International trade and the U.S. automobile industry: current research, issues, and questions

James Levinsohn *

Department of Economics, University of Michigan, Ann Arbor, MI 48109, USA

Abstract

This paper has dual goals. On the one hand, it attempts a review of recent econometric work on the U.S. automobile industry. Three studies of this industry are reviewed, compared, and contrasted. These are studies by Koujianou Goldberg [Koujianou Goldberg, P., 1993, U.S. trade policy implications of a new approach to modelling the demand for automobiles, mimeo (Princeton University)], Feenstra and Levinsohn [Feenstra, R. and J. Levinsohn, 1994, Estimating markups and market conduct with multidimensional product attributes, Review of Economic Studies, forthcoming] and Berry et al. [Berry, S., J. Levinsohn and A. Pakes, 1994, Automobile prices in market equilibrium, Econometrica, forthcoming]. On the other hand, a number of current policy issues facing the industry are also reviewed. These issues include the Voluntary Export Restraint applying to exports from Japan, the NAFTA, anti-dumping law, and industry globalization. The paper provides a broad overview, then, of econometric and policy issues as they relate to the U.S. automobile industry.

Keywords: International trade policy; Automobiles

JEL classification: F1; L9

1. Introduction

One of the rewards of researching the U.S. automobile industry is that there is seldom a lack of interesting questions. The industry has a way of staying in the news. This may be due to the sheer size of the industry. It is America's largest manufacturing industry with annual sales topping $150 billion in 1989. In that year, automobiles were also at the top of the list of categories of imports by value. With domestic production so large and imports also so significant, it is perhaps inevitable that the automobile industry dominates many discussions of U.S. trade, environmental, and industrial policy.

This paper has dual goals. On the one hand, it provides a summary of recent econometric work on the U.S. automobile industry. The last few years have seen at least three major studies of the industry and these will be reviewed, compared, and contrasted. On the other hand, a number of current policy issues face the industry and these too will be reviewed. These two goals are unfortunately less well integrated than one might hope. In an ideal world, econometric work would address the important current policy issues. This is not the case, though, as the econometric studies have each involved years of research while the policy world moves much more quickly.

The organization of the paper is as follows. In Section 2 I review three large econometric studies of the U.S. automobile industry. These are studies by Koujianou Goldberg (1992), by Feenstra and Levinsohn (1994), and by Berry et al. (1994). These studies address some policy issues and these are briefly discussed. In the end, though, there are several current policy issues that are either not addressed at all or not adequately addressed in these studies. The third section of the paper discusses modelling issues that arise when considering export restraints and the North American Free Trade Agreement. Some other policy issues confronting the automobile industry are also discussed. Conclusions and are gathered in Section 4 as well as suggested topics for future research.

2. Recent econometric studies of the U.S. automobile industry

One common theme that binds the studies by Koujianou Goldberg (KG), Feenstra and Levinsohn (FL), and Berry, Levinsohn and Pakes (BLP) is their focus on estimating equations at the product or nameplate level. This is in contrast to more aggregate approaches that attempt to model the total demand for autos. The aggregate approach only requires one elasticity of demand for an average automobile. As such, cross-price elasticities between *types* of automobiles are a non-issue. The studies by KG, FL, and BLP, on the other hand, devote much of their intellectual energy to estimating the

many own-price and cross-price elasticities that arise when one uses product-level data. For example, with about 150 nameplates sold in the U.S. in 1990, there are over 22,000 elasticities to be estimated for 1990 alone. With sufficient data, one could, in principle, estimate all these elasticities for every year. In practice, this is not possible due to data limitations. Because data limitations play a large role in determining the methods used, it is worthwhile to discuss some of these limitations.

To estimate demand elasticities for automobiles, one typically requires data on the products available and on the sales of those products. The source of this data is almost always the Automotive News Market Data Book – an annual publication that lists the engineering characteristics of all products offered in a model year and the sales of all products during the previous year. [1] For many of these characteristics, this data is pretty good. Examples include engine specifications, vehicle weight, wheelbase, legroom, and the like. These data are typically given for the base model. In two cases, though, the data are problematic, and this is the case with price and sales. The price given is suggested list price for the base model. This often differs from the price the consumer actually pays, the transaction price. There is no simple solution to this problem. One approach is to assume that discounts vary according to vehicle classes and just posit the discount. Another approach is to just use the list price. Proprietary transactions prices data exist, but these are difficult to obtain and even then they do not solve the problem. Usually, the transaction price *exceeds* the base model list price, presumably due to the purchase of vehicle options that are added onto the base price. Unless one has both the transaction price *and* vehicle characteristics of the actual car bought, there is still a problem. Sales figures are also problematic. This is because many vehicles are leased or bought for fleets. The typical approach is to assume that either leasing or fleet sales are representative of the average end-use consumer, and hence use these sales in the demand estimation. Nonetheless, at the nameplate level, data, while decent, are not perfect. Researchers in this field typically note the problems and do their best with what is available.

A source of consumer-level data is found in the Consumer Expenditure Survey (CES). This is a consumer survey which asks, among many other things, what car the household purchased (if any) during the year of the survey. This is helpful data, but not nearly enough to get at the over twenty thousand elasticities since it surveys about 6000 households per year and only about 10 percent or 600 of these buy a car in any given year. With over 150 models, this gives an average of 4 observations per model, but many models in fact have no observed purchases in this data set.

---

[1] For some years, this data is available from the author on request. Send all requests by e-mail.

One other source of data is the U.S. consumer magazine "Consumer Reports." This publication publishes a widely read reliability rating for almost every model. This is updated annually. While this data is possibly useful in the cross-section, it is not nearly as useful if used over multiple years as in a time-series. This is because the data given by a scale of 1 to 5 and a 1 in one year is not comparable to a 1 in another year. Because cars have become much more reliable over the last few years, the absolute difference between the highest and lowest scores has narrowed a great deal. Using this data in a time-series, then, may be problematic.

With the data that *is* available, one then tackles the problem of estimating demand elasticities. This is a non-trivial job given the sparseness of available data, but it is a crucially important task if one wishes to undertake many types of policy analysis.

