RESEARCH SEMINAR IN INTERNATIONAL ECONOMICS

School of Public Policy
University of Michigan
Ann Arbor, Michigan 48109-1220

Discussion Paper No. 393

Voluntary Export Restraints on Automobiles: Evaluating a Strategic Trade Policy

Steven Berry
Yale University
National Bureau of Economic Research

James Levinsohn
University of Michigan
National Bureau of Economic Research

Ariel Pakes
Yale University
National Bureau of Economic Research

March 19, 1997

Voluntary Export Restraints on Automobiles: Evaluating a Strategic Trade Policy

by

Steven Berry
Yale University
National Bureau of Economic Research
James Levinsohn
University of Michigan
National Bureau of Economic Research
and
Ariel Pakes
Yale University
National Bureau of Economic Research

Current version: March 19, 1997

Address. Berry and Pakes: Department of Economics, 37 Hillhouse Ave., Yale University, New Haven, CT 06520; Levinsohn: Department of Economics, University of Michigan, Ann Arbor, MI 48109; Internet: SteveB@econ.yale.edu, Ariel@econ.yale.edu, and JamesL@umich.edu

Voluntary Export Restraints on Automobiles: Evaluating a Strategic Trade Policy

Steven Berry
Yale University
National Bureau of Economic Research
James Levinsohn
University of Michigan
National Bureau of Economic Research
and
Ariel Pakes
Yale University
National Bureau of Economic Research

1. Introduction.

In May, 1981, a voluntary export restraint (VER) was placed on exports of automobiles from Japan to the United States. As trade policies go, this one was important. The automobile industry is the largest manufacturing industry in the United States and the initiation of the VER captured headlines in the popular press. At about the same time, though to much less fanfare, international trade theorists were obtaining (then) startling results from models of international trade in imperfectly competitive markets. These models suggested that in imperfectly competitive markets, an activist trade policy might enhance national welfare. In this paper, we provide some empirical evidence on whether the these new theoretical possibilities might actually apply to the policy of VERs.

In so doing, we address the following "big-picture" questions. First, did the VERs matter? That is, did they raise prices and, if so, by how much? Second, how much did the VERs benefit the domestic producers and how much did they hurt Japanese producers? Also, how were European firms affected by the policy? Third, were the VERs sound domestic public policy and, if not, could they have been if they had been implemented differently? Our answers are at odds with much of the existing empirical literature on the VERs. In particular, the point estimates of our model imply that: 1) The VERs did *not* significantly raise prices when they were first initiated, but they were responsible for higher prices of Japanese cars in the later 1980's and that accounting for direct foreign investment by the Japanese auto producers into the U.S does not really change this conclusion; 2) Summing over the years for which the VER's were binding, the VERs increased the

We are grateful to Jagdish Bhagwati, Alan Deardorff, Robert Feenstra, Gene Grossman, Mustafa Mohatarem, Dani Rodrik, Gary Saxonhouse, Steve Stern, Marina V.N. Whitman, Frank Wolak, and two anonymous referees for helpful comments. We gratefully acknowledge funding from National Science Foundation Grant SES-9122672.

profits of U.S. producers by about ten billion (1983) dollars, and this estimate has a standard error of about seven billion dollars. We also find that U.S. producers responded to the VERs by selling more cars, but they did not significantly raise prices as it was typically the price sensitive consumer who switched from Japanese to domestic cars; and 3) The VERs resulted in moderate net welfare losses to the U.S. (our point estimate of the loss is close to \$3 billion, but it has a standard error of \$7.5 billion.)

We also compute what would have happened to U.S. welfare had the VERs instead been implemented as tariffs or quotas. However, this calculation requires us to assume that the tariffs would not cause any change in the cars marketed in the U.S., or lead to trade retaliation of any form. Under these questionable assumptions, replacing the VERs with a tariff would have enhanced U.S. welfare by about 8.3 billion (1983) dollars with a standard error of 8.3 billion dollars, leaving open the possibility that strategic trade policy could have actually worked. This change in welfare is comprised of three components– the above mentioned increase in domestic profits, the foregone tariff revenue, and the change in consumer welfare. We estimate that the revenue foregone by using a VER instead of a tariff was 11.2 billion dollars (with a standard error of 3.1 billion dollars.) This foregone revenue almost equals the loss in consumer welfare of 13.1 billion dollars (with a standard error of 2.5 billion dollars.)

This paper has, of necessity, a large methodological component. This is due to some large discrepancies between the standard theoretical models and the actual structure of the automobile market. While theory is typically constructed around models with two countries, symmetric firms each producing one product, a constant elasticity of demand between differentiated products (or homogeneous products), a representative consumer with a love of variety, and observed marginal cost, empirical work must confront a very different situation. In the case of the U.S. automobile market, there are multiple firms of vastly different sizes, almost all of which produce multiple products. These firms are from about a half dozen different countries. There are, in any given year, roughly 20,000 unknown elasticities and they are not equal. These elasticities play a key role in determining the Nash equilibrium prices firms charge. There are over 90 million households potentially in the market and they are quite heterogeneous. Finally, marginal cost is unobserved. Dealing carefully with these facts and constraints, while still obtaining explicit guidance from an equilibrium oligopoly model, requires new methodological tools, which we take largely from Berry, Levinsohn and Pakes (1995, henceforth BLP.)

As in any policy analysis of the VERs, in order to arrive at our conclusions we have to make a host of very detailed assumptions about functional form and behavior. We are explicit on exactly what these assumptions are, hence allowing other researchers to evaluate and expand our analysis. We also provide extensive sensitivity analyses investigating how changes in these assumptions impact results.

This paper is organized into 7 sections. In section 2, we review some of the existing empirical literature examining the VERs on automobiles. In section 3, we outline the underlying theoretical model used here to evaluate the VERs, while section 4 discusses the methodology used to estimate this model. Section 5 presents a discussion of policy details, the data, and the base case results while section 6 is focussed on determining how robust our results are to several alternative theoretical and econometric specifications. Conclusions and caveats are gathered in section 7.

2. The Previous Literature.

At the most general level, we hope this paper might contribute to the debate on the applicability of the insights of the strategic trade policy literature. On the one hand, some of the economists most responsible for the development of the theory of strategic trade policy have argued eloquently against its use in the public policy arena. See, for example, Paul Krugman's (1994) *Peddling Prosperity*. On the other hand, the insights from the the strategic trade policy literature appear to have struck a chord with some currently powerful policymakers and advisors.

Since the early theoretical models are now over a decade old, one might have expected that there would be several econometric studies investigating exactly this question in a multitude of industries. We know of no econometric studies of strategic trade policy. This absence is documented in the recent review of empirical studies of trade policy by Robert Feenstra (1995). As noted in Feenstra's survey, the empirical studies of strategic trade policy have been simulation models in which simple theoretical models are parameterized and experiments run.

While we know of no econometric studies investigating the efficacy of an implemented (possibly) strategic trade policy, there have been several studies of international trade and the U.S. automobile industry. While a complete survey of this literature is beyond the scope of this paper, we provide an overview of some of this work. (See Levinsohn (1994) for an extended survey.)

Some of the first studies of the effects of VERs on the U.S. automobile industry were by Robert Feenstra (1984) and (1988). These studies focused on the phenomenon now referred to as quality upgrading. Feenstra documented that when the VERs were implemented, the list prices of Japanese cars as well as the base-model characteristics of those cars increased. Using data from 1979 to 1985, [1] he showed that some of the observed price increases in Japanese cars could be accounted

[1] Not all the Feenstra papers used all these years of data, but Feenstra (1988) uses all years.

for by corresponding increases in "quality," such as more horsepower, larger vehicle size, and the like. Hence, if one only looked at the change in prices, without adjusting for the concurrent change in quality, one would over-estimate the price rise due to the VERs.

