Volume Title: Trade Policy Issues and Empirical Analysis

Volume Author/Editor: Robert E. Baldwin, ed.

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-03607-3

Volume URL: http://www.nber.org/books/bald88-2

Publication Date: 1988

Chapter Title: Empirics of Taxes on Differentiated Products: The Case of Tariffs in the U.S. Automobile Industry

Chapter Author: James Levinsohn

Chapter URL: http://www.nber.org/chapters/c5846

Chapter pages in book: (p. 9 - 44)

I Measuring Trade Policy Effects under Imperfectly Competitive Market Conditions

This Page Intentionally Left Blank

2 Empirics of Taxes on Differentiated Products: The Case of Tariffs in the U.S. Automobile Industry

James Levinsohn

Recent theoretical advances in the industrial organization literature have provided insight into modeling the demand for differentiated products. Lancaster (1979) introduced and developed what he termed the "characteristics approach" to modeling the demand for differentiated products, while Dixit and Stiglitz (1977) pioneered what has come to be known as the "love of variety" approach to the subject. Both approaches have been applied to international trade theory. The result has been a heightened awareness of the role that product differentiation plays in trade theory. This work is presented in Helpman and Krugman (1985).

There have thus far been relatively few empirical applications of the new theories of trade. In this chapter, I present a new technique for econometrically estimating the demand for differentiated products. I adopt a Lancasterian approach to product differentiation and use theoretical results from this approach to solve several empirical problems. I then apply the technique to the demand for automobiles.

The estimates derived from this method allow me to analyze various trade and industrial policies for the U.S. automobile industry. For example, what would be the effect of a tariff applied only to Japanese imports on the total automobile import demand? Would domestically produced auto sales replace the Japanese imports or might German and Swedish imports rise, leaving total imports relatively constant? Some economists have argued for a tariff on all small foreign cars. Such a tax does not discriminate by country of origin and hence is not

James Levinsohn is assistant professor in the Department of Economics at the University of Michigan, Ann Arbor.

inconsistent with General Agreement on Tariffs and Trade (GATT) rules. As small foreign cars became more costly, would domestic car sales rise substantially or would the United States just trade imports of small cars for imports of larger cars? Optimal industrial policy toward the U.S. auto industry may involve subsidies to domestic producers, thereby possibly lowering the price of domestic autos.[1] Or perhaps government policy may involve subsidizing only one producer (e.g., Chrysler). What effects would these policies have on demand for different types of foreign and domestic automobiles?

All of these questions are, in a formal sense, quite similar. Each considers the effect of a tax placed on a subset of a group of differentiated products. Parameters needed to answer questions such as those posed above are own-price and appropriately defined cross-price elasticities of demand. Any analysis of the taxation of differentiated products must estimate (or use existing estimates of) these demand elasticities. The approach developed in this chapter provides a utility-consistent technique for deriving these elasticities. While I apply the methodology to issues of trade and industrial policy in the U.S. automobile industry, I believe that the general approach will have wider application. The methodology could, for example, be used to estimate demand elasticities in other differentiated products industries such as microcomputers, audio-video equipment, lumber, and steel. All of these industries have been the subject of recent trade policy debate.

This chapter is a first attempt at solving some of the empirical issues associated with the analysis of taxation of differentiated products. While the chapter provides some answers, it also raises a number of microeconomic and econometric issues for future research.

In section 2.1, I provide a brief critical review of the literature. Section 2.2 develops the methodology that is then applied in section 2.3. Using the demand system estimated in section 2.3, section 2.4 addresses many of the policy concerns posed in this introduction. Section 2.5 concludes with a brief summary.

2.1 A Brief Summary of the Literature

In theory, estimating the demand system for a set of differentiated products is no different than estimating a demand system for several homogeneous products. A typical estimated equation in such a system would regress quantity of a good demanded on its own price, the prices of the other differentiated or homogeneous products, and several other variables such as income and personal and demographic characteristics. Food is a good example of a set of differentiated products whose demand functions are nicely estimated by standard techniques. Recent work based on Deaton and Muellbauer's almost-ideal demand system provides excellent examples of this approach.[2]

For many sets of differentiated products, though, standard techniques are inapplicable. In the case of automobiles, there are over one hundred models available and few models are available for more than four consecutive years. The standard techniques would imply a system of, say, one hundred equations with ninety-nine cross-price effects. With so few years of data, the system is not estimable with any degree of accuracy. In the case of video cassette recorders or microcomputers, technology changes so quickly that no more than two or three years' data are likely to be available.

Several approaches to these problems have been taken in the empirical literature. Almost all have been applied to the automobile industry—at least partly because data are relatively plentiful. I accordingly focus on this body of research.

The easiest way around the problems posed by product differentiation is to ignore the issue. Not surprisingly, this was the approach first adopted. Work by Suits in 1958 used time series of total quantity of autos sold, average auto price, and real disposable income to arrive at aggregate demand elasticities. While it is surely unfair to judge the econometric methods of thirty years ago by the standards permitted by today's computing technology, Suits's approach is incapable of addressing the issues raised in this chapter's introduction. Surprisingly, research as recent as Toder (1978) uses elasticities imputed in part from Suits's original work when analyzing current automobile trade and industrial policy. Tarr and Morkre (1984) and Dixit (1987) in turn use elasticities derived from Toder.

Time-series techniques, even modern ones, are not applicable to investigating the effects of trade policy in the U.S. automobile industry because both products and tastes have changed significantly over the period of estimation (approximately the last twenty years). A 1965 Toyota is not the same car as a 1985 Toyota. As Toyotas change, the meaning of a single (constant) elasticity of demand for Toyotas becomes unclear. Tastes for autos and the characteristics that comprise them have also changed. While it may be theoretically possible to control for the reputation effects and network externalities that are responsible for this shift of tastes, it is not easy to do so in practice.

The most recent comprehensive study of the U.S. demand for automobiles is reported in Toder's (with Cardell and Burton) *Trade Policy and the U.S. Automobile Industry.* In that book, demand elasticities are estimated using three methodologies. Because most studies of the welfare effects of trade policy in the auto industry have used elasticity estimates from Toder, it is worthwhile taking a close look at these alternative approaches. Each is discussed in turn.

Toder's first approach is a time-series analysis. This work is more sophisticated than earlier work in that it introduces hedonic price indexes. Toder estimates the following regression:

$$ln\left(\frac{F}{D}\right) = \alpha_0 + \alpha_1 \, ln\left(\frac{P_f}{P_d}\right) + \alpha_2\bar{Z},$$

where F/D is the foreign domestic auto sales ratio. P_f/P_d is the ratio of foreign to domestic hedonic prices. $\bar{Z}$ is a vector of exogenous variables.

The estimation uses annual data from 1960 to 1974. Estimates of α_1 ranged from -0.9 to -1.7, depending on the Z vector. The coefficient α_1 is the elasticity of substitution in demand. Using the estimate of α_1 and older estimates of total market demand elasticities, conventional price elasticities of demand can be derived.

There are at least four problems with this approach. First, as mentioned above, tastes seem to have changed over time, since casual empiricism suggests that a foreign car in 1960 was viewed very differently from one in 1974. As tastes vary over time, the economic relevance of the estimates of the elasticity of substitution in demand is called into question. Second, older estimates of the total market elasticity of demand are required to convert Toder's results into standard price elasticities of demand. While Toder used hedonic price indexes, the older studies did not. As cars are not homogeneous products, it is unclear exactly what the results of the older studies by Suits and others mean. Also, the older studies were conducted before auto imports were an empirically relevant phenomenon. Using these older out-of-sample market elasticities to derive the standard elasticities of demand may yield very misleading results. Third, even if the time series would yield accurate estimates, the agglomeration of all foreign cars prevents the analysis of taxes applied to only a subset of foreign autos. Fourth, regressing relative demands on relative hedonic prices does not follow from either a Lancasterian or Dixit-Stiglitz model of product differentiation. The choice of using relative demands and relative prices of domestic and foreign goods allows Toder, like all his predecessors, to estimate a single equation instead of a complete demand system. Toder's implicit assumption that an otherwise homogeneous good is differentiated only by country of origin is termed the Armington assumption. This assumption makes little sense from a consumer theory viewpoint unless there is some basis for supposing that goods are homogoeneous within countries but not across countries. Toder's first approach is, then, a utility inconsistent approach to modeling demand for differentiated products.

Toder's second method employed a cross-sectional approach to the demand-estimation problem. Toder used transport costs to introduce cross-sectional price variation. The units of observation were each of the continental United States. Here the regression estimated was

$$\frac{F}{D} = \alpha_0 + \alpha_1\left(\frac{P_f}{P_d}\right) + \alpha_2 PC2029 + \alpha_3 PCI + \alpha_4 PGAS,$$

where F/D is the ratio of foreign to domestic new car sales, P_f/P_d is the ratio of *delivered* foreign to domestic list prices, *PC*2029 is the percentage of population between ages 20 and 29, and *PGAS* is the price of gasoline.

While the problems of time-varying parameters are not present in this cross-sectional approach, this method still relies on previously derived market elasticities to construct conventional price elasticities of demand. The cross-sectional methodology yielded generally unsatisfactory results. This is not surprising, since one might suspect that variables other than transport costs, gasoline costs, per capita income, and the percentage of the population in their twenty's explain why foreign cars are more predominant in New Jersey or California than they are in Michigan or Indiana.

Toder's third approach is by far the most innovative. Although computationally complex, the intuition behind this methodology—termed a hedonic market share model—is straightforward.