Demand elasticities determine not only how consumers will respond to changes in prices, they also play a key role on the supply side of the model. Each of the three studies model the U.S. automobile market as imperfectly competitive. As such, demand elasticities enter firms' pricing decisions. The markup of price over marginal cost, a key input to firm profits, is a function of these elasticities. Hence both the consumer and firm side of the industry are effected by demand elasticities. In equilibrium, each firm chooses its prices to maximize profits, given its own costs, the demands for it products, and the prices (or quantities) of rival products.

While demand elasticities are at the heart of all of the three studies, each confronts the problem of how to estimate so many parameters from so little data from a different perspective. The FL study is quite different from the other two, so it will be reviewed first and separately. The KG and BLP studies, while different in important ways, are still more closely related to one another. They are reviewed after FL.

2.1. A review of Feenstra and Levinsohn

A seminal study of the demand for automobiles is Bresnahan (1981). In that study, Bresnahan modeled automobiles differentiated in a single dimension. Individual models were lined up along a line according to their "quality." There was a distribution of consumers over this line and the demand for a particular model was given by integrating over the distribution of consumers where the integrands were the points at which a consumer became indifferent between two neighboring products. This was an original and sophisticated econometric implementation of a Hotelling model. Bresnahan's paper played an important motivating role in all of the three studies reviewed in this paper. FL stuck with Bresnahan's spatial approach and extended it to multiple dimensions. That is, while Bresnahan modeled products as lying along a quality line, FL model products as lying in a quality space – the

dimensions of this space are given by various product attributes. For example, instead of using "quality" to line up models on a line, FL use automobile characteristics such as weight, horsepower, air conditioning, and so on to plot models as points in a *n*-dimensional space, where *n* is the number of characteristics used. Like Bresnahan, FL assume a uniform distribution of consumers over this space. Product demand is given by integrating over the "market space" of each model. Demand elasticities are given by computing how this market space changes with changes in prices. FL are reasonably successful at the latter, but experience difficulty with the former. That is, their methods yield a rich set of demand elasticities, but are not very good at explaining the level of demand. For this reason, they confine much of their analysis to estimating the pricing side of the automobile market. Markups over marginal cost depend on the demand elasticities and not on the level of demand.

By adopting a multi-dimensional spatial approach to estimating the demand for automobiles, FL estimate which products in characteristics space neighbor one another. Unlike the Hotelling model using one dimension, products can have more than two neighbors. Indeed, when products are modelled as having 5 dimensions, FL find that cars have on average about 6 to 10 neighboring products. (The number varies with different specifications and tolerances in their algorithms.) Distance between models plays an important role in FL's analysis. Loosely speaking, products with lots of close neighbors have smaller markups, while products with few and distant neighbors have larger markups. This is merely reflecting the notion that products with several close neighbors will have relatively elastic demands while the opposite is true for models with distant neighbors.

Much of FL's analysis is devoted to how one should measure the notion of "distance." They show that for the utility function they adopt, distance is measured as Euclidean distance once characteristics are appropriately transformed, and they find that the harmonic mean of distances from a model's neighbors plays an important role in computing demand elasticities. To determine the cross-price elasticities between models, they must first map out the market space of each model in characteristics space. The size of the border between neighboring market spaces is a key input to determining cross-price and own-price elasticities. Intuitively, cars sharing a large border will have more elastic cross-price responses. Cars that do not border one another in characteristics space have zero cross-price responses.

FL's methodology is a mix of non-linear estimation and grid-search methods. They use grid search methods to map out each product's market space in a five-dimensional space, and then use results from this search in a more standard non-linear estimation of firm's pricing decisions. Using 1987 data, they find that markups range from about zero for the Ford Escort to a few hundred dollars for cars like the Toyota Tercel to around fifteen or twenty

thousand dollars for the high-end Mercedes. Compared to the markups estimated by KG and BLP, FL find markups that are somewhat smaller. The pattern of mark-ups, though, is intuitive and in line with that found in the other studies.

There are many equilibrium concepts and FL experiment with several. In particular, they estimate the pricing equation that would obtain if all firms played Cournot (simultaneously set quantities) or if they played Bertrand (simultaneously set prices). FL also estimate cases in which a firm's conjecture about its rivals' responses will depend on the nationality of the rival. This is an empirical implementation of an idea formalized by Krishna (1989), since when a firm is subject to a quantitative restraint, all firms will behave in a way consistent with this constraint. They then test which hypothesized strategies of firms are most consistent with the data. This is an approach first applied by Bresnahan (1987). FL find that specifications in which all firms believe that Japanese and European sales to the U.S. are fixed fit the data better than more traditional and straightforward models of Cournot or Bertrand conduct. This finding, if correct, has important implications for the effects of future U.S. trade policies, and these are discussed in Section 3.

2.2. A review of traditional discrete choice methods with an application to automobile demand

While the FL approach was a spatial one, KG and BLP adopt a different discrete choice framework. Because both KG and BLP can be thought of as either applications of or extensions to a traditional discrete choice framework, it is useful to briefly review that framework.

Like the approach outlined in the review of FL, the traditional discrete choice framework also begins with a utility function and derives demands. The way in which one goes from a utility function to the demand function, though, are quite different. Suppose the utility that individual i gets from product j, U_{ij} is given by:

$$U_{ij} = f(\nu_i, p_j, x_j, \zeta_j; \theta) \tag{1}$$

where ν_i is a vector of individual characteristics, (for example, income and family size), p_j is the price of product j, x_j is a vector of observed (by the econometrician) product characteristics, ζ_j are unobserved (by the econometrician) characteristics, and θ is a vector of parameters to be estimated.