Avinash Dixit (1988) constructed a simple simulation model of the U.S. automobile industry in which there were two types of products, U.S. and Japanese. Assuming linear inverse demands and constant marginal cost, Dixit calibrated his model to perfectly fit data that were aggregated in this way. This was done for the industry in 1979 and again for 1980. Drawing on elasticities and estimates of marginal cost from various sources, Dixit computed the *optimal* strategic trade policy and compared the welfare gain this would have yielded relative to the simpler policy of levying a standard Most Favored Nation tariff of 2.9 percent. Dixit found that the gains from employing strategic trade policy would have been very small– on the order of 17 to 300 million dollars depending on the policy tools adopted and the parameters selected.

Elias Dinopoulos and Mordechai Kreinin (1988) treat the U.S. automobile industry as a homogeneous product perfectly competitive industry with linear supply and demand schedules and compute the triangles that comprise the deadweight loss from the quality-adjusted price increase the VER induced.

A more recent and more sophisticated empirical investigation of the effect of the automobile VERs on the United States is Pinelopi Goldberg (1995).[2] In that paper, Goldberg estimates a structural oligopoly model of the U.S. automobile industry using both product-level data and consumer level data from the Consumer Expenditure Survey. Her annual data cover 1983 to 1987. Goldberg first estimates a logit-based demand system from the consumer data in the CES. This yields demand elasticities that feed into the oligopolistic firms' profit maximizing first order conditions. These first order conditions result from multi-product firms maximizing profits in a Bertrand fashion. Goldberg finds that the VERs were binding in 1983, 1984, and again but much less so in 1987. A principal message of Goldberg's paper is that the main effect of the VERs came immediately after they were imposed and that in later years the policy had little or no effect. Goldberg reports on the profit shifting aspect of the trade policy, but notes that "the objective of our analysis is not to compute national welfare, but to assess the quota impact on prices, production and market shares..." We return to her conclusions after presenting our results.

We address the broader question of whether the VERs were sound U.S. public policy. In particular, when the entire picture of U.S. firm profits, consumer welfare, and government revenues

[2] A less technical paper that also addresses many of these issues is Goldberg (1994).

are considered, who were the winners, who were the losers, and what was the magnitude of these gains and losses? To address these questions, we use a structural model of static oligopoly. This model is presented in the next section.

3. A Model of VERs in Oligopoly

To proceed we need a model of demand and supply for the new car market. The model we use has four primitives; i) a distribution for consumer utility functions, ii) a distribution for producer cost functions, iii) a specification for the rules governing the impacts of the VER's, and iv) a behavioral assumption which determines equilibrium. We take our specification for the distribution of the utility and cost surfaces from our earlier work (BLP, 1995) which we review briefly now. We next provide our specification for the VER's and then consider alternative equilibrium concepts.

Utility and Demand

Our demand system is obtained by explicitly aggregating over the discrete choices of individuals with different characteristics.[3] The utility that a consumer derives from a given choice depends upon the interaction between the consumer's characteristics, to be denoted by ν, and the product's characteristics. Thus the preference for a car of a particular size may depend on family size, while price tradeoffs may depend on family income. We distinguish between three kinds of product characteristics; those that are observed by the econometrician but determined before the current period (such as horsepower and vehicle size) to be denoted by x, price, or p, which is also observed but may be changed in every period, and unobserved (by us) product characteristics, denoted by ξ. The vector ξ is meant to take account of characteristics that are observed by market participants, but are either inherently difficult to measure (such as "prestige") or are potentially measurable but are not included in our specifications (usually because of a lack of data).

The consumer has J+1 choices. She can choose to purchase one of the J cars marketed, or she can choose not to purchase a new car. We let the (indirect) utility derived by consumer i from choosing alternative j be

$$U(\nu_i, p_j, x_j, \xi_j; \theta),$$

where θ is a vector of parameters to be estimated. Consumer i chooses alternative j if and only if

$$U(\nu_i, p_j, x_j, \xi_j; \theta) \geq U(\nu_i, p_r, x_r, \xi_r; \theta), \text{ for } \quad r = 0, 1, ..., J,$$

[3] For a discussion of the advantages of demand systems obtained in this way, and a review of the relevant literature, see BLP, 1995, and the literature cited therein.

where alternatives $r = 1, ..., J$ represent purchases of the competing differentiated products. Alternative zero, or the outside alternative, represents the option of not purchasing any of those products and allocating all expenditures to other commodities. It is the presence of this alternative that allows us to model changes in the total quantity of automobile purchases.

Let $A_j(\theta)$ be the set of values of ν that induce the choice of good j when the parameter vector is θ:

$$A_j(\theta) = \{\nu : U(\nu, p_j, x_j, \xi_j; \theta) \geq U(\nu, p_r, x_r, \xi_r; \theta), \text{ for } r = 0, 1, ..., J\}. \tag{1}$$

The market share, s_j, of a product is given by computing the fraction of the population with $\nu \in A_j$. That is,

$$s_j(p, x, \xi; \theta) = \int_{\nu \epsilon A_j(\theta)} P_0(d\nu), \tag{2}$$

where P_0 provides the distribution of ν.

A note on functional forms is appropriate here. Computational constraints have frequently induced the traditional discrete choice literature to analyze models in which utility is additively separable into a component that depends only on product-level attributes, say δ_j, and a disturbance, say ϵ_{ij}; i.e. $U(\nu_i, p_j, x_j, \xi_j; \theta) = \delta_j + \epsilon_{i,j}$. The $\epsilon_{i,j}$ are assumed to be independently and identically distributed across choices, as the specification then enables one to compute market shares from the solution to a unidimensional integral (if, in addition, the ϵ are distributed multivariate extreme value, the needed integral has an analytic form). However, the computational simplicity that these assumptions produce comes at a large cost. These assumptions result in a model which, no matter the parameter estimates (or the precise values of the δ_j), implies that when consumers substitute away from one product they will not substitute towards products with similar characteristics, but rather to products with large market shares; a fact which leads to counterintuitive cross-price elasticities (see BLP,1995).[4]

To enable richer substitution patterns we allow different consumers to have different intensities of preferences for different characteristics. We do this in a tractable way via a random coefficients

[4] Related properties of the standard assumptions have been noted by several authors and have led to several alternative modeling assumptions. Probably the most well known of the modifications is the nested logit. In the nested logit the researcher provides an *a priori* classification of products into groups and then has substitution patterns constrained only between members of the same group and between a member of one group and members of any other group (see Cardell, 1991, for an intuitive discussion). An alternative, and one which is closer to our specification, is the random coefficients model used by Hausman and Wise, 1978. This specification does not produce an analytic integral for the shares. However, if the dimension of the random coefficients is small enough (as it was in the Hausman and Wise case), numerical integration can be used to solve for those shares.

utility specification. The utility function for consumer i, considering products indexed by j, is

$$\begin{aligned} u_{ij} &= x_j\bar{\beta} + \xi_j - \alpha_i p_j + \Sigma_k \sigma_k x_{jk} \nu_{ik} + \epsilon_{ij} \quad \text{for } j = 1, ..., J, \quad \text{while} \\ u_{i0} &= \sigma_0 \nu_{i0} + \epsilon_{i0}. \end{aligned} \tag{3}$$

The ϵ_{ij} are traditional *i.i.d.* extreme value ("logit") draws, which capture an idiosyncratic taste of this consumer for this product. The term $x_j\bar{\beta} + \xi_j$, where $\bar{\beta}$ is a parameter to be estimated, is common to all consumers. This term allows the mean level of utility to vary with observed and unobserved characteristics. Consumers then have a distribution of tastes for each of the product characteristics. For each characteristic k, consumer i has a taste ν_{ik}, which is drawn from an *i.i.d.* standard normal. The parameters σ_k capture the variance in consumer tastes. Similarly, the parameter σ_0 captures additional variance in consumers' tastes for the outside good. Because the outside good is in fact a broad category including, *e.g.*, all used cars and public transport, we expect the idiosyncratic variance for this alternative to be larger than the variance for the "inside" goods.