The model requires only one year's data on sales, list prices, and characteristics of automobiles. Let $\{\alpha^i_n\}$ be the set of marginal rates of substitution between N characteristics and price. Toder posits a log-normal distribution of $\{\alpha^i\}$ across consumers. Next, he estimates coefficients, β, which form a vector of sufficient statistics for the probability distribution of the α's. Let S be the vector of actual shares of auto sales by model. He next chooses β to maximize the likelihood of observing S. In brief, the technique selects statistics describing a distribution of consumer's utility functions that reproduce as nearly as possible the actual market shares observed.

Toder then applies the estimated taste distribution to a new set of available models (differing from the old set by price) to generate a new market share distribution. In this sense, the model simulates the relevant elasticities. Unlike the previous two approaches, the hedonic market share model can, in principle, predict market share elasticities for any subset of models. In practice, only an elasticity of substitution in demand between all foreign and all domestic cars is estimated. This yielded coefficients of -2.3 and -2.1, depending on the price increase simulated.

There are at least three major problems with this approach, the first two being interrelated.

1. The model is computationally quite difficult. Toder uses five characteristic variables to estimate the taste distribution. Calculating the maximum likelihood estimates for β requires a fifth-order numerical integration between each iteration of the likelihood function maximization. The cost of such computational techniques can be prohibitive. Also, some experts at numerical analysis question the accuracy of such a high-order integration of a complicated distribution function.

2. More importantly, this technique does not yield standard errors. For policy analysis, point estimates without standard errors are of limited use. Without the standard errors, it is impossible to know whether and how well the data fit the model.

3. The results of this technique hinge critically on the choice of the distribution function of tastes. Toder used a log-normal distribution. The choice of the distribution function is completely arbitrary, yet possibly key to the results. While all nonrobust estimation methods are subject to this critique, the problem is compounded here by the lack of standard errors of the estimates. Without the standard errors, it is especially difficult to ascertain whether the distribution function of tastes chosen fits the model.

Bresnahan (1981) also models the demand for automobiles. Using sophisticated econometric techniques, he accounts for product differentiation and avoids the pitfalls of time-series analysis. His goal, though, is more ambitious than just a model of automobile demand, as he focuses on the issue of departures from marginal cost pricing in the automobile industry. Because he looks at a broader range of issues than just the demand side of the model, his results are not disaggregated enough to analyze the questions posed in the introduction of this chapter. While he does not estimate elasticities per se, estimated parameters can be manipulated to give an industry demand elasticity (a proportionate change of all prices) of .25 and an elasticity for the *average* product (one price changes and all others are constant) of 3.2. Bresnahan is forthright about the restrictive assumptions that he requires on the demand side of his model. The most serious of these is the assumption that the density of consumer tastes is uniform (as opposed to Toder's log-normal assumption). Bresnahan's methodology also is computationally complex, and, like Toder's hedonic market share model, it does not yield estimates of standard errors. Bresnahan, though, approximates the variances of parameter estimates in four ways. Although variances depend on the approximation used, this does give some feel for how well the data fit the model. In short, Bresnahan's method is carefully developed, but it is not suitable for addressing the types of issues raised in the introduction of this chapter.

Finally, a number of studies of automobile demand investigate whether a car is purchased at all, and if so, how many are purchased. These studies are fairly common in mode-of-transportation studies. Methods used range from simple logit to multinomial logit to multinomial probit. A technically sophisticated example of this approach is found in Train (1986).[3] These studies ask a set of questions that are for the most part only tangentially related to questions about the demand effect of taxes on differentiated products. As such, their results are not very useful to the issues with which I am concerned.[4]

2.2 Methodology

In this section, I explain my approach to the estimation of demand for differentiated products. This is done in two steps. In step 1, the demand function to be estimated is derived. I avoid many of the pitfalls of previous approaches by relying on results from Lancasterian consumer theory. In step 2, I explain how the insights offered by Lancasterian consumer theory are empirically implemented.

Step 1: I avoid the problems associated with time-series analysis by using only three years of data—1983 through 1985.[5] Three years of time-series data, though, leaves few degrees of freedom. The much-needed additional price-quantity variation is introduced by using a cross section of (the same) 100 models of automobiles for each year. The data are a time-series cross section, or panel, consisting of 300 observations.[6]

While using panel data instead of only time-series introduces additional price-quantity variation, it also poses some problems. It may be wrong to regress quantity on price since, across observations, the good is not the same. I address these problems using results from the characteristics approach to product differentiation.

In the Lancasterian model of product differentiation, a good is represented by its bundle of characteristics. Different models of the good contain different bundles of these characteristics. With this view of product differentiation, as tastes vary across consumers, demands for a model, given its price, will vary with the model's characteristics bundle. Because products are identified by their bundle of characteristics, it is appropriate to control for the cross-sectional variation in models by including in the demand function those characteristics that differentiate models.[7]

Lancaster hence posits that the quantity demanded of a model depends on its own price and characteristics and on the price and characteristics of competing models. In log-linear form, this implies

$$ln\ Q_{it} = \alpha_0 + \alpha_1 ln\ P_{it} + \alpha_2 ln\ \bar{P}_{jt} + \beta' X_{it} + \Gamma' X_{jt},$$

where Q_{it} is the quantity demanded of model i in year t, P_{it} is the price of model i in year t, $\bar{P}_{jt}$ is the vector of prices of substitutes to a model with sales Q_{it}, X_{it} is a characteristics vector of model i in year t, and X_{jt} is a characteristics vector of model j in year t.

I posit that the above model may be subject to country-of-origin specific errors, and hence use a fixed effects model.[8] Allowing also for time-dependent shifts of demand gives

$$\begin{aligned}(1A)\quad ln\ Q_{it} = {} & \alpha_0 + \alpha_1 ln\ P_{it} + \alpha_2 ln\ \bar{P}_{jt} + \beta' X_{it} + \Gamma' X_{jt} \\ & + \alpha_3 JAPAN_i + \alpha_4 GERMAN_i + \alpha_5 SWEDE_i + \alpha_6' T_t,\end{aligned}$$

where $JAPAN_i = 1$ if model i is Japanese, $GERMAN_i = 1$ if model i is German, $SWEDE_i = 1$ if model i is Swedish, and T_t is a time dummy for year t. Equation (1A) is consistent with a Lancasterian approach to consumer demand for autos.

Somewhat surprisingly, Lancaster's work does not discuss the hedonic price literature, which posits that the price of a good is a linear combination of the implicit prices of the attributes of the good. Thus in equation (1A), X_{it} would be highly collinear with P_{it}. An analogous relationship holds for X_{jt} and P_{jt}. According to the hedonic approach, the price of a good already contains information about the qualities of the good. Hence, estimating equation (1A) merely introduces severe multicollinearity. Instead, the hedonic hypothesis argues in favor of estimating the following demand function.

$$\text{(1B)}\quad \ln Q_{it} = \alpha_1 + \alpha_2 \ln P_{it} + \alpha_2 \ln \bar{P}_{jt} + \alpha_3 JAPAN_i + \alpha_4 GERMAN_i + \alpha_5 SWEDE_i + \alpha_6' T_t.$$

I econometrically consider both equations (1A) and (1B). In doing so, I assume the consumer takes as given all independent variables.

The functional form of the demand function should follow from the density of consumers over characteristics space. Formally, demand for a model is given by integrating the density of consumers over the neighborhood of the model. Making the link between distribution of consumers and functional form of demand is a difficult question that I do not address. Rather, I consider equations (1A) and (1B) as convenient statistical approximations of demand.

In standard consumer theory, with 100 models, 99 models could serve as substitutes for model i, and thus 99 prices would appear in $\bar{P}_{jt}$. This would imply 9,900 cross-price terms to be estimated in the standard demand system. This is not feasible with only three years of data. Again, I rely on the theory of product differentiation to, in effect, place many zero restrictions on the vector α_2.

The earliest work on product differentiation by Hotelling (1929) arranged products along a line. In figure 2.1, model B competes for customers with models A and C but not with any other models.

Here, models A and C are termed "neighbors" of model B, whereas the other models (D, E, etc.) were not. Were there 100 models arranged along the spectrum, this setup would imply 97 zero restrictions on the vector α_2 for good B. Only the price of B and the prices of its neighbors, A and C, would enter the demand function for B.

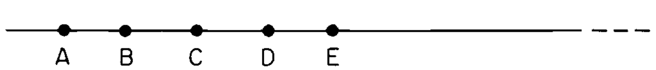

Fig. 2.1 Hotelling product differentiation

Lancaster extends the Hotelling model to allow products to differ across more than one dimension. Lancaster posits that each good is a bundle of several characteristics. In this case, if there are n products, each product may have up to $n-1$ neighbors and all have at least one neighbor.[9] I rely on the Lancasterian approach to product differentiation to determine endogenously which products compete with each other for consumers. This, in turn, allows me to place zero restrictions on α_2 in a utility-consistent manner.

Step 2: Empirically determining the neighbors for each product is complicated by the fact that while characteristics of the goods are observed, individual consumer tastes over these characteristics are not. I adopt an approach to this problem that is based in part on a methodology developed by Feenstra (1986).[10]

The first task in any Lancasterian model is to define the metric in characteristics space that is to be used to determine how far apart any two products are. To this end, let $x = (x_1, x_2, \ldots, x_n) > 0$ be a vector of physical characteristics that differ across models and X^n be the n-dimensional space in which products are differentiated. Let $\Theta = (\Theta_1, \Theta_2, \ldots, \Theta_n)$ represent the vector of taste parameters for a particular individual.