Then consumer i chooses good j if and only if

$$U(\nu_i, p_j, x_j, \zeta_j; \theta) \geqslant U(\nu_i, p_r, x_r, \zeta_r; \theta), \forall r = 1, 2, \ldots J \tag{2}$$

where alternatives $r = 1, \ldots J$ represent purchases of the competing differen-

tiated products. The set A_j defines the set of values for ν that induce the choice of good j. Hence,

$$A_j = \left\{ \nu : U(\nu, p_j, x_j, \zeta_j; \theta) \geqslant U(\nu, p_r, x_r, \zeta_r; \theta) \forall r = 1, 2 \ldots J \right\}. \tag{3}$$

If $P_0(d\nu)$ is the density of ν in the population, the market share of good j, as a function of the characteristics of all the goods in the market, is given by

$$s_j(p, x, \theta) = \int_{\nu \in A_j} P_0(d\nu). \tag{4}$$

The traditional discrete choice model ignores unobserved product characteristics, hence,

$$U(\nu_i, p_j, x_j; \theta) = x_j \beta - \alpha p_j + \epsilon_{ij}. \tag{5}$$

The mean of the ϵ vector is assumed to be zero so that $x_j \beta - \alpha p_j$ is the mean of U_{ij}. It is also assumed that differences in the distribution of ϵ_{ij} across j are independent of the observed characteristics of products. This is very important, for it makes it much easier to compute the integral that gives market shares. When the ϵ are distributed multivariate extreme value, there is a closed form solution for this integral, and there is no need to simulate over consumers. This is the oft-used logit case. Since both BLP and KG use this case as a foundation on which to build, I briefly review estimation of this model, for its computational simplicity probably goes a long way towards explaining its popularity.

When the ϵ_{ij} in (5) have an extreme value (Weibull) distribution with a distribution function of $\exp(-\exp(-\epsilon))$, the integral given by (4) has the following solution.

$$s_j(p, x, \theta) = \frac{e^{x_j \beta - \alpha p_j}}{\sum_{j=0}^{J} e^{x_j \beta - \alpha p_j}}. \tag{6}$$

Allowing the outside good to be the option of not purchasing a car ($j = 0$) and normalizing the mean utility of this option to zero, we have

$$s_j = \frac{e^{x_j \beta - \alpha p_j}}{1 + \sum_{j=1}^{J} e^{x_j \beta - \alpha p_j}}. \tag{7}$$

Note that the market share of the outside good is given by:

$$s_0 = \frac{1}{1 + \sum_{j=1}^{J} e^{x_j \beta - \alpha p_j}} \tag{8}$$

hence the utility function parameters can be expressed as:

$$\ln(s_j) - \ln(s_0) = x_j\beta - \alpha p_j. \tag{9}$$

Market shares (including s_0) are data so long as one can measure the size of the potential market. For automobiles, this is typically taken to be the number of households in the United States.

If one is willing to overlook the somewhat embarrassing problem that (9) ought to fit the data exactly (since one has already integrated out over ϵ), (9) can be run as a linear regression. Eq. (9) will be used in the discussion of the approaches taken by BLP below.

It is useful, at this point, to draw some very simple comparisons between the traditional discrete choice approach and the FL approach. Both approaches essentially integrate over heterogeneous consumers to determine the demand for each differentiated product. Because FL adopt a spatial approach, it will not be the case that every model (spatially) neighbors every other model. That is, some cross-price responses are zero. In the traditional discrete choice approach, one integrates over consumers but since the ϵ are i.i.d. and unbounded, each product will have a non-zero cross-price response to each other product. That is, in the traditional discrete choice framework, there will be a draw of ϵ, the linear and random component to utility, that will lead a consumer to choose any one of the alternatives.

In the logit case, the responses to price changes are easy to compute. In particular, the own price semi-elasticity (how market share would change with a given change in the level of price) is given by $-\alpha(1 - s_j)$, where α is the coefficient on price and s_j is the market share of product j. Cross-price responses are similarly simple. These are given by αs_k, where the cross price response is the change in market share of product k with an x-dollar price change in product j. These formulae for the all-important price responses have their advantages. In particular, they are very simple, depending only on a single estimated parameter (α). Market shares, the s's, are data. Furthermore, one can estimate α using standard linear methods. Simplicity, though, comes at a cost. For 1990, with about 150 models available, the parameter α is determining over 20,000 own- and cross-price responses. This obviously implies some strong restrictions that are imposed and which cannot be tested.

It is straightforward to show, for example, that in the simple logit case, consumers who might no longer purchase their first choice, due to, for example, its price rising, will tend to substitute toward the products with the largest market shares. Hence, if the Honda Accord has the largest market share, the cross price elasticity between the Accord and any other car will be greater than the cross price elasticity between that other car and any third car. For example, the cross-price elasticity between the Accord (a mid-level family sedan) and a Porsche 911 will be greater than the cross-price elasticity between the Porsche and another high-end sports car (such as the Mazda

RX-7). While our intuition suggests that consumers who substitute away from the Porsche might be more likely to opt for the Mazda RX-7 than towards a family sedan, the traditional logit approach ensures this can *never* be the case. That is, no data will allow this reasonable cross-price response to come through in the standard logit results. This property of the logit model has led researchers to try to amend the traditional logit approach in ways that might permit more reasonable demand responses. Koujianou Goldberg takes one approach, Berry, Levinsohn, and Pakes another. Each are discussed below.

2.3. A review of Koujianou Goldberg

In "Product differentiation and oligopoly in international markets: The case of the U.S. automobile industry," KG uses a combination of consumer level data from the Consumer Expenditure Survey and a sequential modeling of the consumer's decision process to avoid some of the pitfalls of the standard logit approach. Her work has three main components. First, she models the demand side of the automobile market. Second, using results from the demand side, she estimates price-marginal cost mark-ups conditional on a Bertrand pricing assumption. Third, she uses the results from her estimated model of the auto industry to address several important policy concerns. In "U.S. Trade Policy Implications of a New Approach to Modelling the Demand for Automobiles," she extends the policy analysis to look at more issues, but maintains the same basic econometric approach. Rather, than discussing these papers separately, I will treat them as a single body of work. I proceed by first reviewing KG's methodology and then briefly survey the policy-relevant results that follow from her work.

By exploiting consumer-level data from the Consumer Expenditure Survey, KG is able work with a much richer data set than the product-level data used by FL and BLP. In particular, she allows consumer characteristics to enter the calculus of estimating market shares. The approach she adopts is original and, for lack of a better term, I will refer to it as a sequential logit approach. She divides the decision of which car a household purchases into several stages. For example, do they even buy a car in a given year? If so, do they buy a new or used car? If a new car, is it a small, medium, or large car? For each type of car, is it domestic or foreign? Finally, at the end of the sequence, which car, conditional on all the other choices already made, is actually bought?