The term α_i is the consumer's distaste for price increases. As in BLP, we assume that the distribution of α_i varies with income. Accordingly, we assume that α_i has a time-varying distribution that is a log-normal approximation to the distribution of income in U.S. households in each year. If y_i is a draw from this log-normal income distribution, then

$$\alpha_i = \frac{\alpha}{y_i},$$

where α is a parameter to be estimated. In this way, price sensitivity is modeled as inversely proportional to income.[5]

Because the utility specification in (3) allows consumers to differ in their preferences for product attributes, consumers who substitute out of, say, a large car, will tend to be consumers who like large cars, and, precisely because of this preference, will substitute disproportionately to other large cars. As a result, the specification in (3) allows for a much richer set of substitution patterns than does vthe traditional logit model.

The random coefficient generalization of the logit model does, however, carry the cost of an increased computational burden. Now, to obtain the market shares implied by the model we will need to evaluate a $k+1$-dimensional integral. As shown in Pakes(1986), this aggregation problem can be solved by simulation.

[5] This functional form for the interaction between income and price can be derived as a first-order Taylor series approximation to the "Cobb-Douglas" utility function used in BLP.

The other novel feature of our model is the allowance for unmeasured product attributes, the ξ_j. Just as with the disturbances in the homogeneous goods supply and demand model, these unobserved characteristics are not integrated out in computing aggregate demand. Hence, they are a real source of difference between the aggregate predictions of the model and the actual data. As one might suspect, however, the ξ_j also generate a differentiated product analogue to the econometric endogeneity problem we are familiar with from the homogeneous goods model. That is, unmeasured characteristics, such as perceived reliability or prestige, are likely to be determinants of and hence correlated with the product's price. If the econometric endogeneity of price is unaccounted for in the estimation algorithm, it will generate inconsistent estimates of the demand elasticities. Berry (1994) suggests using an inversion routine to solve for the ξ, and then instrumental variable techniques to estimate the parameters, and BLP provides a simple way of implementing these suggestions (see below). BLP also shows that the bias generated by the econometric endogeneity of price is likely to be empirically important.[6]

This completes the discussion of the utility side of our model. We now turn our attention to the firm's problem.

Firms, Costs, and Equilibrium Prices

The firm side of the model is straightforward. In any given year, there are F firms, each of which produces some subset of the J products, $\mathcal{J}_f$. The decision of which products (bundles of characteristics) are produced in any year is assumed to be predetermined outside of our model. [7]

Marginal costs are assumed to depend on observed product attributes, country-specific cost shifters such as wages and exchange rates, and an unobserved productivity variable. The product attributes that enter marginal cost may be the same as those that determine utility (though this is not necessary), and the unobserved productivity term may be correlated with the unobserved product attributes (or the ξ_j). Note that we assume that marginal costs are independent of output levels. The decision to model a product's marginal cost as constant is the result of our data limitations. We do not observe worldwide output of foreign models and this, not just sales in the U.S., is what marginal cost might vary against (see the discussion in BLP). In addition, almost all

[6] As an example, when we do not account for the endogeneity of price, several products are estimated to face inelastic demands; this is problematic in an oligopoly model.

[7] Modeling the firm's decision of which products to produce conditional on its beliefs about the products other firms will produce and the state of future demand in a multi-dimensional differentiated products oligopoly is an important and very difficult problem that is beyond the scope of this paper.

researchers since Bresnahan (1981) have adopted the constant marginal cost assumption.[8] Using a logarithmic specification then, the marginal cost of product j is written as:

$$ln(mc)_j = w_j\gamma + \omega_j, \tag{4}$$

where γ is a vector of parameters to be estimated, w_j is a vector of observed marginal cost shifters, and ω is the unobserved productivity term.

To move from demand and costs to industry equilibrium requires two modeling decisions. First, how should the VER be modeled? Second, what is the equilibrium concept – Cournot, Bertrand, or something yet different?

When Japan "voluntarily" agreed to reduce automobile exports in May, 1981, the agreement pertained to total exports from Japan. These were to be limited to 1.68 million units (a figure that increased in later years.) The Ministry of Trade and Industry (MITI) in Japan then essentially divided this limit across the Japanese automakers. It has been suggested that a firm's allocation depended in various ways on past sales or market shares, and this is surely true, but there is not a (publicly available) hard and fast formula used by MITI.

Modeling the VER raises several issues. There is a large literature discussing tariff-quota equivalences or non-equivalences in the presence of imperfect competition, and the lessons from that literature might, at first glance, appear relevant here. For example, Bhagwati (1969) showed that in a linear monopoly model, tariffs and quotas might be non-equivalent. In an oligopoly setting, Krishna (1989) has demonstrated that when firms compete by setting quantities (as in Cournot), the quota and an appropriately set specific tariff are equivalent, in that they yield the same equilibrium. This is not the case when firms set prices. Krishna notes that with a VER or quota on the foreign firm, the home firm's best response function is discontinuous, and there need not be an equilibrium in pure strategies.

However, in light of how the VER was actually implemented, we believe that the target levels of exports MITI allocated to the firms should not be viewed as firm specific quotas. Failure to meet the target presumably impacted negatively on the firm's relationship with MITI and probably on the firm's future allocations. It did *not* prevent an additional unit from being exported. (It is often claimed that Suburu and Honda exceeded their allocations in early years of the VER.) Rather, the firm would have to evaluate these costs and decide on a course of action. As a result we choose to

[8] The importance of the constant marginal cost assumption in the analysis of trade policy in the auto industry is explored, using a partially calibrated model, in Fuss, Murphy, and Waverman (1992).

model the impact of the firm specific limits as a tax on exports in excess of that limit. The tax rate is the implicit unit cost of exceeding MITI's limits, and becomes a parameter to be estimated.

For simplicity, we begin with the case in which the VER is implemented as an implicit tax on every unit exported. If the tax per unit is denoted by λ, the firm's profits are given by

$$\pi_f = \sum_{j \in \mathcal{J}_f} (p_j - mc_j - \lambda VER_j) M \ s_j(p, x, \xi; \theta) - \sum_{j \in \mathcal{J}_f} Fixed\ Costs_j, \tag{5}$$

where M denotes the market size and VER is a dummy variable that is set to one if the car is subject to the tax.