I assume that all individuals have the same form of utility function, namely, CES, but that individuals differ in their vector of tastes Θ. Then, an individual's utility is given by[11]

$$U(x,\Theta) = \sum_{i=1}^{n} \Theta_i x_i^\delta. \tag{2}$$

The parameter δ is related to the elasticity of substitution between characteristics, σ, that is,

$$\sigma = \frac{1}{\delta - 1}.$$

The twin constraints of utility increasing in x and concavity of utility in x imply $\sigma \in (0, -1)$. This range of σ is perhaps overly restrictive for the case of substitutability of auto characteristics. In order to permit $\sigma \in (0, -\infty)$, I take a Box-Cox transformation of equation (2). This yields

$$U(x,\Theta) = \sum_{i=1}^{n} \Theta_i \tilde{x}_i^\delta,$$

$$\text{where } \tilde{x}_i^\delta = (x_i^\delta - 1)/\delta, \ 0 \text{ not equal to } \delta < 1,$$

$$\text{and } \tilde{x}_i^\delta = 1n\, x_i, \text{ if } \delta = 0.$$

As I will be working with the case of less than perfect substitutability between characteristics, I will, for notational simplicity, henceforth use the (still CES) utility function:

$$U(x,\Theta) = \sum_{i=1}^{n} \Theta_i\,(x_i^{\delta} - 1)/\delta. \tag{3}$$

The price of a model depends on its characteristics. I specify the functional form for $P(x)$. In particular,

$$P(x) = \exp(\alpha + \beta' x), \tag{4}$$

where $\alpha > 0$ and $\beta = (\beta_1\ ,\ \beta_2\ , \ldots ,\ \beta_n) > 0$ are parameters.

Denoting the homogeneous numeraire good by N and exogenous income by Y, the consumer's problem is to

$$\max_{x,N}\ U(x,\Theta) + N \tag{5}$$

subject to $P(x) + N \leq Y$. The additively separable form of the utility function in equation (5) and the linearity in N imply that the optimal choice of auto characteristics is independent of income. The first-order conditions for equation (5) imply

$$\Theta_i x_i^{\delta - 1} = \beta_i \exp(\beta' x + \alpha) \tag{6}$$

at an optimum x^*.

Equation (6) can be solved for the unobservable taste parameters in terms of observables. As in Feenstra (1986), it proves useful to do so. This yields

$$\Theta_i = (x_i^*)^{1-\delta}\, \beta_i \exp(\beta' x^* + \alpha). \tag{7}$$

I next define a consumer surplus function $S(x,x^*) = U(x,x^*) - p(x)$. This function gives the surplus associated with a model having characteristics vector x if the consumer's optimal choice is described by x^*.

Simple substitution gives

$$S(x,x^*) = \exp(\beta' x^* + \alpha)\, \Sigma \left[\frac{B_i}{\delta}\right] (x_i^* b)^{1-\delta}\, (x_i^{\delta} - 1) - \exp(\beta' x + \alpha). \tag{8}$$

It is easy to verify that S is maximized when $x = x^*$. This surplus function will serve as the metric for measuring distance in characteristics space.

Having defined the metric, I turn now to the task of using this metric to determine which products compete with one another, that is, which are neighbors. While there are many models of automobiles, and hence many available bundles of characteristics, there is not a continuum of products available on the market. Thus, a consumer may find that her optimal model, x^*, does not exist in the market. In this case, the consumer receives less surplus than she would if x^* had been available. In figure 2.2, I illustrate an iso-surplus contour for a typical consumer for the case of two characteristics. In the figure, $S(x,x^*)$ is constant along any contour and $S(x,x^*)$ decreases as one moves away from x^*. Thus, the consumer whose optimal characteristics bundle is x^* is indifferent between point A which entails slightly more horsepower and less weight and point B which gives relatively much more horsepower and a heavier auto.

Two models A and B would be neighbors if there is any consumer who is indifferent between A and B who prefers these two to all other models. Graphically, in figure 2.2, A and B would *not* be neighbors if there existed a model such as C.

Different consumers may have different ideal models. Because of this, there are many iso-surplus contours that will pass through any two models. In figure 2.3, individual 1 has an optimal choice of x^*_1, and A and B lie on the same iso-surplus contour—S_1. Another consumer, individual 2, has an optimal choice of x^*_2. For this consumer, A and B also lie on the same iso-surplus contour (S_2). The analogous story applies to consumer 3 whose optimal choice is x^*_3.

An ideal algorithm for determining neighbors would proceed in steps. For every possible pair of models in the sample, one would conduct a detailed grid search in characteristics space. At every point in the grid search, one would pose the following question: Is the consumer whose ideal model is this point in characteristics space indifferent to the two potential neighbors? If the answer is no, move on to the next point on the grid and repeat the question. If the answer is yes, ask if any of the other ninety-eight models in the sample give higher surplus than the

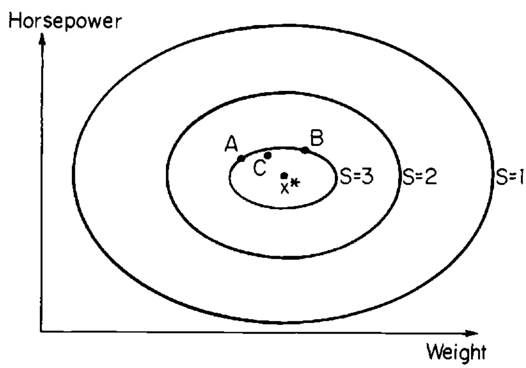

Fig. 2.2 Iso-surplus contours

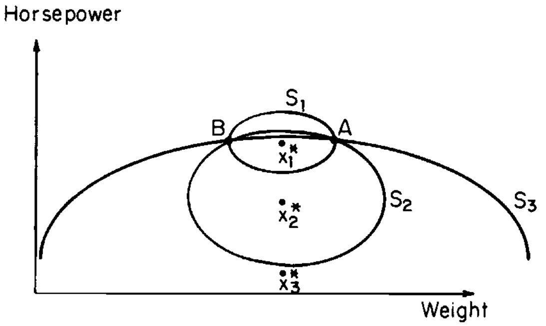

Fig. 2.3 An example of many consumers indifferent to two models A and B

pair being considered. If the answer here is no, the pair of potential neighbors are indeed neighbors.

This algorithm will determine which multidimensionally differentiated products are neighbors. Unfortunately, the algorithm is computationally infeasible for the case of automobiles. This is because I find that at least five characteristics are necessary to adequately account for differentiation between autos. The algorithm described above, then, would require many five-dimensional grid searches entailing many calculations at each point in each search. This is too expensive on a mainframe computer and too time-consuming on an advanced personal computer.

I refine the above definition of neighbors. (Two models were neighbors, it will be recalled, if there existed a consumer indifferent between them and who preferred them to all other available models.) Amending this definition allows me to derive a computationally feasible method for determining neighbors to each model in my sample. I take the *smallest* iso-surplus contour containing the potential neighbors as the basis for comparison. In figure 2.3, this is S_1—the surplus that consumer 1 obtains. This is akin to saying that it is the preferences of the consumer whose optimal bundle is most similar to the potential neighbors that, on the margin, matter. In figure 2.3, then, when I ask if A and B are neighbors, I use the preferences of consumer 1 and then look for a point such as C that lies within S_1. If a point such as C exists, A and B are not neighbors. This method is economically sound if it will always be the case that if consumer 1 has a model preferred to A and B, so will all other consumers. There are examples in which this will not be true, and this issue will be discussed in detail. First, though, it is convenient to state a working definition of "neighbors."

Definition: Models A and B are considered neighbors if, for the smallest iso-surplus contour containing both of them,

$$S(x_a,x^*)\ [=S\ (x_b,x^*)] > S(x_c,x^*) \text{ for all models } c.$$

This is, I believe, an economically intuitive and computationally straightforward definition of neighbors. It is not a perfect definition for at least two reasons. I discuss each in turn.

The first problem with the definition concerns identifying the x^* that defines the highest surplus associated with indifference between models A and B. Recall that x^* is a consumer's optimal choice of characteristics and as such is not observed. I posit that x^* is the midpoint of a line drawn between two potential neighbors, A and B, where the surplus function provides the metric. Since a model is represented by a vector of its characteristics, I find x^* by varying Ω from 0 to 1 until $x^* = \Omega x_A + (1-\Omega)x_B$ and $S(x_A,x^*) = S(x_B,x^*)$. If iso-surplus contours were proper ellipsoids, the x^* defined in the above linear fashion would indeed identify the smallest iso-surplus contour containing A and B. Insofar as the iso-surplus contours defined by equation (8) are not proper ellipses, defining x^* as the midsurplus point on the line between points A and B may not yield the smallest contour containing A and B.

There are two possible responses to this critique. First, the iso-surplus contours defined by equation (8) are, in fact, not too different from ellipses for the case of automobiles. Iso-surplus contours derived from data are drawn in weight-horsepower space in figure 2.4. Due to

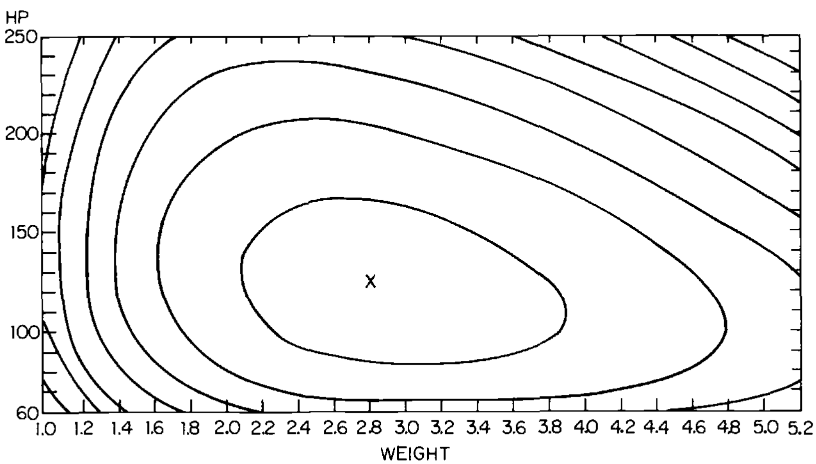

Fig. 2.4 Iso-surplus contours with Box-Cox CES utility and log-linear hedonic price function (weight versus HP)

the symmetry of equation (8), contours are similarly shaped in the shape of any two characteristics. Second, if x^* was poorly defined by drawing a line between A and B, one would expect the method to yield nonsensical sets of neighbors. I show in the next section that this is not the case.[12]

A second problem is that this definition of neighbors, which uses the smallest iso-surplus contour as the basis for comparison, may falsely reject potential neighbors. This is demonstrated in figure 2.5.