Although at each stage of the sequential logit, a traditional logit approach is adopted, purchase probabilities will be conditional on the probabilities of each of the steps that lead to the final stage. For this reason, at the final stage, the estimated elasticities do not exhibit the counter-intuitive properties that a more traditional (single-stage) logit approach would exhibit. It is because KG has the data at the consumer-level that she has enough observa-

tions to take the sequential logit approach. The degrees of freedom are required because she must have enough observations at each node of the decision tree to estimate the probabilities at that node. Absent consumer-level data, this would not be feasible. In the end, KG estimates demand elasticities and cross-price elasticities for all models.

KG's demand elasticities are reported by model-class, not nameplate, so they are not directly comparable to FL's (or BLP's) elasticities. The pattern she reports, though, is similar to that reported by FL. In particular, domestic intermediate, domestic standard, and foreign intermediate and compact cars have more elastic demands than luxury or sports cars. Because KG uses consumer-level data with information on repeat purchasers, she is able to demonstrate that repeat buyers have less elastic demands than do first time buyers. This is an intuitively appealing result, and KG's data and methods allow her to quantify it.

Some elasticities are reported at the nameplate level. These tend to be lower (less elastic) than FL's or BLP's estimates. KG notes this fact and, in the end, it is left for the reader to decide if the elasticities are reasonable.

These demand responses feed into the pricing side of the model. There, she posits that marginal cost is a function of observable attributes and the markup is a function of the (already estimated) demand elasticities. Since her estimated elasticities are comparatively low (in absolute value), she finds that mark-ups are fairly high. Her markups average about 35 percent, while she cites industry estimates to be more in the range of 15 percent. Part of the issue here is just what represents *marginal* cost in an industry with tremendous fixed and quasi-fixed costs. Markups on a nameplate level are not reported, although a pattern of higher markups for sports and luxury cars than economy cars is reported.

Unlike FL or BLP, KG then takes her estimated model of the automobile industry and investigates some policy implications of her estimates. In "Product differentiation and international trade...", KG investigates whether the quotas on automobile imports from Japan were binding during 1983–87 and the effects the quotas had on the market equilibrium. She also investigates exchange rate pass-through. That is, to what extent did foreign producers adjust their price in dollars as the exchange rate between their home currency and dollars fluctuated. In "U.S. trade policy implications of a new approach to modelling the demand for automobiles," KG revisits the voluntary export restraint negotiated between the U.S. and Japan and investigates the predicted effects of the North American Free Trade Agreement (NAFTA).

KG's estimated pricing equations take the VER on Japanese automobiles into account. Her methodology estimates a Lagrange multiplier which gives the shadow price of the quota. By assuming that the Japanese government

allocated the quota optimally, she estimates only one such multiplier for each year. By evaluating the estimated value of this multiplier (along with its standard error), KG is able to quantify the effect of the VER. She finds that the VER was binding in 1983 and 1984, but "had weaker effects in subsequent years." She estimates that in 1983 and 1984, the VER led to a drop in Japanese sales to the U.S. of about 340,000 cars per year. The effect on American producers was to increase sales by about 100,000 per year.

KG also investigates the possible effects of the NAFTA on the U.S. automobile market. These sort of predictions are more daring, since it may not be too long before the experiment starts to yield real data. KG models the effects of the NAFTA as a decrease in the marginal costs of non-U.S. producers. The admittedly arbitrary figure used for her experiments is a three percent decrease in foreign producers marginal costs as they move production to Mexico to bypass U.S. tariffs. Her estimates imply that the cost savings from producing in Mexico are not passed on to consumers in the form of lower prices, but rather firms' variable profits increase.

2.4. A review of Berry, Levinsohn, and Pakes

Berry et al., like KG, adopt a discrete choice framework. Their methodological focus is on obtaining reasonable demand elasticities using product-level data. They introduce two methodological innovations to the standard discrete choice literature. First, drawing heavily on Berry (1994), they allow for product characteristics that give consumers utility, but which are not observed by the econometrician. Second, they adjust the utility function in order to allow reasonable substitution patterns. Each of these is briefly outlined in turn.

Recall that if we allow for unobserved product characteristics (which yield utility), Eq. (5) can be amended to give:

$$U(\nu_i, p_j, x_j, \zeta_j; \theta) = x_j\beta - \alpha p_j + \xi_j + \epsilon_{ij}, \tag{10}$$

where ξ_j gives the mean utility of the unobserved component. BLP emphasize two reasons for explicitly modelling the unobserved component of utility. First, absent that component, when one aggregates across consumers, in (5), there is no residual term remaining as seen in (9). That is, there is no source of error in the model and the model, taken literally, ought to fit the data exactly. More importantly, the component of utility that is unobserved by the econometrician but which yields utility is almost certainly positively correlated with a car's price. In other words, cars with higher, say, "quality", also cost more. Certainly, it is hard to explain the high price and substantial demand for a Mercedes or Lexus when one only looks at a few engineering characteristics such as weight, horsepower, size, and the like. It is straightforward to show that a positive correlation between the unobserved component

of utility, ξ, and price biases elasticities toward zero. This has severe implications for an equilibrium model of the automobile industry. Estimated demands will be less elastic than true demands, and this implies price – marginal cost markups that are higher than the true markups. In extreme cases, the estimated elasticities may be inelastic. This wreaks havoc on an oligopoly model as it implies negative marginal costs or infinite profits.

The usual response to this econometric endogeneity is to use instrumental variables methods. Instruments need to be correlated with a model's price but independent of the model's unobserved characteristics. In an oligopoly model in which firms' characteristics are taken as given, the characteristics of other models are appropriate instruments for price. The correction can have fairly dramatic effects. BLP, using 20 years of data spanning 1971 to 1990 find that in a simple logit model, ignoring the endogeneity of prices led to about two thirds of the elasticities being estimated as inelastic. When instrumental variables methods were used, coefficients had more intuitive signs and only about one percent of the elasticities were estimated as being inelastic.