Initially assume that the equilibrium is Nash in prices, i.e. at equilibrium each firm is setting each of its product prices to maximize total firm profits conditional on the prices charged by the other firms and the characteristics of all the cars marketed. Provided such an equilibrium exists, the resulting prices must satisfy the first order conditions:

$$s_j(p, x, \xi; \theta) + \Sigma_{r \in \mathcal{J}_f} (p_r - mc_r - \lambda VER) \frac{\partial s_r(p, x, \xi; \theta)}{\partial p_j} = 0. \tag{6}$$

In the simple case where there is one product per firm, equation (6) sets a price equal to marginal cost plus the tax (where applicable) plus a markup equal to the inverse of the elasticity of demand for that product. For our multi-product firms the markup is more complicated as the firm takes account of the effect of a change in the price of one of its product on the profits earned from all of its products. In particular if we let the vector of markups for the multi-product firm case be $b(p, x, \xi; \theta)$, then

$$b(p, x, \xi; \theta) \equiv \Delta(p, x, \xi; \theta)^{-1} s(p, x, \xi; \theta), \tag{7}$$

where Δ is a J by J matrix whose (j, r) element is given by:

$$\Delta_{jr} = \begin{cases} \frac{-\partial s_r}{\partial p_j}, & \text{if r and j are produced by the same firm;} \\ 0, & \text{otherwise.} \end{cases}$$

Given the markups, or $b(p, x, \xi; \theta)$, and our model for marginal costs, (4), the first order conditions can be rearranged to yield

$$ln(mc_j) = ln(p_j - b_j(p, x, \xi; \theta) - \lambda VER_j) = w_j \gamma + \omega_j. \tag{8}$$

Note that in (8), the VER, as modeled, looks like a specific (as opposed to an *ad valorum*) tariff. That is, the VER raises prices by an amount in excess of cost plus markup. It is this aspect of the VER that may have led firms to adjust their product mix by upgrading (as documented empirically by Feenstra, and as modeled theoretically by Das and Donnenfeld, 1987, and Krishna, 1987).

The first-order condition in (8) is restrictive in several ways. First, it assumes that the same tax is placed on each firm. It has been suggested that since the VERs were allocated according to a formula that placed heavy weight on past market shares, it penalized the smaller upstart firms more heavily. Honda, in particular, claimed that they were more constrained in the early years of the VER, while other firms were less so. To investigate this possibility, our robustness analysis includes runs that estimate separate tax rates for large and small Japanese firms (where the division is admittedly somewhat arbitrary).

Note, however, that the first-order condition in (8) does not require that the tax be placed on each unit produced, but only on the marginal units. MITI might exempt some initial level of production from any political pressure. For our purposes, the level of the exemption might vary across firms, as long as the marginal tax rate was the same. Depending on how we modeled exemptions, they might once again place a discontinuity in the firms' reaction functions which might in turn lead to existence problems. We assume that either the exemptions do not cause problems or else that the tax rate is in fact applied to all units of production.

We also investigate the robustness of our results to the assumption that equilibrium is Nash in prices. The effect of any change in the equilibrium assumption will be to change the definition of the markups, or $b(p, x, \xi; \theta)$, in equation (7). One familiar alternative to our Bertrand assumption (Nash in prices) is to assume that firms play a Cournot game (Nash in quantities). The problem with this is that few, if any, industry observers seem to believe that, in the automobile industry, firms really set quantities and let the Walrasian auctioneer set the prices that clear markets. From Bresnahan (1981) on, researchers have modeled imperfect competition in the automobile industry in a Bertrand fashion. One might, however, posit a Nash game in which Japanese firms set quantities (subject to the export limits set by MITI), but the rest of the firms set prices. This is an approach empirically adopted by Feenstra and Levinsohn (1995) and coined Mixed Nash. Another possibility is that the VER somehow "taught" the Japanese firms to collude, and these colluding firms played a Bertrand game with the rest of the world. In section 6, we examine the robustness of our results by estimating the model under the Cournot, the Mixed Nash, and the collusion assumptions.[9]

[9] Readers interested in the derivation of the Mixed Nash first order conditions and the resulting markups are referred to Appendix I of the NBER working paper version of this paper. The markups from the Cournot game are familiar from the previous literature.

In concluding, we would like to stress that our estimates do not assume the VER raised prices in every year. If it had no effect on prices in a particular year, we ought to estimate a λ which is within estimation error of zero in that year.

This completes the discussion of the theory underlying our structural model. The key parameters to be estimated are those characterizing the distribution of tastes in the population, $\bar{\beta}$, σ, and α, those determining marginal costs γ, and the tax rates associated with the VERs, the λ's. The parameters on the demand side will permit us to evaluate how consumer welfare changes with the VER. These plus the cost side parameters allow us to estimate the effect of the VERs on the distribution of profits. The λ's measure the implicit tax on Japanese cars and allow us to compute the revenue foregone by the implementation of a VER (modeled essentially as an export tax by Japan) instead of a tariff imposed by the U.S. (assuming a tariff could be implemented without changing any of the other details of the problem, including the cars that are marketed in the U.S.). One needs these pieces of information, or something very close to them, to evaluate this strategic trade policy.

4. Estimation and Computation

We closely follow the estimation methods detailed in BLP. Here we outline those methods referring the interested reader to BLP for details.

Overview. As in an OLS or two-stage least squares estimation procedure, we base our estimates on a set of moment restrictions. In particular, we assume that the unobservables defined by the model, evaluated at the true values of the parameters, are mean independent of a set of exogenous instruments, z. Formally,

$$E[\xi_j(\theta_0) \mid z] = E[\omega_j(\theta_0) \mid z] = 0, \tag{9}$$

Equation (9) implies that the unobservables are uncorrelated with any function, $H_j(\cdot)$, of the instruments. Defining

$$G^J(\theta) = \frac{1}{J}\sum_{j=1}^{J} E\left[H_j(z)\begin{pmatrix}\xi_j(\theta)\\ \omega_j(\theta)\end{pmatrix}\right], \tag{10}$$

equation (10) implies

$$G^J(\theta_0) = 0.$$

Following the literature on Generalized Method of Moments (GMM) (Hansen, 1982) then, we choose as our estimate of θ that value that comes "closest" to setting the sample analog of the moments in equation (9) to zero. This sample analogue is

$$G_J(\theta) = \frac{1}{J}\sum_{j=1}^{J} H_j(z)\begin{pmatrix}\xi_j(\theta)\\ \omega_j(\theta)\end{pmatrix}. \tag{11}$$

The GMM estimator then minimizes

$$\|G_J(\theta)\|_{A_J}, \tag{12}$$

where for any vector y, $\|y\|_{A_J} = y'A_J y$, and where the matrix A_J converges in probability to some positive definite matrix A (we use the sample analogue of $EG_J(\theta_1)G_J(\theta_1)'$, where θ_1 is an initial consistent estimate of θ_0, as our A_J). Under suitable regularity conditions this estimate is consistent and asymptotically normal with covariance matrix detailed below.

To make use of the method, we must be able to calculate the unobservables as functions of the data at different values of the parameter vector. BLP provides a simple method for doing this computation and we follow this method exactly.

We turn next to the choice of instruments, z.

Instruments. The estimation method as outlined requires us to find a vector of observables, the z vector, that are mean independent of the unobservables (and are in that sense "econometrically exogenous"), and then use functions of them, the $H_j(z)$, as instruments. Since all the equilibrium notions discussed above imply that the p and q of every product are functions of the (ξ, ω) pairs of all products, we do not want to place price and quantity in the z vector. This is precisely the same reasoning that leads to the use of instruments for price and quantity in the analysis of demand and supply in homogeneous goods markets.