Suppose there are only three models, A, B, and C. My definition of neighbors rules out A and B as potential neighbors, since the iso-surplus contour drawn is the smallest containing A and B, and C is preferred to A and B. Yet for a consumer whose optimum is $x^{*\prime}$, A and B are neighbors. My method for determining neighbors, though, will never account for the preferences of a consumer with an optimal choice of $x^{*\prime}$ in figure 2.5. Because I find the optimum bundle by drawing a line between two models, and do so for all pairs in the sample, I will never account for the preferences of a consumer whose optimum bundle lies outside the outermost envelope of available models. The preferences of these consumers are ignored. In figure 2.5, this envelope is defined by the triangle ABC—an area that does not include $x^{*\prime}$.

For the automobile market, this problem is not likely to be an empirically important one because, in a market with as many models as the auto market, it is unlikely that there are many consumers whose ideal lies outside this outer envelope. Were this the case, one would expect such profitable market niches to be readily filled.

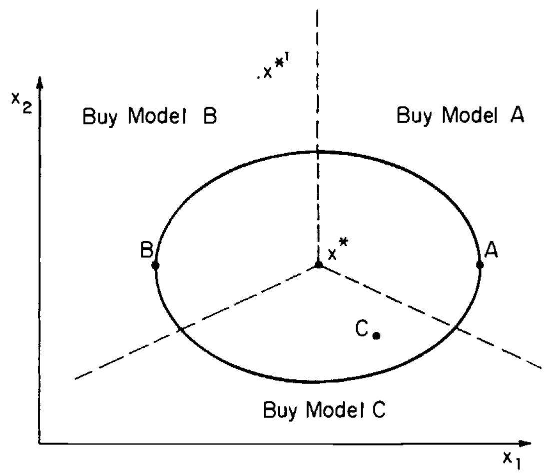

Fig. 2.5 An example of preferences for which the model will not account

The algorithm for finding neighbors, then, is as follows.

Step 1: Find x^* such that $S(x_1,x^*) = S(x_2,x^*)$ using the above-described linear method.

Step 2: See if there exists a model j not equal to 1,2 such that $S(x_1,x^*) < S(x_j,x^*)$. Models 1 and 2 are neighbors if no such j exists in the sample.

Step 3: Repeat the above two steps for all possible pairs in the sample. This algorithm ensures that if 1 is a neighbor of 2, then 2 is a neighbor of 1. If 3 is a neighbor to 2, though, it need not be a neighbor to (2's neighbor) 1. The number of neighbors a model has depends on its characteristics and the characteristics of the other models in the sample. The actual number of neighbors for each model is endogenous and will differ across models.

This procedure yields the neighbor(s) to every model in the sample. I use these neighbors as the elements of P_{jt} in the demand equation (1). Conversely, models that are not neighbors are assumed to have no cross-price effect in equation (1).

This concludes the description of the methodology. In this section, I have explained how I use results from a Lancasterian model of product differentiation to derive an estimable demand function. The resulting demand function circumvents many of the myriad problems that plagued earlier attempts to estimate the demand for differentiated products—specifically automobiles.

2.3 Data and Results

The data set comprises almost all automobile models sold in the calendar years 1983–85. Specialty models with annual sales of under 4000 were excluded (e.g., Ferrari and Rolls Royce). Models not produced for all of each of the three years were also deleted. This allows me to avoid the problems that would be posed by a model that is introduced in October and hence has very low annual sales for the calendar (as opposed to model) year. A similar, though less severe, problem would exist for models withdrawn after October. Models included in the sample are given in table 2.1. Each model/year observation consists of the following variables for the base model: (1) sales by Nameplate, (2) suggested retail list price, (3) wheelbase, (4) length, (5) width, (6) height, (7) weight, (8) headroom, (9) legroom, (10) number of engine cylinders, (11) engine displacement, (12) fuel injection or carburation, (13) manual or automatic transmission, (14) power or manual steering, (15) power or manual brakes, (16) air conditioning as standard, (17) horsepower (HP), (18) turning radius, and (19) country of origin. All data were collected from annual issues of *Automotive News Market Data Book*.

Table 2.1 **Models Used in the Sample**

Toyota Tercel	Saab 900 S	Buick Skyhawk
Toyota Corolla	Saab 900 Turbo	Buick Skylark
Toyota Celica	Porsche 944	Buick Century
Toyota Camry	Porsche 911	Buick Regal
Toyota Cressida	Isuzu I-mark	Buick LeSabre
Toyota Supra	Isuzu Impulse	Buick Electra
Nissan Sentra	Peugeot 505	Buick Riveria
Nissan Maxima	American Motors Alliance	Cadillac Cimarron
Nissan 300zx	American Motors Eagle	Cadillac Seville
Nissan 200SX	Plymouth Horizon	Cadillac DeVille
Nissan Stanza	Plymouth Turismo	Cadillac ElDorado
Nissan Pulsar	Plymouth Reliant	Chevrolet Chevette
Honda Accord	Plymouth GF	Chevrolet Cavalier
Honda Civic1.5S	Plymouth LeBaron	Chevrolet Citation
Mazda 626	Plymouth NewYorker/5thAvenue	Chevrolet Camaro
Mazda RX-7	Dodge Omni	Chevrolet Celebrity
Mazda GLC	Dodge Charger	Chevrolet Corvette
Subaru DL/GL	Dodge Aries	Chevrolet MonteCarlo
Chry/Ply Colt	Dodge Dodge600	Chevrolet Chevrolet
Volvo DL	Dodge Diplomat	Oldsmobile Firenza
Volvo 760 GLE	Ford EXP	Oldsmobile Cutlass/Sierra
VW Jetta	Ford Escort	Oldsmobile Cutlass/Supreme
VW Quantum	Ford Mustang	Oldsmobile Olds88
BMW 320/318	Ford T-Bird	Oldsmobile Olds98
BMW 530/528	Ford LTD	Oldsmobile Toronado
BMW 733	Ford CrownVictoria	Pontiac 1000
Mercedes 300D	Mercury Lynx	Pontiac Sunbird
Mercedes 300SD	Mercury Cougar/XR7	Pontiac Firebird
Mercedes 190E	Mercury Capri	Pontiac 6000
Audi 5000	Mercury Marquis	Pontiac Bonneville
Audi 4000	Mercury GrandMarquis	Pontiac GrandPrix
Mitsubishi Tredia	Lincoln Continental	Volkswagen Rabbit
Mitsubishi Cordia	Lincoln MarkVii	
Mitsubishi Starion	Lincoln Lincoln	

Some variables of economic significance are absent from the above list. In particular, data on the incomes of consumers and on the actual transaction prices are not available. The suggested list price of the base model for P_{it} is used.[13] This introduces systematic bias insofar as some models consistently sell for more or less than list price. For some Japanese models, this may have been the case in my sample.[14]

Neighbors are computed for the 1984 models. I assume that product characteristics do not change so much that neighbors change over the sample period; I will relax and test this assumption in future work. Indeed, computing neighbors for each year provides an alternative test of Feenstra's (1985) upgrading results. Here, differential upgrading would take the form of changing neighborhoods over time.

I begin by estimating the hedonic price equation $P(x)$. Like most researchers before me,[15] I find that the functional form of $P(x)$ that best fits the data is equation (4):[16]

$$P(x) = \exp(\alpha + \beta' x).$$

I find that a linear combination of the following five characteristics accounts for almost 90 percent of the variation of $P(x)$—weight, horsepower, and dummies for power steering, air conditioning, and foreign made. Dummy variables take the value of 2 if a car is foreign and if air and power steering are standard, and a value of 1 otherwise. This differs from the usual 1-0 convention because some dummies are raised to negative powers. The only effect of this change is to alter the constant term in the hedonic regression. Numerical experiments show that this has no effect on the determination of neighbors. I estimate the log of equation (4) to give

$$\begin{aligned} \ln P = {} & \underset{(.123)}{.215} + \underset{(.056)}{.209}\ \text{Weight} + \underset{(.0009)}{.0045}\ \text{HP} \\ & + \underset{(.052)}{.1261}\ \text{PS} + \underset{(.050)}{.4703}\ \text{Air} + \underset{(.044)}{.161}\ \text{Foreign}. \end{aligned} \tag{9}$$

Standard errors are in parentheses. 100 observations. $R_2 = .885$.

It is useful to view dummy variables here as proxies for various degrees of luxury and/or quality. Hence an optimal choice of characteristics, x^*, may involve .5 units of air conditioning. This just means that the consumer would prefer less luxury than is imposed by the all-or-nothing choice of air conditioning but more than is afforded by a no-air model. The coefficients in equation (9) are used to parameterize the surplus function of equation (8). While the coefficients are subject to measurement error, their very small standard errors argue that neglecting this error is unlikely to be an empirically relevant omission.