The other methodological focus of BLP is on obtaining reasonable substitution patterns while still maintaining a structural model based on heterogeneous agents utility functions. Eq. (10) allows for unobserved characteristics, but the random component, ϵ_{ij} still enters utility in an additively separable way. This means that when the price of a car increases, consumers will tend to substitute towards cars that have large market shares, and these need not be cars with similar attributes. The particular distribution of the ϵ_{ij} does not effect this troubling property. Logit (multivariate extreme value) and probit (normal) estimators will exhibit this property. BLP address this problem by allowing individual-specific components of utility to interact with observed product characteristics. The distribution of the utility obtained from consuming product j can still be decomposed into a mean and a deviation from that mean. Allowing for unobserved product characteristics, the mean is now given by

$$\delta_j \equiv x_j \beta - \alpha p_j + \xi_j,$$

while the deviation from that mean is

$$\mu_{ij} = \sum_k \sigma_k x_{jk} \nu_{ik} + \epsilon_{ij},$$

where σ_k is the variance of marginal utilities of characteristic k. Now the properties of the distribution of the deviation from the mean utility depend on the interaction between consumer preferences for different characteristics and the characteristics of the product. As a result, consumers who have a preference for, say, horsepower, will tend to attach high utility to all high horsepower cars, and this will induce substitution effects between good that have similar characteristics.

Using a simulation estimator, BLP estimate the mean utility, β_k, for each product characteristic as well as the variance of that utility. [2] BLP combine the demand side of the market with a standard log linear marginal cost function embedded in the pricing equation. This yields a simultaneous system since demand elasticities, which depend on the level of demand, enter the oligopolists' pricing equations. Their estimated system is a market equilibrium in which (simulated) heterogeneous consumers maximize utility from either purchasing a variety of automobile (or not purchasing at all) and firms are maximizing profits over the set of products they offer. In equilibrium, prices are such that the prices firms optimally choose yield demands that are consistent with the demand elasticities used by firms in their profit maximization. An advantage of BLP is that, unlike FL and KG, their model is an equilibrium model so that one could impose a tariff, quota, or product regulation, and then resolve for the entire market equilibrium.

The results in BLP stress the importance of the two methodological innovations discussed above. They show that correcting for the endogeneity of price makes a very large difference to the elasticities estimated and the markups those imply. They also show that interacting consumer characteristics with product characteristics yields much more plausible substitution patterns than a traditional discrete choice model would imply. They make this point by estimating both a logit and their own specification using identical data. Elasticities are more elastic than those reported in KG, and this is consistent with BLP's treatment of the endogeneity of price. Correspondingly, BLP's estimated mark-ups are lower than KG's. Exact comparisons of BLP's and KG's mark-ups are misleading, though, since KG and BLP treat dealer mark-ups differently and use slightly different definitions of price. Also BLP report markups on a model-by-model basis, while KG's results are focused on averages for classes of models. FL and BLP use identical definitions and FL find markups that are somewhat lower than BLP. In FL, each model has a non-zero cross-price response with only a subset of all available products. Those with a non-zero cross-price response are termed neighbors. In BLP, every model in a given year has a non-zero cross-price response with every other. Still, the list of models with the largest cross-price responses in BLP corresponds closely with those models that FL call neighbors.

While the original BLP paper was focused on understanding equilibrium in the U.S. automobile industry and developing a methodology to do so, more recent work in progress has started to investigate policy experiments. I report here the flavor of some of these findings. I would like to stress, though, that

---

[2] Estimation is non-trivial and the interested reader is referred to the BLP paper.

the results on which I base this summary are preliminary and subject to change. BLP and KG each investigate the effects of the VERs in the 1980's. As reported above, KG finds that the VERs were binding in the initial years, and then became less so. BLP find otherwise. Their results suggest that the VERs were most binding in 1986–88. By 1990, they again were not binding. They find that the VERs accounted for an average extra-markup of Japanese cars of about $900 during these three years. BLP's results are somewhat striking as they challenge the conventional wisdom that the VER raised prices of Japanese cars in the early 1980's. (This "wisdom" is apparently more conventional in the media and among academic economists than among those working in the industry. Discussions with economists in the industry suggest that BLP's results reinforce what they already believed.) Possible explanations for the divergence between BLP's results and those of KG as well as earlier work by Feenstra (1988) are discussed in the next section of the paper. There I leave reviews of the three studies of demand in the automobile industry and concentrate on some recent and current policy issues.

3. Policy issues: evidence and unanswered questions

Two policy issues that have dominated discussions about the U.S. automobile industry are imports from Japan (and restraints thereon) and the effects of the NAFTA on the industry. These are discussed in turn. The section concludes with a discussion of some less prominent policy issues in the automobile industry.

3.1. Voluntary export restraints

There are several studies on the effect of the VERs on the U.S. automobile industry. Feenstra (1988) documents the phenomenon of "upgrading." That is, after the VER was applied, Japanese firms, constrained by the number of cars they could export to the U.S.. altered the characteristics of the cars that were exported, and the new products were more up-market.

KG found that the export restraints were binding and resulted in a rise in the price of Japanese cars, conditional on their attributes, in the first years of the VER. BLP found that the VERs did not result in significantly higher Japanese prices until 1986. One possible explanation for the difference between the two sets of results is that BLP allow for trends in costs by country of origin, and then ask whether the VER led to a deviation from the trend that would be expected to occur absent any VER. Over the 20 year sample period, Japanese prices were trending upward. Ignoring the upward trend and then only allowing prices to deviate during the VER years will tend to overstate the effect of the VER. While this might explain some of the

difference in results, an important puzzle remains. Typically, when a firm is subject to a binding quota, prices rise. Comparison of import data with allocated quota shares indicate that most of the Japanese firms appear to have been filling their export allocations. Why, then, didn't Japanese prices rise in BLP's results?

BLP suggest three factors that might have dampened the price increase that would have otherwise occurred. First, Japanese firms decided to sell more expensive (and higher "quality") products in the U.S. This is the phenomenon documented by Feenstra. Selling more expensive products, whose demand is smaller, may have allowed the Japanese to meet the quota without raising prices conditional on those product characteristics. Second, there was substantial production by Japanese firms in the U.S. This production by transplants was not subject to the quota. Hence, while the quota may have appeared to be binding, the marginal car was being produced in the U.S. where production was not constrained. This effect may have been quantitatively important. By 1986, around 300,000 cars were being produced by Japanese firms in the U.S. For a quota of around 1.68 to 1.85 million, 0.3 million *on the margin* may be very important. Third, perhaps Japanese firms could have increased prices to maximize static profits, but they were willing to sacrifice short term profits in order to maintain a reputation among price-sensitive consumers, whom they knew they would be able to profitably serve as more transplant production came on-line in the later 1980's.