As in the analysis of homogeneous goods markets we look for observables that shift the demand and cost functions to use as the components of z. In the differentiated products framework these include the characteristics of *all* the products marketed (their size, fuel efficiency, acceleration, etc.), or the observed x vectors, as well as the variables, such as wage rates, that determine costs conditional on product characteristics, or the components of the observed w vectors that are not included in x.[10]

Note that the observed characteristics of all the products marketed in a given year are included in z, and the value of the instrument for any given product, the $H_j(\cdot)$, can be any function of z. In oligopolistic differentiated products markets the price of each good depends on the characteristics and prices of all goods marketed (thus markups will be lower for products which have many competitors with similar characteristics). As a result the value of the efficient instrument

[10] Of course just as in the homogeneous product model, to the degree that there are unobserved cost and demand factors that are correlated with our observed characteristics, our parameter estimates will be inconsistent. Indeed, once we start considering dynamic models in which product characteristics are endogenous, the restrictions we are currently using for identification become questionable. As a result we are exploring alternative identifying assumptions in our current work (see the discussion in BLP).

for any given product will be a function of the x and w vectors of *all* the products marketed (see Chamberlin, 1986, for a discussion of efficient instruments given conditional moment restrictions.) In the appendix, we develop an easy to compute approximation to the efficient instruments; these are used in our estimates.

Panel Data. The data set we actually use is not a single cross section, but a panel data set that follows car models over all years they are marketed. It is likely that the demand and cost disturbances of a given model are more similar across years than are the disturbances of different models. Correlation in the disturbances of a given model marketed in different years will affect the variance-covariance matrix of our parameter estimates. As a result, we use estimators that treat the sum of the moment restrictions of a given model over time as a single observation from an exchangeable population of car models. That is, replacing product index j by indices for model m and year t, we define the sample moment condition associated with a single model as

$$g_m(\theta) \equiv \sum_t H_{mt}(z) \binom{\xi_{mt}(\theta)}{\omega_{mt}(\theta)}$$

and then obtain our GMM estimator by minimizing our quadratic form in the average of these moment conditions across models. As noted in BLP, this is not likely to be the most efficient method for dealing with correlation across years for a given model, but it does produce standard errors that allow for arbitrary correlation across years for a given model and arbitrary heteroscedasticity across models.[11]

5. Policy Details, Data, Results, and Interpretation

This section begins with a discussion of the details of how the VER worked as they relate to implementing our procedures, and then turns to the available data and some of its more important features. Next we discuss the variables included in the utility function (3), and the marginal cost function (4). The results of our base case scenario are presented next, and the section concludes with interpretation of these results.

[11] Unlike BLP the standard errors we present here do not correct for simulation error in the computed market shares. We were able to increase the number of simulation draws to the extent that this error should not be important.

Some facts about the VERs

Moving from the oligopoly model described in section 3 to the data requires a more detailed discussion of exactly how the VER worked. As noted in the introduction, the VER was initiated in May 1981 and at that point total exports were limited to 1.68 million cars. In 1984, this figure increased to 1.85 million. In 1985, Japan voluntarily agreed to extend its already nominally voluntary export restraint, and from 1985 through early 1992, exports were limited to 2.30 million. Following President Bush's visit to Japan, the allocation was reduced back to 1.65 million in 1992. The VER was formally lifted in 1994.

A reasonable first pass at the data might include figures on firm-level allocations and shipments. However even if this data were available it would not suffice for the questions of interest. For example, one might note that firms just met their allocation, but it could still be that the quota was just barely binding, hence Japanese prices might not rise appreciably. On the other hand, it could be that some firms met their allocations, and some did not, and the overall effect might be ambiguous. Yet again, it could be that firms did not sell their entire allocations because they were worried about possible repercussions of inadvertently exceeding the limits. Finally, it could be that firms faced continual pressure from MITI to limit exports to the U.S. and, while MITI might have been hesitant to commit to a lower aggregate limit, it may have pressured firms in subtle ways to keep prices high and sales low. The bottom line is that data on allocations and sales are less informative than one might initially guess, and this is why a structural model is especially useful.[12]

The VER was structured such that cars produced by Japanese firms in the United States did not count against the VER. This production via direct foreign investment (dfi) was an empirically important phenomenon. Beginning with Honda's Marysville plant in 1982, Japanese firms responded to the VER by producing in the U.S. By 1990, Honda, Nissan, Toyota, Mazda, and Mitsubishi were producing in the U.S.. In our base case, the VER dummy variable was set to zero for all Japanese models that had production facilities in the U.S., although the profits accruing to these

[12] The situation is actually much worse than the previous discussion indicates as reliable figures on the allocations are simply not available. Professor Gary Saxonhouse kindly provided the data, attributed to MITI, that he has on allocations and shipments. They indicate that from 1981 to 1986 *every* firm managed to hit its allocation *exactly* and no firm ever missed by even one vehicle. We find these figures simply not credible, as they appear manufactured more for political purposes than for econometric analyses. In this context we note that though it is hard for us to verify the MITI figures, we have made some rough calculations. Difficulties arise mainly because our sales data are by calendar year while the MITI figures are by VER-year (May through April), and the MITI figures refer to shipments and these need not equal sales, although over time these two should more or less even out. Though the reader should keep these caveats in mind, when we did investigate we found that the MITI figures do not mesh well with the actual sales figures.

models were classified as Japanese profits. For cars produced in both Japan and the U.S. (and prominent examples of this for the latter part of our sample period are the Honda Accord and the Toyota Camry), this amounts to assuming that the marginal car sold was produced in the U.S.[13] We experiment with the assumption that the marginal car was produced in Japan, and hence that the VER dummy should be set to one for these models, in section 6.

The VER was also structured such that cars imported from Japan and sold under a U.S. brand were counted against the VER. These so-called captive imports were cars usually produced by Mitsubishi, Suzuki, and Isuzu and sold under the Dodge/Chrysler or Geo labels by Chrysler and General Motors respectively. In the estimation, we carefully account for these captive imports as their quantities are significant. In the sensitivity analyses, we experiment with ignoring captive imports and see if our policy conclusions are altered. It is unclear whether the profits from these cars should accrue to their Japanese manufacturers or the U.S. firms whose name they bear. We somewhat arbitrarily assume that profits accrue to the U.S. firm in this case, although the truth is surely somewhere between these two polar cases.

We now turn to a discussion of the data used in the estimation.

Data

All of our product-level data are obtained from the Automotive News Market Data Book (annual issues). These data include information on most engineering specifications of the automobiles marketed. The data span the period 1971 to 1990. In terms of the theory presented in Section 2, these data comprise the product attributes. They include continuous characteristics such as the car's horsepower, weight, length, width, wheelbase, engine displacement, and EPA miles per gallon rating. The data also include binary variables such as whether air conditioning, power steering, power brakes, and automatic transmission are standard equipment. Each model is in fact available in many variants (termed trim levels) and the list of standard equipment and specifications typically varies across trim levels. In order to keep the number of products computationally manageable, we include only the base model for each nameplate. It is important, then, that the price variable be that which also applies to the base model, and this is done.