The only remaining unknown in the surplus function is the parameter δ which is related to the degree of concavity of the utility function. Recall that the elasticity of substitution, $\sigma = 1/(\delta - 1)$. This parameter is not identifiable with the data available. Following Feenstra (1986), I posit many different values for δ and replicate the entire methodology from the beginning for each of these. Fortunately, I find that the choice of δ over a wide range of plausible values does not affect the qualitative results. I consider values of $\delta = .5, -1, -3, -6$, and -8. Only at values of -8 and below do results change substantially. That is, the choice of neighbors is mostly unaffected until $\delta = -8$. At -8, neighbors become much more numerous and, to a degree, counterintuitive.

Once δ has been specified, I compute neighbors for every model using the 1984 data.[17] The results for $\delta = -3$ are given in table 2.2. Table 2.2, for example, tells us that the neighbors of the Honda Accord, model 13, are the Toyota Camry, Nissan Stanza, Mazda 626, Mitsubishi Tredia and Cordia, Chevrolet Cavalier, and Pontiac Sunbird. An intuitive way of interpreting table 2.2 is to note that it answers the question, What other autos did the consumer consider before she decided to purchase the one actually selected?

Table 2.2 **Neighbors of Models ($\delta = -3$)**

Model No.	Model Name	No. of Neighbors	Model Nos. of Neighbors
1	Toyota Tercel	6	17 18 19 42 44 80
2	Corolla	2	18 39
3	Celica	4	4 10 16 45
4	Camry	8	3 11 13 15 45 52 81 82
5	Cressida	4	6 21 37 67
6	Supra	3	5 9 37
7	Nissan Sentra	4	12 17 19 100
8	Maxima	3	25 29 36
9	300zx	3	6 26 38
10	200SX	8	3 11 16 20 24 31 41 53
11	Stanza	5	4 10 12 31 53
12	Pulsar	8	7 11 17 18 22 23 31 100
13	Honda Accord	5	4 15 32 33 81
14	Civic1.5S	1	100
15	Mazda 626	9	4 13 23 32 57 69 81 82 88
16	RX-7	4	3 10 20 34
17	GLC	7	7 12 1 18 19 56 100
18	Subaru DL/GL	8	2 12 17 1 22 51 61 94
19	Chry/Ply Colt	3	7 17 1
20	Volvo DL	10	10 16 24 30 34 35 41 54 62 98
21	760 GLE	4	5 25 27 37
22	VW Jetta	6	12 18 39 55 61 100
23	Quantum	5	12 15 31 40 76
24	BMW 320/318	6	10 20 31 35 40 76
25	530/528	4	8 21 27 92
26	733	5	9 28 38 68 85
27	Mercedes 300D	8	21 25 28 74 75 79 92 93
28	300SD	5	26 27 74 77 78
29	190E	2	8 35
30	Audi 5000	3	20 35 40
31	4000	6	10 11 12 23 24 76
32	Mitsubishi Tredia	5	13 15 33 81 95
33	Cordia	3	13 32 39
34	Starion	6	16 20 36 54 60 62
35	Saab 900 S	4	20 24 29 30
36	900 Turbo	3	8 34 37

Table 2.2 (continued)

Model No.	Model Name	No. of Neighbors	Model Nos. of Neighbors
37	Porsche 944	4	5 6 21 36
38	911	3	9 26 85
39	Isuzu I-mark	4	2 22 33 55
40	Impulse	5	23 24 30 41 76
41	Peugeot 505	7	10 20 40 43 59 83 98
42	Alliance	2	1 80
43	Eagle	5	41 59 76 87 99
44	Horizon	3	1 50 94
45	Turismo	4	3 4 46 53
46	Reliant	3	4 45 52
47	Plymouth GF	7	20 34 49 54 73 91 92
48	LeBaron	5	10 53 70 76 84
49	NewYorker/5thAvenue	5	47 60 73 75 92
50	Omni	3	44 1 51
51	Charger	2	18 50
52	Aries	4	4 46 53 82
53	Dodge600	6	10 11 45 48 52 70
54	Diplomat	7	20 34 47 58 73 91 92
55	EXP	4	22 39 61 95
56	Escort	5	17 18 22 61 80
57	Mustang	5	15 63 70 71 82
58	T-Bird	6	20 34 54 62 92 98
59	LTD	3	41 43 64
60	CrownVictoria	6	34 49 65 67 75 92
61	Lynx	7	17 18 22 55 56 80 94
62	Cougar/XR7	4	20 34 58 92
63	Capri	3	3 57 71
64	Marquis	3	59 76 83
65	GrandMarquis	3	60 67 75
66	Continental	4	7 68 77 79
67	MarkVii	4	5 65 66 79
68	Lincoln	4	26 66 78 85
69	Skyhawk	5	15 32 81 88 95
70	Skylark	6	48 53 57 71 82 84
71	Century	5	57 63 70 84 89
72	Regal	4	41 76 86 98
73	LeSabre	3	47 49 91
74	Electra	5	27 28 77 79 93
75	Riviera	5	27 49 65 92 93
76	Cimarron	11	23 24 31 40 43 48 83 92 96 97 98
77	Seville	5	28 66 74 78 79
78	Cadillac DV	3	28 68 77
79	ElDorado	7	27 28 66 67 74 77 93
80	Chevette	5	42 56 61 1 94
81	Cavalier	6	4 13 15 32 69 82
82	Citation	6	4 15 52 57 70 81
83	Camaro	5	41 64 76 96 98
84	Celebrity	5	48 70 71 76 89

Table 2.2 (continued)

Model No.	Model Name	No. of Neighbors	Model Nos. of Neighbors
85	Corvette	3	26 38 68
86	MonteCarlo	5	20 41 72 76 90
87	Chevrolet	4	43 91 92 99
88	Firenza	3	15 69 70 95
89	Cutlass/Sierra	3	71 84 97
90	Cutlass/Supreme	5	41 43 76 86 99
91	Olds88	4	54 73 87 92
92	Olds98	8	25 27 49 54 58 60 75 76
93	Toronado	4	27 74 75 79
94	1000	5	18 44 61 80 1
95	Sunbird	5	15 32 55 69 88
96	Firebird	3	76 83 97
97	6000	3	76 89 96
98	Bonneville	6	20 41 58 72 76 83
99	GrandPrix	4	41 43 87 90
100	Rabbit	5	7 12 14 17 22

In addition to varying δ, another type of sensitivity test was conducted in calculating neighbors. Because iso-surplus contours are not perfect ellipses, the linear method of finding the optimal model x^* is, as noted above, only an approximation. I used another approximation and retested for neighbors. This other approximation was based on finding x^* such that consumers whose ideal models were A and B were equally dissatisfied with x^*. This approximation yielded the same qualitative results as the linear approximation of x^*.

The next step in the methodology is to estimate the demand functions given in equations (1A) and (1B). Models have, on average, about 6 neighbors. With 100 models, this implies 600 cross-price terms to be estimated. While this is certainly an improvement over the previous 9,900 terms, the demand functions are still not accurately estimable with only 300 observations. I take the mean price of neighbors as the observation for P_{jt}. Similarly, I take the mean characteristics of neighbors as the observation for X_{jt}. Because the demand functions use the log of P_{jt}, it matters that the average of the logs is not the log of the averages. Numerical experiments show that this approximation does not affect results. There are other specifications for P_{jt}. Recall that the estimated demand equation is just a convenient statistical representation. Perhaps P_{jt} should be the average price of neighbors weighted by their sales. This representation of P_{jt} yields the same qualitative results, but standard errors on the parameters in the demand function are larger.

I estimate equations (1A) and (1B) using ordinary least squares (OLS). Because equation (1B) is nested within equation (1A), a straightforward F-test is used to test which specification should be used. That is, I test to see if own and neighbors' mean characteristics are jointly statistically significant.[18] For all values of δ tested, the data cannot reject the hypothesis that own and neighbors' characteristics are jointly insignificant.

The existence of multicollinearity in equation (1A) is confirmed by collinearity diagnostics following the approach of Belsley, Kuh, and Welsch (1980). Singular-value decomposition analysis indicates multicollinearity. The SVD analysis does not indicate that the data matrix is so ill conditioned as to suggest numerical error in the estimates. Due to the multicollinearity in equation (1A), estimated standard errors are inflated. This biases the F-test toward rejecting joint statistical significance of own and neighbors' mean characteristics. I nonetheless accept the results of the F-tests and use equation (1B) as the demand function in the analysis that follows. Table 2.7 presents the results of using instead equation (1A). As the collinearity diagnostics indicated, results are very similar to those obtained using equation (1B) (and given in table 2.4), except that standard errors are inflated.