These three factors are important considerations as one ponders what might happen if import restraints were to be tightened in the future. If each explanation suggests a dampened price response to import restraints in the past, it would be wrong to suggest this might re-occur. First, Japanese firms have, with the introduction of the Lexus (Toyota), Infiniti (Nissan), and Acura (Honda) line, already moved about as far up-market as is feasible. It is hard to imagine that they could make a shift of similar magnitude were they faced with another trade restraint. Second, transplant production has grown substantially. This has dual implications. First, political pressures are such that none of the large Japanese firms are currently expected to further expand transplant production beyond those plants that are already on the drawing boards. This is in spite of the fact that the rise of the Japanese yen similarly enhances the economics of transplant production by Japanese firms. While transplant production may increase in the 1990's most of this increase will probably involve European firms, not Japanese firms. For example, BMW and Mercedes have announced plant sites in the U.S., while Audi is beginning a search of its own. In the case of European firms, it is the appreciation of the German mark that has led them to expand production to North America.

Second, U.S. policy makers are aware of the transplant production in a way that they were not in 1980. I am not sure it really crossed policy-makers

minds to include transplant production in a sales quota. This is not the case now, and future protection may take the form of a sales restraint, not just an export restraint. Such policy measures have been floated in the past two years in the U.S. Congress and the issue is not going away.

In the end, we are left with diverging results concerning the effects of the VERs of the 1980's, but there is good reason to believe that firms' responses, whatever they may have been, may not be a good guide for predicting future responses to trade protection.

3.2. The North American free trade agreement

The NAFTA looms large in many discussions of trade policy and the auto industry. Consensus is a rare commodity. Koujianou Goldberg, as well as Berry et al. (1993), have investigated probably effects of the NAFTA.

KG models NAFTA as a decrease in the marginal cost of U.S. firms and/or European and/or Japanese firms as firms from these countries move their operations to Mexico. Marginal cost is assumed to fall by three percent – a figure said to equal current tariffs on imported automobiles. KG is careful to point out that the three percent figure does not come from any data on labor costs and should be interpreted with some caution. She finds that change in price, given a decrease of three percent in marginal cost, is quite small. The effect on sales of the same reduction in marginal cost is somewhat larger, ranging up to about a five percent increase for some scenarios. The actual figures differ by class of cars and depend on which subset of producers move some operations to Mexico. KG concludes by noting that her model predicts that U.S. consumers will not benefit much since prices will not fall much, although firm profits will increase due to the lower costs and higher sales. It is important to keep in mind that the demand elasticities that figure into the markup from marginal cost to price are U.S. demand elasticities, so her experiment is investigating what happens to prices, sales, and profits of firms producing *in Mexico for the U.S. market.*

This analysis is probably a very good prediction of what would happen with a depreciation of the Yen or Mark (for Japanese or German producers respectively.) The vagaries of the NAFTA agreement are such, though, that it is not necessarily as valid a prediction of what might happen with NAFTA.

Before discussing the likely impacts of trade liberalization in the automobile industry, it is useful to first review what protection already exists for U.S.–Mexican automobile trade. As KG noted, tariffs on automobiles (and parts) from Mexico into the U.S. are small. The three percent figure cited by KG is certainly in the ball park. Most of the existing protection concerns trade going the opposite direction – from the U.S. into Mexico. Here, nominal tariffs are still not very large. (There are conflicting estimates, but they seem to be in the 10 to 20 percent range.) Focussing on tariffs, though,

is very misleading. Until quite recently, imports of automobiles from the U.S. were basically banned. The exception concerned a rule that allowed firms to more or less balance trade on a firm-by-firm basis. This permitted some automobiles to be exported to Mexico, but the numbers were small. The extent of the protection is revealed by a comparison of prices. Berry, Grilli, and Lopez-de-Silanes (BGL) note that in 1990 the average price of a compact car in Mexico was US$30,574 while the average price of a luxury car was US$45,383. What makes these numbers so astounding is the definition of a compact and luxury car. Compacts include cars like the VW Jetta or Ford Topaz, while luxury cars include cars like the Buick Century or Ford Taurus, but no U.S. style luxury cars like the Cadillac or Mercedes. These figures indicate that while nominal tariffs are low, there are *very* large non-tariff barriers protecting the Mexican car industry. Estimates of 200 to 250 percent do not seem unreasonable.

The reasons for the high implicit protection lie in the history of the Mexican automobile industry. Since 1962, there have been a number of Automotive Decrees. These are reviewed in BGL. These decrees were motivated by Mexican policy makers concerned with a large sectoral trade deficit in automobiles. The exact policies that currently account for the very large differential between U.S. and Mexican auto prices (for essentially the same car) are myriad and complex. One effect of these decrees is that only a few foreign firms were allowed (in 1964) to produce automobiles in Mexico. These firms now go by the names of General Motors, Ford, Chrysler, Nissan and Volkswagen. This historical fact has very important implications for modelling the probable effects of NAFTA.

The U.S. automobile producers lobbied hard and successfully for a provision in the NAFTA (Chapter 4, Article 403 in the December 1992 draft) that effectively prohibits firms which are not already producing cars in Mexico from now commencing production for export back into the U.S. Hence, predictions that entail Japanese or European firms (except for VW) producing in Mexico for export to the U.S. are inconsistent with the way NAFTA is structured. VW already produces the new 1994 Jetta and Golf in Mexico for export to the U.S., hence it seems unlikely the NAFTA will have a big effect on the one European producer which is exempted. Nissan is the sole Japanese firm that may be able to export in the U.S., but they have already invested heavily in the U.S. That fact, in combination with their flagging U.S. market share, suggests to me that Nissan is not about to make huge capacity investments in Mexico in order to bypass a three percent tariff and export freely into the U.S. The numbers just do not add up. NAFTA is more likely to have a larger impact on the other Japanese firms since they *cannot* produce in Mexico even if they wanted to. And with the recent rise of the yen and the risk of taxes on U.S. production by the Japanese firms, they might want to.