We have list prices for each product. This is not ideal, but we think it is the best that can be done with our present data sources. The alternative is something akin to the average transaction price, where the average is taken for all purchases of a given nameplate. Such data are in fact

[13] For a more detailed examination of how dfi works in a model of oligopoly and quotas, see Levinsohn (1989).

available (but are proprietary) for many, though not all, models in the later years of our sample. It turns out that transactions prices for a given model are almost always higher than its list price. This is because very few cars are actually purchased without any options, and the purchase of options drives up the transaction price. Without detailed information on the relationship between options and transaction prices, the transactions prices are of limited use.[14]

We also make use of some macroeconomic data. These variables include exchange rates, consumer price deflators (in order to put all prices into real terms), the prime interest rate, the Gross National Product, and foreign wages. These are obtained from annual issues of the Economic Report of the President and the OECD Main Economic Indicators. Finally, we require information about the number of households and the distribution of income in the United States. These data are obtained from the Current Population Survey.[15]

We next consider some general trends in key variables. Table 1 provides some market averages, while Table 2 focuses more narrowly on trends in U.S. and Japanese competition. Table 1 lists the number of models, average sales and real price, and four key attributes for 1971-1990. It is clear that the number of models climbed fairly steadily until 1988, while the average sales per model declined. The deflated price of automobiles has risen steadily since 1974, although a noticeably larger than average blip appears in 1981, the year the VERs were initiated, and then again in 1982. Note also, however, that a smaller blip in prices occurred in 1980, a year before the introduction of the VER's, and there is an equally large series of increases in real prices between 1985-1987. Moreover, an almost *identical* series of increases occur in the variable, "Air" which provides the fraction of models in which air conditioning was standard equipment, and this suggests that the price increases may not be "pure price increases" but rather may reflect quality upgrading.

A measure of acceleration is given by horsepower divided by weight. This variable declined during the 1970's and rose during the 1980's. Vehicle size, measured as length times width has generally fallen. Cars have become better equipped, and this is proxied by the inclusion of air conditioning as standard equipment. In 1971, no car had it, while almost one third did by 1990. Finally, we include a measure of the cost of driving: miles driven on one dollar's worth of gas. This variable has generally trended upwards, although the oil shocks are apparent. An important message to take from Table 1 is that most of the variables exhibit significant trends, some well before the VERs, and we will want to account for this phenomenon in our empirical work.

[14] For some cursory evidence on the average transactions prices, see Table 1 and accompanying discussion in the NBER working paper version of this paper.

[15] All of our data are available on request by electronic mail. To obtain the data, send a request by e-mail to JamesL@umich.edu. The data will be sent by e-mail as a MIME attachment. The programs are similarly available.

The first two columns of Table 2 compare sales weighted average real list prices of Japanese and domestic cars. From 1973 to 1979, prices of domestic vehicles stayed relatively constant. Either coinciding with the imposition of the VER in 1981, or one year prior to it, U.S. prices started to increase, and they continued to increase steadily throughout the rest of the sample. Japanese prices, on the other hand, began a fairly steady climb in 1976, several years *prior* to the VERs. Indeed, the largest annual jump in Japanese prices occurred between 1977 and 1978, well before the imposition of the VER. This suggests the possible importance of using data prior to the VERs when investigating the effects of the VER. Put another way, if Table 2 began with 1981 data, it would appear that the VER had very strong influences on Japanese prices. When we note that these prices were increasing prior to 1981, the evidence becomes less clear. The last four columns of Table 2 give sales and market shares. Prior to the imposition of the VER, the Japanese market share was rising, from 5.7 percent in 1971 to 21.3 percent in 1981. This was mostly at the expense of U.S. market share which fell from 86.6 to 74.0 percent, a fact that led some (but not all) of the Big Three auto makers to press for import relief.

One message suggested by Tables 1 and 2 is that there were many trends in the industry both pre- and post-1981. Prices and quantities do seem to change around 1981, but they exhibit as large or larger changes both before and after, and around 1981 we also seem to see a large change in the product mix.

To throw further light on the issues related to the VER, we consider a simple OLS hedonic regression of prices against characteristics and a combination of trends and time dummies (Table 3). The regressors include four vehicle attributes (horsepower/weight, size, miles per dollar (MP$G), and air conditioning as standard), separate trends for the US (the omitted region), Europe, and Japan, as well as dummy variables for each of the three regions, the lagged and current exchange rate, and the current exchange rate interacted with region dummies. Appended to this list of regressors are year-specific dummy variables for Japan (the VER dummies) and the U.S.(the DOM dummies). The estimated regression had 2217 observations and an R^2 of .815.

All included vehicle characteristics except MP$ contribute positively to ln(price) in a precise way. The coefficient on MP$ is negative and significant. Region dummy variables suggest that, conditional on other included characteristics, European products sell at a premium. The precisely estimated coefficient on the overall trend indicates that prices are trending upwards. We pick up very little exchange rate pass-through except in the case of the German DM.

The coefficients on the VER and DOM dummy variables address a key question at hand: what was the relationship between the advent of the VERs and prices? The estimated coefficients on the

VER dummies in Table 3 are all *negative* and some are significantly so. While we are hesitant to draw conclusions from a hedonic regression, these results are nonetheless surprising in light of what seems to be the common wisdom. After accounting for trends and changes in vehicle characteristics, Japanese prices *fell* or at least did not seem to rise during the VER years. If the VER had the expected effect of increasing Japanese prices, then perhaps any fall in Japanese prices would have been greater absent the VER. During the same period, the coefficients on the domestic dummy variables are usually positive. The bottom line is that simple least squares analysis yields puzzling results, but, due to the lack of any underlying theory, it is hard to know what to make of them. We turn now to results from the estimated structural model.

Results

Recall that the structural parameters to be estimated are the means and variances of the distribution of the taste parameters in the utility function, the parameters of the cost function, and the implicit taxes associated with the VERs. We estimate means and variances of the tastes for: horsepower divided by weight (HP/WT), vehicle size, whether air conditioning is standard (AIR), miles driven on one dollar's worth of gasoline (MP$), and for the utility associated with the outside alternative (the constant). We have experimented with other vehicle attributes and, in BLP, we report that the estimated elasticities and resulting markups are robust to reasonable changes. One variable that does *not* appear in our list of attributes is a measure of reliability as given by a Consumers' Report rating. While we have such data for several years, it has severe problems in a time series context since ratings are relative to other vehicles in a given year. Hence, the definition of the variable is changing year by year. Moreover inclusion of the reliability index never seemed to matter. We note that the problems caused by not including more characteristics are somewhat attenuated by the fact that the model explicitly allows for characteristics not included in the specification (our unobserved characteristics).

On the cost-side, we include a constant as well as the following vehicle attributes: ln(HP/WT), ln(SIZE), and AIR. We include region dummies for Europe and Japan, as well as trends for the U.S., Europe, and Japan. Finally, we also include the log of the exchange rate of the exporting country (lagged one year) and the log of the wage rate in the producing country. We experimented with the contemporaneous exchange rate and found its effect was always about zero and imprecisely estimated.

We include VER dummies for each year since 1981, the year the policy was implemented. These dummy variables are set to one if the VER applies to that automobile model. As noted above, our base case assumes Japanese models produced in the U.S. did not count against the VER, while captive imports did. Note that this implies that Japanese wages and the yen to dollar exchange rate are determinants of costs for captive imports while U.S. wages are determinants of costs for the Japanese models produced in the U.S.

The estimates for our base case and their standard errors are given in Table 4. We begin with a discussion of the demand side parameters. When interpreting these parameters, it is important to keep in mind that demand for a particular car is driven by the maximum, and not by the mean, of the utilities heterogeneous consumers place on that car. Hence, there are two ways to explain why cars with, say, high HP/WT are popular. Either a high mean for the distribution of tastes for HP/WT or a large variance of tastes will have a tendency to increase the share of consumers who buy cars with large values of HP/WT. The results in Table 4 show that the means ($\bar{\beta}$'s) are all highly significant. The standard deviations of the taste parameters for Size and MP$ are also significant. The magnitudes of the standard deviations suggest that relative to their means, there is the most variance in the value of MP$.

