. The standard errors of $\tilde{\lambda}$, however, must be corrected to reflect the fact that β was estimated in a prior stage.

Let Ω_2 be defined like Ω_1 with \hat{v} replaced by \hat{u} and $H_{2\beta}$ be the $n \times p_1$ matrix $\partial f_2 / \partial \beta$, evaluated at $\tilde{\beta}$. Then, if u and v are uncorrelated,³⁴

$$\begin{aligned} \tilde{\Sigma}_\lambda &= (X_2^T X_2)^{-1} (X_2^T \Omega_2 X_2) (X_2^T X_2)^{-1} \\ &\quad + (X_2^T X_2)^{-1} (X_2^T H_{2\beta} \tilde{\Sigma}_\beta H_{2\beta}^T X_2) (X_2^T X_2)^{-1}. \end{aligned}$$

Hypothesis Tests

Two sorts of excess margins, θ , are estimated. Hypotheses concerning the econometrically estimated parameters, $\hat{\theta}$, can be tested using standard techniques. In addition, the implicit variables, $\tilde{\theta}$, which are nonlinear functions of estimated parameters, are themselves random variables that can be the subject of tests. Two methods of testing hypotheses concerning implicit variables are used. The first and simpler of the two is based on the fact that any sequence of i.i.d. variates with uniformly bounded moments greater than two, whether they are estimates or not, have a limiting normal distribution.³⁵ Unfortunately, the notion that the estimates, $\tilde{\theta}_i$, $i = 1, \dots, n$, are independent across i , even in the limit, is questionable. If they are dependent, their

³⁴ v is an unobserved demand factor, whereas u is an unobserved cost factor. The assumption that they are uncorrelated is thus not unreasonable. The formula is similar to equation (8) in Newey (1984) for the uncorrelated case. The difference is that his first stage estimation is exactly identified.

³⁵ This follows from the Lindberg theorem (e.g., Doob [1953], theorem 4.2).

standard errors will in general be larger and rejection of the null less likely. When the null is not rejected, only this test is used. The second test, which is used when the null is rejected by the first, is a parametric bootstrap. In particular, repeated draws from the estimated joint distribution of the parameters are performed, the desired quantity is calculated, and a bootstrap distribution is generated.