OLS estimates of equation (1B) are presented in table 2.3. In table 2.3, equation 1 presents estimates of the demand function excluding any

Table 2.3 Estimated Automobile Demand Functions

Equation	(1)	(2) $\delta = .5$	(3) $\delta = -1.0$	(4) $\delta = -3.0$	(5) $\delta = -6.0$	(6) $\delta = -8.0$
CONSTANT	6.085	5.814	6.087	6.041	6.004	5.772
	(.278)	(.276)	(.277)	(.270)	(.267)	(.411)
LOGLIST	−.7942	−1.814	−1.333	−2.076	−2.271	−.912
	(.119)	(.254)	(.319)	(.313)	(.311)	(.165)
LOGPN		1.112	.522	1.250	1.444	.237
		(.246)	(.287)	(.284)	(.282)	(.290)
D84	.114	.111	.118	.122	.120	.108
	(.116)	(.113)	(.116)	(.113)	(.112)	(.116)
D85	.1672	.161	.173	.178	.179	.155
	(.117)	(.113)	(.116)	(.113)	(.112)	(.117)
SWEDE	−1.321	−1.350	−1.316	−1.228	−1.181	−1.302
	(.253)	(.245)	(.252)	(.246)	(.244)	(.254)
JAPAN	−.554	−.607	−.570	−.578	−.594	−.529
	(.117)	(.113)	(.116)	(.113)	(.112)	(.119)
GERMAN	−1.01	−.615	−.843	−.537	−.424	−.928
	(.117)	(.183)	(.189)	(.193)	(.196)	(.183)
R^2	.3954	.4349	.4022	.4330	.4450	.3976

Notes: Standard errors in parentheses. Dependent variable is LOGSALE. Variable definitions: LOGSALE = log of sales in 1000s; LOGLIST = log of the list price in \$1000s; LOGPN = log of the average price of the neighbors in \$000; D84 = 1 if the year is 1984, 0 otherwise; D85 = 1 if the year is 1985, 0 otherwise; SWEDE = 1 if the car is Swedish, 0 otherwise; JAPAN = 1 if the car is Japanese, 0 otherwise; and GERMAN = 1 if the car is German, 0 otherwise.

cross-price effects. This equation is roughly a panel data version of the older time-series studies that neglected cross-price effects. Equation 1 gives a highly significant total market elasticity of demand of $-.794$. This estimate is in line with existing, older estimates. Equation 1, though, is misspecified, as cross-price effects are omitted.

Equations 2–6 in table 2.3 give estimates when the demand function includes the price of neighbors, hence allowing for the possibility of substitution. Varying δ from .5 to -6 affects the significance of the parameters on own and neighbors' price, but the point estimates are fairly constant. (Recall that the choice of δ only enters the demand function via its effect on the determination of the set of neighbors.) For $\delta = .5, -1, -3$ and -6, the coefficient on the neighbors' price is highly significant. For these values of δ, the coefficient on own price is somewhat stable across equations and is highly significant.

For values of δ between -1 and -6, the total market elasticity $(\alpha_1 + \alpha_2)$ varies from $-.81$ to $-.83$, all of which are statistically significant at the 90 percent level. As theory would lead one to expect, allowing for substitutability leads to a more elastic own-price elasticity. This is evidenced by own-price elasticities (α_1) greater in absolute value than the coefficient of $-.794$ in equation 1.

In sum, the "neighbors" approach to restricting the dimensionality of the demand function in conjunction with a short panel of data seems to fit the data remarkably well. I have completed some sensitivity analyses in the spirit of Leamer (1985). These ad hoc specification tests include using other hedonic characteristics to control for cross-sectional variation. The results have been exceptionally robust to such tests.

2.4 On the Empirics of Taxation Schemes for Differentiated Products

The methodology by which the demand functions in table 2.3 were derived was based on Lancasterian consumer theory. That theory tells us that not all differentiated products need be substitutes. It also tells us to group products according to their characteristics and not only, as the Armington assumption implies, according to their country of origin.[19] The elasticities that are estimated in the equations of table 2.3, then, are the relevant ones from the vantage point of consumer theory.

Trade policy, though, typically taxes a good based on its country of origin. The analysis of trade policy issues requires trade elasticities. I derive these elasticities from the estimates of the demand system provided in sections 2.2 and 2.3. This is accomplished by perturbing the system on whatever margin trade policy operates to simulate the elasticity relevant to the study of trade taxes. This approach is more likely to give valid elasticities than direct estimation of import demand equations (see, for example Leamer and Stern 1970) because it is based on a utility-consistent framework for demand.

Suppose, for example, that policymakers wish to know how the demand for domestic autos changes when a tariff is applied to *all* auto imports. To derive this elasticity, I increase the price of all foreign cars by one percent—my proxy for a small change. This increases the demand for models of domestic autos that have as neighbors some foreign model. Summing the new demand for all domestic autos gives the information needed to construct the relevant elasticity.

This approach requires a caveat. I have nothing to say about the effect of large taxes. This is because the estimated demand system is only a local representation of demand. The system may behave quite differently at a point far from the initial situation. This is a standard warning in the empirical tax analysis literature. Also, here, large taxes may change the neighbors of a model. I assume that the taxes I consider are small enough that neighbors do not change. Preliminary numerical experiments indicate that this is indeed the case for the one percent price changes I consider.

In table 2.4, I give a wide variety of elasticities corresponding to various policy scenarios. For each elasticity, I also give its standard error. This statistic is computable given the variance-covariance matrix of the estimates of the initial demand equation. These elasticities all are simulated using the demand equation 4 found in table 2.3. That is, δ from

Table 2.4 **Elasticities of Demand by Country of Origin Using Equation (1B) ($\delta = -3$)**

	Quantity Change				
Price Change	Domestic Autos	All Imports	Japanese Imports	German Imports	Swedish Imports
All domestic	−1.187	.225	.213	.258	.076
	(.146)	(.051)	(.048)	(.058)	(.017)
All foreign	.367	−1.045	−1.030	−1.078	−.897
	(.084)	(.129)	(.128)	(.132)	(.118)
Japanese	.187	−.663	−1.43	.393	.300
	(.042)	(.081)	(.187)	(.089)	(.068)
German	.112	−.279	.317	−1.717	.745
	(.025)	(.036)	(.072)	(.240)	(.169)
Swedish	.024	−.064	.071	.247	−1.971
	(.005)	(.011)	(.016)	(.056)	(.292)
All foreign weighing < 2300 lbs.	.096	−.376	−.550	−.199	0.0
	(.021)	(.046)	(.067)	(.025)	—

Note: Standard errors in parentheses.

the utility function is set to -3. I take this as a central case for expositional purposes. Tables 2.5–2.7 present the same elasticities when the entire methodology is conducted using other values of δ. Those tables show that results remain qualitatively similar for a range of δ's.

Table 2.4 is easily interpreted. The table shows, for example, that the elasticity of demand for domestically produced automobiles with respect to the price of Japanese autos is .187. That is, a one percent increase in the price of all Japanese cars (via a tariff, perhaps) yields a .187 percent increase in demand for domestically produced autos. Were such a price increase applied to all imported autos, demand for domestically produced autos would rise instead by .367 percent. This example illustrates an error present in earlier studies of U.S.-Japanese auto trade policy. These studies used an imputed elasticity of demand for domestic autos with respect to a *foreign* price change because there were no estimates available of elasticities of domestic demand with respect to a change in only the Japanese price. Table 2.4 tells us that this error leads one to believe that demand for domestic autos is *twice* as responsive to a small tariff on Japanese cars as is actually the case. The difference arises due to substitution by American consumers away from Japanese cars toward other foreign cars not affected by the trade policy. These results serve as another indication of the ineffectiveness of selective protection.

Suppose that the purpose of trade or industrial policy in the U.S. automobile industry is to increase demand for domestically produced autos. Table 2.4 shows that a tax on all imports has less than half the effect on domestic demand that an equal subsidy on domestic models would have (.367 versus -1.187). (Consequences for government revenue are, of course, quite different.) An increase in a tariff on Swedish autos has very little effect on domestic demand; the relevant elasticity is .034. This is because most of the neighbors to Swedish autos are also foreign.

Suppose that the purpose of trade taxes is to reduce imports from a specific country. Then table 2.4 shows that a tax on only Swedish cars reduces Swedish imports by relatively less than the same tariff on German autos. Swedish cars are the most elastically demanded import, followed by German models, then Japanese models (-1.97 versus -1.71 versus -1.43). This is because Japanese models have many Japanese neighbors, while this is not the case for Swedish models. Indeed, most neighbors to Swedish models are German, as evidenced by the relatively high cross-price elasticities between German and Swedish autos.

Perhaps contrary to prior beliefs, a tax on all imports would have roughly the same relative impact on Japanese, German, and Swedish producers.

Some economists have argued for a tax on all small foreign cars instead of a tax on Japanese autos. Such a tax would not discriminate on the basis of country of origin and is viewed more kindly by GATT. I arbitrarily define small cars to be those weighing under 2,300 pounds. (For purposes of comparison, a Toyota Tercel weighs 1,985 lbs., a Honda Accord 2,187 lbs., and a Saab 900 2,612 lbs.) While such a broadly based tax might make the U.S. trade representative's job easier, the policy is only half (.096 versus .187) as effective at increasing demand for domestically produced autos as a direct tax on all imports. Swedish producers are totally unaffected by such a tax since no Swedish export to the United States weighs less than 2,300 pounds (there is a reason Volvos are so safe), and no Swedish car has a neighbor weighing less than 2,300 pounds.

It is possible to investigate the effects of various other trade and industrial policies using table 2.4. The above scenarios provide only a beginning.

2.5 Summary

This chapter has developed a new methodology for investigating empirically the effects of taxes on differentiated products. The approach adopted a Lancasterian, utility-consistent view of product differentiation. Using this approach, I calculated which multidimensionally differentiated products were neighbors. This information proved a useful basis for decreasing the dimensionality of the demand-estimation problem. Using a panel of one hundred automobile models over three years, a demand function was estimated. This yielded quite reasonable and statistically significant demand elasticities.

Recognizing that tax policy often acts on a different margin than consumer theory, the demand elasticities necessary for tax policy analyses were simulated. This provided the first estimated set of such elasticities. These elasticities provide some insight into a number of possible policy scenarios.

The methodology developed in this chapter provides ample opportunities for Leamer-type ad hoc specification tests. Results using different elasticities of substitution are presented in tables 2.5–2.7. These results appear robust.

The elasticities estimated and given in table 2.4 are well suited to simulation analyses of strategic trade and industrial policies concerning the U.S. automobile industry. This is the subject of ongoing research.