These factors led BGL to predict that the big effects of the NAFTA will be related to the increase in demand *by Mexicans* as the price of automobiles falls and their income rises with trade liberalization. Hence, BGL argue, a narrow focus on the U.S. market misses most of the real effect of the NAFTA on the industry. As demand in Mexico increases, quite dramatically perhaps, U.S. firms are well situated to capitalize on this and increase their sales south of the border. This line of argument seems quite on target.

There has been a great deal of media attention on automobile industry jobs moving to Mexico. The UAW (the auto workers union) has steadfastly opposed the NAFTA on the grounds that it will shift employment to lower-wage Mexico. Ross Perot based much of his unsuccessful 1992 presidential campaign on similar fears. He continues to lobby vigorously against the NAFTA on the same grounds. BGL note that there have been few constraints, prior to NAFTA, on U.S. production in Mexico for export to the U.S. U.S. firms could have moved to Mexico for export to the U.S. if they wanted to and apparently they did not want to. (The broad exceptions to this are the production of upholstery and wire harnesses – two very labor intensive production operations and Ford's plant at Hermisillo. These tasks have already moved to Mexico.) It is not obvious, then, why they would suddenly do so now.

Servicing the new demand from Mexico may involve expanding operations in Mexico, but this is not the same as moving jobs currently in the U.S. to Mexico. Rather, that decision involves where to add new jobs. BGL note that "While some new plants may locate in Mexico to produce models with have particular appeal in Mexico, we should note that apart from the VW Beetle the models currently produced and sold in Mexico are very similar to other North American models. Therefore, much if not most of the increased Mexican demand could be served from plants with serve the overall North American market. As we argue above, if these plants were likely to move en masse to Mexico, they would have done so already" (p. 40).

BGL estimate a constant elasticity demand system at the model class (instead of nameplate) level. They estimate that a price decrease to U.S. levels about doubles demand for cars in Mexico. Furthermore, if Mexican national income grows by five percent a year for the next five years, the income effect this entails would also greatly increase demand. The combination of U.S. level prices and strong growth, estimate BGL, could lead to a seven-fold increase in demand over the next five years. This is the large effect of the NAFTA on the automobile industry and it involves higher U.S. firm profits and happier Mexican consumers. BGL note that their demand analysis is basic and that the methods of BLP or KG are well suited to a more detailed study. This work remains to be done.

All of the studies ignore the demand for used cars and U.S.–Mexican trade therein. The NAFTA, though, does address the issue of international

trade in used cars. Specifically, liberalization here is still far in the future. Imports of used cars into Mexico remain prohibited for 15 years after signing and are, even then, only gradually liberalized. Initially, only ten year old cars will be allowed in, then nine year old cars the following year and so forth. For the reasonably near future, then, international trade in used cars remains quite minimal. One can surmise that this prohibition is in the interest of firms that would otherwise sell new cars in Mexico, while consumers are made worse off. The calculus, though, is not an obvious one for automobile firms. While allowing used cars to be exported from the U.S. into Mexico will decrease the demand for new cars in Mexico, it may have the effect of simultaneously increasing new car sales in the U.S. and Canada. If there is increased demand for used cars, their price will rise and this will lower the expected user-cost of a new car. Realizing that they can get a higher price for their used car, American and Canadian consumers would be more likely to purchase a new car, hence increasing sales. Automobile manufacturers must balance the lost sales of new cars in Mexico with the increased sales in the rest of North America when used car trade is allowed. Given the prohibition on trade in used cars for the near future, though, this is an academic exercise, albeit not one that anyone has actually undertaken.

3.3. Other policy issues effecting the U.S. Automobile industry

While trade restraints and the NAFTA are the two policy issues with the highest visibility, there are other important policy issues effecting the automobile industry. These include the role of anti-dumping law, a gasoline tax, environmental regulations, and the globalization of the industry.

3.4. Anti-dumping law

Anti-dumping law has probably had a significant impact on U.S.– Japanese automobile trade. This is despite the absence of anti- dumping petitions filed by U.S. automobile firms. Recent theoretical work in the International Trade literature has suggested that anti-dumping law can work to facilitate collusion in an industry. The story is basically as follows: Oligopolistic firms would like to raise prices but, in equilibrium, do not do so for fear of being undercut by a competitor and losing (profitable) sales. An anti-dumping petition acts as a signal (and credible threat) which leads foreign firms to increase their prices. U.S. firms then do likewise and the new equilibrium entails higher prices and profits. This story works even in the absence of a finding of dumping, and indeed evidence suggests that many anti-dumping petitions are withdrawn prior to any ruling. Whether this phenomenon has effected the automobile industry is pure speculation. And I speculate that it has.

In late 1992 and early 1993, the U.S. auto firms (or at least two of the big three) were reported to be considering filing an anti- dumping suit against Japanese firms. Had the reported petition actually been filed, it would have represented the largest anti- dumping claim ever filed. Instead, in early 1993, Japanese prices rose and the discussion of filing faded away. There are several reasons, the strength of the yen being the foremost, why Japanese prices might have risen anyway. Nonetheless, the timing suggests that the threat of an anti-dumping suit played a role in Japanese pricing decisions.

This issue is not going away. As I write this, today's (October 25.1993) issue of Automotive News, a U.S. trade weekly, quotes the new Chairman-elect of Ford, Alex Trotman, as saying: "Dumping is going on. The yen has increased 22 percent (since the beginning of the year), but prices have moved by only about 4 percent. There is a big discrepancy between the yen and pricing." Trotman said Ford, GM, and Chrysler have agreed to hold off on filing a dumping complaint "to allow the discussions with the Japanese government a little more time." KG, as well as Feenstra et al. (1993) have demonstrated that the one-to-one exchange rate pass-through to prices implied by Trotman is not a new phenomenon. Imperfect pass-through has been the rule.