On the cost-side, we find that each attribute contributes positively to marginal cost and almost all of their coefficients are quite precisely estimated. Japanese and European cars cost more to produce and transport, even after conditioning on wages and exchange rates. Domestic marginal costs are trending upwards, while Japanese and European marginal costs are trending slightly downwards. The elasticity of marginal cost with respect to wages is just over a third, not unreasonable for a production process with so large a materials component, while exchange rate pass-through is about zero. This last result is somewhat surprising, but experimentation suggests that it is robust. Exchange rates just do not seem to matter much. This finding contrasts to other estimates of exchange rate pass-through (see Feenstra, Gagnon, and Knetter (1993)), but our estimates are based on on more disaggregated data and on a more detailed model of the industry.

There are several ways to interpret the magnitude of the utility and cost parameters. One way which is easy to understand and captures the information on both the utility and cost sides of the model is to examine price-marginal cost markups. These markups depend on the demand elasticities implied by the $\bar{\beta}$'s and σ's as well as the marginal cost function parameters (all of which have been jointly estimated.) A representative sample of these markups for a handful of 1990 models representing the quality spectrum is presented in Table 5.[16] These estimates appear quite reasonable and are generally in line with other studies. The standard errors of the markups are

[16] All 2217 markups are available on request.

presented in column 4 and imply that the markups are quite precisely estimated. (A discussion of how the standard errors are computed is given below in the "Implications" subsection.)

The coefficients on the VER dummies address the following question: Suppose the VER was instead implemented as a specific tax on Japanese automobiles, *and* no other aspect of the model changed. What is the level of that tax that would generate equilibrium prices equal to those we observe when we have the VERs? A coefficient (or tax) of zero, would imply that the VER was not binding, while larger values correspond to a larger implicit tax. These coefficients are given in the bottom panel of Table 4.

In 1981, 1982, and 1983, the point estimates are about zero with a standard error between \$187 and \$248. In these years, the point estimates imply that the VER had almost no effect on prices, and we cannot reject that the effect was nil. In 1984 and 1985, the point estimates of the implicit tax rise to \$403 and \$361 respectively, but these estimates have standard errors of \$243 and \$303. Again, we cannot reject the hypothesis that the VER was not binding, although it should be noted that two standard errors encompass a wide range of implicit taxes; *i.e.* while we cannot reject that the VER had no effect in 1984 and 1985, neither can we reject that the implicit tax was in the range of \$600-\$800. We adopt as our null hypothesis, though, the absence of any price effect of the VER and are unable to reject this null for 1981-85. It is perhaps not surprising that the VERs had no effect in 1981, as they were not implemented until mid-year. However, the lack of any effect on equilibrium prices in 1982 and 1983 is likely to be surprising to some observers. Goldberg, for example, finds a large effect of the VER in 1983, the first year of her sample. Nonetheless, our result is robust to the many different variants of our model we have run.

Moreover, the available raw data are consistent with our results. The figures in table 2 indicate that total Japanese sales in the U.S. were below the VER limit in every year until 1986, the first year we estimate a significant VER dummy. It should be stressed that the export limits themselves are not used at all in our estimation algorithm, and hence provide some independent support for our results. We note again, though, the differences between calendar year and VER year and between sales and shipments that make this comparison problematic. Further, the figures in Table 2 have not have not been adjusted for the nuances imposed by dfi and captive imports.[17]

There are several reasons why we find the VER did not initially bind. The most important of them is that demand was low when the VERs were initially implemented. In 1981 the U.S. was

[17] When we make our best guess of the number of vehicles that count against the VER, we find that in *no* year did sales about equal the VER limit, although in 1983 sales were close to the limit (due to a surge in captive imports) and in 1986 sales fell only about 130,000 short of the 2.3 million limit. In most other years our guess was noticeably below the limits.

both in the midst of a recession, and had a prime interest rate over 18 percent. The prime rate did not fall to below 10 percent until 1985, and as late as 1984 it was over 12 percent. This type of economic environment affects an industry as cyclical as the automobile industry very adversely. Thus, a simple interpretation of the insignificant estimates of the VER dummy parameters for 1981-1983 is that in the middle of a severe recession, the VERs were set at a level that did not bind. Indeed, the VERs may well have been agreed to by the Japanese precisely because the Japanese realized that the promise of export restraints at the agreed level was both politically expedient and economically inexpensive at the time the agreement was made. We return to the impact of macro variables on our results in the robustness discussion below.

In 1986, the VER begins to have a statistically significant effect on prices in that we can no longer reject that the implicit tax was zero. In 1986, the point estimate of the implicit tax is $675 (with a standard error of $307). With an average price of Japanese cars at about $8,200, the VER is equivalent to about a 8.2 percent tax per Japanese car. (Recall the tax is specific, so it is much larger in percentage terms for inexpensive cars and less for costly ones.) The largest effects of the VERs are from 1987 to 1989, and this is again consistent with the notion that business cycles matter in this industry. During these years, the VER was equivalent to a tax of between $1277 (with a standard error of $458) and $1558 (with a standard error of $353.) In 1990, the estimated implicit tax falls to a still hefty $1063. Our estimate of the effect of the VER in 1990, though, is not very robust and should be interpreted with caution. (For a more extensive discussion of this point, see section 6.)

These are large effects and, by 1990, are somewhat surprising. For example, even with the fore-mentioned problems in comparing shipments or sales data to quota allocations, Nissan was surely not exporting its allocation at the end of our sample. Many industry observers have noted that although the VERs were still in effect in 1990 (they remained so until 1994), they were not important due to the increased direct foreign investment by the Japanese into the U.S. Our base case results suggest otherwise. What might be going on here? There are multiple mutually non-exclusive explanations. Note that the VER dummies enter the firms' first order conditions such that it captures price increases above those explained by marginal cost (including region dummies and region-specific trends) and the mark-up. A significant VER dummy would occur if Japanese firms were induced, either by MITI, or by the U.S. or by cartelization to keep prices high and sales low relative to the no-VER Bertrand equilibrium. Indeed dynamic models involving political variables and/or cartel behavior could be built to rationalize this process. Another possible explanation is that while some firms may not have been constrained by the VER, others were. For example, while Nissan probably was not constrained, Mitsubishi (due to the many captive imports supplied to Chrysler) almost certainly was. Indeed, one reason exports under the VER were increased in the mid-1980's was probably the increase in captive imports. A third explanation is that some of the large estimated VER dummies in the later 1980's and especially 1990 are not always robust to specification testing. We return to a more detailed examination of these alternatives below.

Thus far, all description of the VERs has been positive, not normative. Sure, prices went up, but this is not all that surprising (though the timing and magnitude of the rises might be.) Insights from the strategic trade policy theoretical literature suggest that the profit-enhancing effect of the VER might make protection welfare enhancing in spite of the concurrent loss of consumer welfare. We turn now to a fuller investigation of the implications of our estimates on both profits and on consumer welfare.

Implications

In order to investigate the effects of the VER on profits and consumer welfare, we need to know what the industry equilibrium would have been in the absence of the VER. To determine that equilibrium, we set λ (the implicit tax) to zero, and solve for the vector of prices and vector of quantities for which the firms' first order conditions hold and for which consumers maximize utility conditional on those prices. This assumes both that the equilibrium without the VER is also Nash in prices and that the equilibrium is unique (or at least that we solve for the relevant one.) It further assumes that the distribution of automobile characteristics would not have changed in the absence of the VER. This last assumption is probably more reasonable in the short run and less so in the longer run, since the time needed to change models is typically measured in years, not months. We only recompute the equilibrium for years in which λ was significantly larger than zero. This is admittedly a somewhat arbitrary choice, but computational constraints played a role in this decision.