Table 2.5 **Elasticities of Demand by Country of Origin Using Equation (1B) ($\delta = -1$)**

Price Change	Quantity Change: Domestic Autos	All Imports	Japanese Imports	German Imports	Swedish Imports
All domestic	−.967	.112	.122	.100	.053
	(.154)	(.062)	(.067)	(.055)	(.029)
All foreign	.162	−.918	−.928	−.906	−.859
	(.089)	(.137)	(.140)	(.134)	(.124)
Japanese	.086	−.559	−1.080	.147	.108
	(.047)	(.089)	(.198)	(.081)	(.059)
German	.042	−.259	.120	−1.173	.343
	(.023)	(.037)	(.066)	(.241)	(.189)
Swedish	.015	−.078	.026	.101	−1.326
	(.008)	(.011)	(.014)	(.056)	(.318)
All foreign weighing < 2300 lbs.	.051	−.338	−.491	−.188	0.0
	(.028)	(.052)	(.076)	(.028)	—

Note: Standard errors in parentheses.

Table 2.6 **Elasticities of Demand by Country of Origin Using Equation (1B) ($\delta = -6$)**

Price Change	Quantity Change: Domestic Autos	All Imports	Japanese Imports	German Imports	Swedish Imports
All domestic	−1.247	.239	.231	.267	.100
	(.083)	(.046)	(.045)	(.052)	(.019)
All foreign	.426	−1.060	−1.052	−1.088	−.922
	(.145)	(.125)	(.125)	(.128)	(.117)
Japanese	.226	−.709	−1.568	.466	.361
	(.044)	(.082)	(.190)	(.091)	(.070)
German	.124	−.243	.427	−1.849	.821
	(.024)	(.035)	(.083)	(.236)	(.160)
Swedish	.026	−.074	.071	.269	−2.144
	(.005)	(.010)	(.014)	(.052)	(.288)
All foreign weighing < 2300 lbs.	.119	−.390	−.581	−.189	0.0
	(.023)	(.045)	(.067)	(.024)	—

Note: Standard errors in parentheses.

Table 2.7 **Elasticities of Demand by Country of Origin Using Equation (1A) ($\delta = -3$)**

Price Change	Quantity Change: Domestic Autos	All Imports	Japanese Imports	German Imports	Swedish Imports
All domestic	−1.412	.201	.191	.230	.067
	(.533)	(.103)	(.584)	(.118)	(.034)
All foreign	.328	−1.285	−1.275	−1.314	−1.152
	(.168)	(.580)	(.098)	(.569)	(.633)
Japanese	.167	−.793	−1.636	.352	.268
	(.085)	(.309)	(.046)	(.180)	(.137)
German	.100	−.358	.284	−1.886	.666
	(.056)	(.195)	(.145)	(.405)	(.342)
Swedish	.021	−.091	.064	.221	−2.113
	(.011)	(.068)	(.033)	(.113)	(.385)
All foreign weighing < 2300 lbs.	.086	−.462	−.672	−.250	0.0
	(.044)	(.206)	(.293)	(.124)	—

Note: Standard errors in parentheses.

Notes

The author is grateful to Angus Deaton, Avinash Dixit, Rob Feenstra, Gene Grossman, Whitney Newey, and Duncan Thomas for helpful comments and suggestions, as well as to the Sloan Foundation and the National Bureau of Economic Research for financial support.

1. See Dixit 1987, which argues this point.
2. See, for example, Deaton and Muellbauer 1980.
3. Other simpler examples of this type of methodology are Johnson 1978 and Cragg and Uhler 1970.
4. Demand for automobiles is the most prevalent example of modeling the demand-differentiated products. I am unaware of any modeling approach for other differentiated products that is not mentioned in this section of the chapter.
5. This is the most recent data available until April 1987.
6. Note that this differs from the usual panel in which goods are the same, but demand is across consumers and over time. Here, the consumers are assumed the same, but goods differ across models, and these models are tracked over time.
7. Actually, it is sufficient to include in the regression those characteristics of which a linear combination accounts for the product differentiation.
8. All foreign models except the Peugeot 505 are produced by either Japan, Germany, or Sweden.
9. This differs from the Dixit-Stiglitz approach to product differentiation. There, all products are neighbors.

10. The approach I use to find neighbors when products are multidimensionally differentiated benefited greatly from discussions with Rob Feenstra. I am very grateful for his many helpful suggestions.

Recent theoretical work by Caplin and Nalebuff (1986) has also addressed the issue of determining neighbors to a good when products are multidimensionally differentiated. They show that if preferences can be represented by a utility function that is Cobb-Douglas in product characteristics and income, there exists a straightforward way of finding neighbors. Using the unit simplex in Cobb-Douglas parameter space, they show that a hyperplane divides all consumers who prefer good x to good y from those who prefer y to x. Because sets of consumers who prefer one model to another (that is, the model's neighborhood) are defined by hyperplanes, finding neighbors is a tractable problem. The tractability comes from the functional form of the utility function. While this is an elegant result, it is not applicable to the automobile market because the utility function that permits the tractability of the problem also implies that all consumers purchase the same value of the most preferred model but differ in quantities purchased. For big-ticket items such as automobiles, this is just not the case.

11. This function over characteristics is sometimes referred to in the literature as a subutility function.

12. While this line of reasoning has strong Bayesian overtones, I do not know another way of getting a feel for the validity of a new methodology. This is another reason why the auto industry is a good candidate to which to apply a new methodology. If my methodology were first applied to lumber and I found clear pine-2 to be a neighbor to grade 3 birch, few economists would have any idea of how well neighbors are defined.

13. This is also the practice adopted by Feenstra 1985. In that paper, he puts forth the argument that for national welfare considerations, dealer markups represent an intracountry transfer.

14. Implicit discounts due to selectively applied low financing rates have also been ignored due to lack of data.

15. The most recent examples are Feenstra 1985 and 1986. Griliches 1971 is a much earlier example.

16. I also estimate this function without logarithms. This functional form yielded a loss of about .20 in the R^2.

17. This procedure is programmed in IBM Profortran for implementation on IBM-compatible personal computers.

18. Throughout this chapter, *statistically significant* means statistically significant from zero at the 90 percent confidence level unless stated otherwise.

Because *FOREIGN* is a near-linear combination of *SWEDE, JAPAN*, and *GERMAN*—the fixed effects, I do not include *FOREIGN* in equation (1A) as an own characteristic.

19. Indeed, demand estimation according to the Armington assumption, using my data set, yields statistically insignificant and nonsensical demand elasticities.

References

Automotive news market data book. Various annual issues. Chicago: Crain Communications.

Belsley, D. A., E. Kuh, and R. E. Welsch. 1980. *Regression diagnostics*. New York: John Wiley.

Bresnahan, T. 1981. Departures from marginal-cost pricing in the American automobile industry. *Journal of Econometrics* 17:201–27.
Caplin, A., and B. Nalebuff. 1986. Multi-dimensional product differentiation and price competition. Mimeo. Princeton University.
Cragg, J., and R. Uhler. 1970. The demand for automobiles. *Canadian Journal of Economics* 3:386–406.
Daganzo, Carlos. 1979. *Multinomial probit: The theory and its applications to demand forecasting*. New York: Academic Press.
Deaton A., and J. Muellbauer. 1980. An almost ideal demand system. *American Economic Review* 70:312–26.
Dixit, A. In press. Optimal trade and industrial policies for the U.S. automobile industry. In R. Feenstra, ed., *Empirical methods for international trade*. Cambridge: MIT Press.
Dixit, A., and J. Stiglitz. 1977. Monopolistic competition and optimum product diversity. *American Economic Review* 67:297–308.
Feenstra, R. 1985. Quality change under trade restraint: Theory and evidence from Japanese autos. Columbia University Department of Economics Paper No. 298.
______. 1987. Gains from trade in differentiated products: Japanese compact trucks. In R. Feenstra, ed., *Empirical methods for international trade*. Cambridge: MIT Press.
Griliches, Z. 1971. Hedonic price indices for automobiles. In Z. Griliches, ed., *Price indexes and quality changes*, 55–87. Cambridge: Harvard University Press.
Helpman, E., and P. Krugman. 1985. *Market structure and foreign trade: Increasing returns, imperfect competition, and the international economy*. Cambridge: MIT Press.
Hotelling, H. 1929. Stability in competition. *Economic Journal* 39:41–57.
Johnson, T. 1978. A cross-section analysis of the demand for new and used automobiles in the United States. *Economic Inquiry* 16:531–48.
Krugman, P. 1980. Scale economies, product differentiation, and the pattern of trade. *American Economic Review* 70:950–59.
Lancaster, Kelvin. 1979. *Variety, equity, and efficiency*. New York: Columbia University Press.
Leamer, E. 1984. Model choice and specification analysis. In Z. Griliches and M. Intriligator, eds., *Handbook of econometrics*, 1:285–330. Amsterdam: North-Holland.
______. 1985. Sensitivity analyses would help. *American Economic Review* 75:308–13.
Leamer, E., and R. Stern. 1970. *Quantitative international economics*. New York: Aldine.
Ohta, M., and Z. Griliches. 1975. Automobile prices revisited: Extensions of the hedonic hypothesis. In N. E. Terleckj, ed., *Household production and consumption*, 325–98. New York: Columbia University Press.
Salop, Steven. 1979. Monopolistic competition with outside goods. *Bell Journal of Economics* 10(1):141–56.
Suits, D. 1958. The demand for automobiles in the United States, 1929–1956. *Review of Economics and Statistics* 40:273–80.
______. 1961. Exploring alternative formulations of automobile demand. *Review of Economics and Statistics* 43:66–69.
Tarr, D. G., and M. E. Morkre, 1984. *Aggregate costs to the United States of tariffs and quotas on imports*. Bureau of Economics Staff Report. Washington, D.C.: Federal Trade Commission.