3.5. A gasoline tax

In President Clinton's recently passed budget, a significant rise in the tax on gasoline was a hotly debated issue. In the end, the gas tax rose, although only modestly. A significant rise in the tax on gasoline has implications for the automobile industry. In particular, it makes driving more expensive and induces substitution towards more fuel efficient cars from the consumer's perspective. From the producers' perspective, such a tax increases the incentive to conduct research and development focussing on fuel efficiency. On first glance, it seems that the static effects of a large gas tax would favor Japanese producers, since their product line is more heavily oriented toward more fuel-efficient vehicles. The experiences of the mid-1970s also support this prediction. U.S. producers did favor the (unsuccessful) large increase in the gas tax, although this is because of an implicit *quid-pro-quo* that the gas tax would replace the complex Corporate Average Fuel Efficiency (CAFE) standards.

3.6. Globalization of the industry

The automobile industry has undergone significant globalization in the last decade. A perhaps unintended result of this is that the efficacy of trade policy has diminished. When a Chrysler product was produced in the U.S. using American parts and a Honda was produced in Japan using Japanese

parts, trade policy was relatively simple. Today, the U.S. firms sell cars produced abroad under their own label. In 1992, Chrysler sold the Colt, Stealth, and Vista. Chevrolet sold the Geo, and Ford sold the Festiva,Tracer and Capri. All of these were imported. On the other hand, Mazda produces in Michigan, Toyota in California and Kentucky, Nissan in Tennesee, Honda in Ohio, and Suburu in Indiana. If this is not sufficiently confusing, the parts trade is also very globalized. For example, Honda's domestic content (the value of parts purchased in the United States as a percentage of total value of the car) is said to exceed Chrysler's domestic content. Chrysler would contend that this is deceiving since many of its "foreign" parts are made by U.S. owned firms operating in Canada. Many of Honda's "domestic" parts are purchased by Japanese firms producing in the U.S. The actual numbers are not relevant here. Rather, the point is that it is becoming very difficult to determine whether a car is imported or domestic. Since trade policy is traditionally oriented toward exactly this distinction, the usual policy tools of tariffs and quotas based on which goods cross the border are no longer especially well suited to the industry. This suggests that unless the demands for protection disappear, new trade policy tools will have to be developed.

In terms of econometric modelling, this internationalization greatly complicates matters. The studies by BLP, KG, and FL all assumed constant marginal cost with respect to output. (BLP report some experiments with decreasing marginal costs, but this is more of a sensitivity analysis than a study of how marginal costs *really* vary with output in the industry.) With constant marginal costs, it is difficult to model the decision of where to produce the marginal car – at home or abroad. Capacity constraints certainly play a role, but these depend on fixed costs, and these in turn are difficult to measure. The bottom line is that globalization, while very real, is thus far mostly ignored in the academic studies.

3.7. Environmental policy

The U.S. auto industry is subject to much regulation that is itself motivated by environmental concerns. The most pervasive current regulations are the CAFE standards. First implemented in 1975, these standards force firms to achieve a sales-weighted fleet- average fuel standard. The standard has generally increased over time, although on one occasion the standard was lowered. The details, especially those involving credits and carry-overs are remarkably complex. The net effect, though, is probably to make fuel efficient cars somewhat less expensive and fuel inefficient cars somewhat more expensive than they would be in the absence of the regulation. One might then conclude that these standards and other environmental regulations on new cars have led to a more fuel efficient and cleaner fleet of vehicles on the road. That this is perhaps wrong points to an important flaw in all the studies I have reviewed.

All the academic studies have focussed on the new car market. This orientation misses some very important and difficult to model dynamic issues. The importance of this omission is best illustrated in connection with the debate concerning raising the CAFE standards. Making a car more fuel efficient costs money and implies higher prices for new cars. Suppose the CAFE standard was raised again. As new cars became more fuel efficient and more expensive, drivers would tend to keep their used cars longer. Older cars have higher emissions and lower fuel efficiencies. The net effect of the regulation could backfire yielding more pollution and worse gas mileage for the cars actually on the road. Modelling the trade-off between new and used cars is very important to understanding how these policies would work. This modelling effort, though, is just starting and there are, as yet, no results to report.

4. Summary

This paper has reviewed three econometric studies of the U.S. automobile market. Each employed different and original methodologies, and these were compared and contrasted. Some of the studies had investigated the policy implications of their approaches and these were discussed. The big policy issues effecting the industry are import/export restraints and the NAFTA. While some research has investigated how these policies have played out in the industry, there are many unresolved questions. I used this paper as an opportunity to speculate on some of these unresolved issues. Time will tell whether the existing research, which is very new, will prove useful and whether my own speculation on policy issues is on target. I am much more confident about the former than the latter.

References

Berry, S., 1994, Estimating discrete choice models of product differentiation, RAND Journal of Economics 25, 242–262.

Berry, S., V. Gilli and F. Lopez-de-Silanes, 1993, The automobile industry and the Mexico–U.S. Free Trade Agreement, in: P. Garber, ed., The Mexico–U.S. free trade agreement (MIT Press, Cambridge, MA).

Berry, S., J. Levinsohn and A. Pakes, 1994, Automobile prices in market equilibrium, Econometrica, forthcoming.

Bresnahan, T., 1981, Departures from marginal-cost pricing in the American automobile industry, Journal of Econometrics 17, 201–227.

Bresnahan, T., 1987, Competition and collusion in the American automobile oligopoly: The 1955 price war, Journal of Industrial Economics 35, 457–482.

Feenstra, R., 1988, Quality change under trade restraints: Theory and evidence from Japanese autos, Quarterly Journal of Economics, 131–146.

Feenstra, R. and J. Levinsohn, 1994, Estimating markups and market conduct with multidimensional product attributes, Review of Economic Studies, forthcoming.

Feenstra, R., J. Gagnon and M. Knetter, 1993, Market share and exchange rate pass-through in the world automobile industry, NBER Working Paper 4399.

Koujianou Goldberg, P., 1993, U.S. trade policy implications of a new approach to modelling the demand for automobiles, mimeo, Princeton University.

Koujianou Goldberg, P., 1992, Product differentiation and oligopoly in international markets: The case of the U.S. automobile industry, mimeo (Princeton University).

Krishna, K., 1989, Trade restrictions as facilitating practices, Journal of International Economics 26, 251–270.