Toder, E., with N. Cardell and E. Burton. 1978. *Trade policy and the U.S. automobile industry*. New York: Praeger.
Train, K. 1986. *Qualitative choice analysis: Theory, econometrics, and an application to automobile demand*. Cambridge: MIT Press.

Comment Raymond Riezman

The trade policy issue addressed in this chapter is how the existence of differentiated products alters our approach to trade policy. Levinsohn outlines a procedure for estimating elasticities for different models of automobiles purchased, but not necessarily produced, in the United States. The Lancasterian characteristics approach is used to reduce the number of cross elasticities considered, thereby making these estimates tractable.

The estimates are then used to conduct a number of interesting policy experiments. For example, if the goal of policy is to increase demand for domestically produced autos, a direct production subsidy is roughly twice as effective as a tariff. This result is an empirical verification of the theme that emerges from the domestic distortions literature, namely, that a policy should be focused as directly as possible at the source of the problem. This finding goes beyond the usual one in which direct subsidies are superior to tariffs because they involve fewer distortions. Here, we also have to worry about the neighbor effects, that is, the extent to which protection increases demand for domestically produced autos depends on whether the banned imports are neighbors of domestically produced autos or other imports.

Generally, it seems that the effectiveness of any trade policy is dependent on what is a neighbor to the good to which the policy is applied. For example, if the goal is to increase domestic production, then taxing imports will work only to the extent that domestic products are neighbors to the imports. If the goal is to reduce the trade deficit, taxing Japanese imports works if these imports are neighbors to U.S.-produced products but will not work if their neighbors are other imported goods. This chapter provides us with a methodology to empirically determine which goods are neighbors.

The basic equation to be estimated (1A) has quantity of a particular model in a certain year as the dependent variable. Independent variables are own price, own characteristics, characteristics and prices of all other models in all other years, time and country dummies. A re-

Raymond Riezman is associate professor of economics at the University of Iowa.

stricted version (1B) which eliminates direct effects of the characteristics is also estimated.

Estimating equation (1A) directly is not tractable. A procedure is developed to eliminate many cross-characteristic effects. First, a Lancasterian utility function is defined. A surplus function is defined which determines a metric in characteristic space. Given this metric, a procedure is outlined to determine which models are neighbors. Nonneighbor models are assumed to have zero cross effects. Once neighbors are determined, equations (1A) and (1B) are estimated. I found the discussion of this procedure difficult to follow. It is particularly puzzling to understand how the heterogeneous tastes are aggregated. Another result needing elaboration is that equation (1B) performs as well as equation (1A).

While it is true, as the author states, that all of these estimates are for "small changes" in tariffs, an interesting extension would be to do simulations to see how the estimates change for finite changes in tariffs. It appears that these estimates would be robust until neighbor relationships start to change. It would be of interest to know how much tariffs can change without substantially altering neighbors. This extension is important since ultimately, the usefulness of this methodology depends on being able to extend the results to finite changes in tariffs.

Data are available on average transaction prices, from consumer publications. It is important to get this information because casual observation suggests that biases exist, namely, Japanese cars sell at list or above, while U.S. autos sell at discounts up to 15 percent.

This chapter has relevance for the Baldwin-Green essay (chap. 7, in this volume) which is concerned with why protection does not always stimulate domestic production. The effect on domestic production of any protectionist policy depends on which goods are neighbors to the protected goods.

Since Levinsohn is concerned with the demand side only, there is little discussion about the industry structure. However, if the industry is competitive it seems that existing methods would work for an appropriately defined characteristic space. If the industry is not competitive, this needs some discussion since there is no general characterization of equilibrium in this kind of model and identification problems seem serious because these noncompetitive firms presumably react to things, such as taste changes, that are causing the demand to change.

Sometimes work such as this is misinterpreted as advocating or encouraging the use of protection. Such a view is quite misguided in this case. In the short run, policymakers operate under many political and institutional constraints. Policy studies such as this one are useful for helping policymakers optimize under the given constraints. Of course, we should also be thinking about long-run issues, but there still is a place for analysis that is useful for current policy problems.

Comment James E. Anderson

This chapter attempts a significant step forward in the estimation of detailed cross-price effects in a demand system for a differentiated products industry. Such cross effects are badly needed for policy analysis, yet very difficult to obtain. One problem is that for an industry with many models, the number of cross effects outruns any possible econometric effort to catch them. James Levinsohn's answer is to develop a method for imposing zero restrictions that is not purely a priori. I think the technique is clever and will undoubtedly be used in future related work. The possible objections to it have been anticipated and discussed by Levinsohn. The zero restrictions are then used to implement a standard type of aggregate demand system estimation. It is here that Levinsohn runs into some significant difficulties.

I have two sets of comments; one on the paper written and one on the paper I think the author should attempt to write. The latter is of course unfair, but I am going to do it anyway, because the problems I raise are significant and tend to vitiate the type of conclusion possible based on the current research strategy. Essentially they involve trying to model the supply side of the auto market.

Below, I critique Levinsohn's demand analysis and then make a case for a simultaneous attempt to model the supply side.

Critique of the Demand Study

Even leaving aside the identification issues of the next part, it is doubtful that Levinsohn's study has pinned down a stable and accurate demand system. I see two major flaws—the failure to treat aggregation issues and the failure to consider the time dimension of the demand for a durable good. I also have one minor comment.

Aggregation

The first flaw concerns aggregation. An aggregate price elasticity set is econometrically estimated, and its values depend on both the underlying utility structure and the aggregation structure. The author develops only the former. The problem is that a great deal of the "action" in determining the elasticities is in the latter. I have little confidence that the estimated elasticities will be accurate when considering policy-induced perturbations.

The difficulty is clearest with the aggregate price elasticity of demand for autos as a group, implied by Levinsohn's estimated set. The underlying consumer theory models an agent who will buy either zero or

James E. Anderson is professor of economics at Boston College.

one automobile. Levinsohn's price elasticity of demand for autos is zero everywhere but at the jump point, where it is undefined. So also in the aggregate: if all consumers are located on one or the other side of the jump point, the elasticity of demand is zero. In the sample data, price movements evidently push some consumers across the jump point, allowing a nonzero elasticity. But all the action is in the distribution of consumers. Without an aggregation theory, there is no way to know how accurate the estimated elasticity is.

Turning to the estimated cross-price elasticities that are the main focus of the chapter, the same difficulty exists. The consumer theory of Levinsohn's program (5) does not define any cross-price effects in autos. The analog to the usual elasticities is the elasticity of demand for a characteristic with respect to a marginal cost p_x. Each model bundles together a given quantity for each characteristic. A rise in the price of a model (proportionately raising all p_x), given an interior solution in program (5), should induce a jump to a new model: an undefined cross-price effect. Realistically, not all possible models are offered, so program (5) has further constraints on it. This means that the effect of a rise in price of model j may not induce a jump in demand for its close substitute model k: a zero elasticity. Aggregation again has all the action in averaging together the two effects. Without a structure for aggregation, there can be little confidence in the results.

The Time Dimension

The second flaw in the demand study is a failure to treat the time dimension of the demand for autos. There are three aspects of this to consider. First, durability and frequency-of-repair data are available and belong in the list of characteristics. There is ample evidence that consumers care about such properties and that manufacturers attempt to respond.

Second, the demand for a durable good must be modeled in an intertemporal setting. This will have a number of implications. For example, initial stocks of autos may be important (do not necessarily disappear into a wealth term). Second, a well-known effect of a rise in the price of a durable is that it shifts up the duration of use. There have been significant changes in the average vintage of U.S. autos in the past fifteen years. Also, such data vary by model. There is no easy way to stuff such phenomena into the author's static model. Presumably a start is for the taste for durability to rise with price, imposing a nonlinearity on program (5), but it is better to confront the problem.

Third, it seems to me that the closest substitute to the cars an agent might buy may often be the car he already owns. There are two aspects to this. First, the choice in program (5) and in the demand system also includes use of the depreciated current auto. Second, knowledge of characteristics is imperfect, and the currently owned auto is the best source of information on the potential new auto from the same source. Manufacturers and dealers certainly act on that belief in their efforts to build and maintain brand loyalty. This suggests another role for the existing stock of autos in demand estimation, standing in for reputational effects.

Minor Comment

The author's technique for identifying close substitutes is the major contribution of the chapter. It would be useful to check this technique with the results of market survey research. One concern of that literature is also the identification of close substitutes, proceeding on the basis of sampling techniques.

The Case for a Supply-Side Model

A basic assumption of the study is that the supply side of the auto market offers a technically fixed price-quality mix to consumers. See program (5), where $p(x)$ is given. Simultaneous equations bias arises when this is not true. Also, policy conclusions based on a tariff levied assuming an unchanging supply side will be falsified by shifts in manufacturers' behavior.

It is idle to complain about identification problems in the absence of a demonstration that they are significant and that feasible methods exist for remedying the situation. The evidence is that for autos these issues are highly significant. Folk wisdom has it that Japanese manufacturers shifted their quality mix in response to the VERs (Voluntary Export Restraints). Robert Feenstra has done some work along these lines suggesting that a large portion of the rise in price was in fact explained by the shift in quality.

It is more dubious whether feasible methods exist for estimating the full structure. Sherwin Rosen (1974) has a full treatment of the structure of perfectly competitive differentiated product markets and some suggestions for estimation. He comments that the problem of monopolistically competitive markets, with discrete differences in brands, is significantly more difficult. I suggest that the sort of cleverness and willingness to heroically simplify that is visible in the author's technique for identifying close substitutes should also be spent on the simultaneous system problem.

Reference

Rosen, Sherwin. 1974. Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55.