When we solve for the equilibrium that would obtain when λ is set to zero, we implicitly are making use of estimated parameters. Since the estimated parameters have standard errors associated with them, so does the new equilibrium. We compute these standard errors when evaluating policy implications of our estimates. Doing so is non-trivial. The ability to put standard errors on policy implications is one great advantage of econometric methods over the calibration methods that are commonly used in evaluating trade policy. However, because the policy implications are typically complicated non-linear transformations of the parameters, computational constraints have limited the extent to which standards errors have been presented.

One solution (the "delta method") is to linearize the policy implications in the parameters. We avoid this linearization and instead take a more direct Monte Carlo approach. To implement this, we take $n = 175$ draws of parameters from the estimated asymptotic normal distribution of the parameters.[18] For each of these draws, we resolve the entire model and then calculate the implied policy implications. The empirical standard deviation of these policy implications, across the n draws, is then a consistent estimate of the true standard error of the policy implications.

We first turn our attention to the profit-shifting side of the story. The effects of the VER on prices and profits are given in Table 6. There, we report the sales-weighted average price of Japanese, American, and European cars as well as profits with and without the VER, the difference between the VER and no VER cases, and the standard error of this difference. These figures are given for each year in which we estimated a statistically significant VER coefficient. As expected, the prices of Japanese cars were driven up by the VER. [19] Note that in a Nash pricing game, when at least some of the products are strategic complements, prices can rise by either more or less than the amount of the tax. Our estimates indicate that both occur.

The issue of strategic complements and substitutes is an important one in this study. In differentiated products price-setting models, it is usual to think of prices as being strategic complements. In these cases, an exogenous rise in a competitor's price will raise own-firm prices. The intuition that price-setting models yield strategic complements comes from linear models in which the competitor's price affects the intercept, but not the slope, of the own-product demand curve. However, in typical discrete choice models both the intercept and the slope change as the rivals prices change: the demand curve shifts out and becomes more price-sensitive. The change in the slope can occur because those customers who shift away from the rival product are those who are more price-sensitive than average. These price elastic consumers might induce a *decrease* in own-firm prices in response to a rival's price increase. Thus, we can obtain either strategic complements or substitutes.[20]

The VER increased Japanese prices fairly dramatically. Prices increased by around \$750 in

[18] We experimented with more draws but found that computational time went up linearly while standard errors remained stable. With substantially fewer draws, estimates became noisy.

[19] Note that since the VERs induce a different combination of cars to be purchased, throughout this table the weights used when the VER is assumed operative are different than the weights when it is not.

[20] It is well-known that prices of products j and k are strategic complements if and only if $\partial^2 \pi_f / \partial p_j \partial p_k > 0$. This cross-price second derivative is

$$\frac{\partial s_j}{\partial p_k} + \sum_{r \in \mathcal{J}_f} (p_r - mc_r)[\frac{\partial^2 s_r}{\partial p_j \partial p_k}]$$

1986 and this figure rose to $1687 in 1987. The increase then fell to around $800 by 1990. These changes in prices are measured with standard errors of $35 or less.

We find that the prices of U.S. autos were little affected by the VER. U.S. prices rose by only about $200 in 1987 and 1988 due to the VER. In other years the increase was less than about $80 and the standard error was never more than $28. Recall that in our model, consumers are heterogeneous. Our results suggest that as Japan raised prices, price sensitive consumers switched to U.S. automobiles, and, as a result, markups did not increase much. However, while prices of domestically produced cars were not much changed due to the VER, sales increased significantly, and this is reflected in the increased profits earned by U.S. firms. The second set of columns in Table 6 indicates that U.S. profits increased by about $3.09 billion in 1987 and by $2.76 billion in 1988. Even in 1986, when we find the VER had a relatively small effect on prices, U.S. profits increased by about $1.6 billion due to the VER. This is the profit-shifting aspect of a strategic trade policy. The standard errors of the difference in profits is large (t-statistics are somewhere between 1 and 2.) Hence, while point estimates suggest that U.S. profits increased, these estimates are not precise. (Since profits depend implicitly on hundreds of elasticities, it may not be that surprising that even if each elasticity is tightly estimated, the change in the level of profits is not that tightly estimated.)

While U.S. profits were much increased by the VER, Japanese profits did not fall a corresponding amount. Our estimates imply that Japanese profits were basically unaffected by the VER. In 1986, point estimates imply that Japanese profits rose by $111 million while in 1988 they fell by $110 million. In other years, the figure is somewhere between these two. These are not large numbers. Neither are they precisely estimated. The standard error of the difference in Japanese profits is on the order of $300-$400 million. Two factors contributed to the relatively small decrease in Japanese profits. First, apparently a large fraction of consumers had relatively inelastic demands for the Japanese models; these consumers preferred paying the increased Japanese prices to shifting their demand to other models. Second, with the VER, as opposed to a tariff, the Japanese firms did not have to pay the implicit tax. Instead they kept the "revenue" such a tax would have generated

---

In our model,

$$\frac{\partial}{\partial p_k}\frac{\partial s_j}{\partial p_j} = \int \left[\frac{\partial s_j(\nu)}{\partial p_j} s_k(\nu) + \frac{\partial s_k(\nu)}{\partial p_j} s_j(\nu)\right] dF(\nu),$$

where recall that ν is the vector of consumer characteristics. Since the first term in the integrand will usually dominate, the integrand will be large and negative when the price-sensitive consumers are likely to shift to good k. If this effect is large enough for products j and k, it will more than compensate for the positive $\frac{\partial s_j}{\partial p_k}$ in the expression for $\partial^2 \pi_f / \partial p_j \partial p_k > 0$.

and this is reflected in the higher prices. VERs are sometimes referred to as bribes to the foreign firm, for Japanese profits might have been lower had the VER instead been implemented as a tariff or regular quota.[21]

The theoretical literature has recognized that a quota (or, in this case, VER) might act to raise industry profits. Our point estimates imply this was indeed the case, although our estimates of the change in profits resulting from the VER have relatively large standard errors.

Profits are only part of the economic welfare equation. Another key component is consumer welfare. We compute the compensating variation in the following way. First take a draw from the estimated distribution of tastes and the distribution of income. This draw can be thought of as a simulated household. Next, compute which product gives the highest utility at the VER (i.e. the actual) prices and the resulting utility. Now find the income which generates the same level of utility at the non-VER prices (i.e. the prices we obtained when we solved for the industry equilibrium in the absence of the VER). The change between this income and the initial draw on the household's income is the compensating variation.[22] To estimate the expected compensating variation for a randomly chosen household, we do this a large number of times and take the average. Multiplying this expectation by the number of households in the economy gives the total compensating variation. The estimates in tables 7 and 8 use 10,000 draws (though we have conducted much of the exercise with 100,000 draws and the results only change in the third decimal point).

Table 7 provides estimates, for 1987, of how household-level compensating variation changes with the imposition of the VER. This table begins to address the question of who bears the burden of the VER. The first two rows look at the economy-wide aggregates. The first row gives the average change in the price of the good actually purchased. There we note that prices rise on average \$18. Most households (about 90 percent) did not purchase a car in a given year, and for these households, the price change was zero. Hence the average figure hides a great deal of variation. The standard deviation of the change in the price of the good purchased under the VER is \$277, while at least one product's price rose by \$2369 and another's fell by \$499. The latter is due to the presence of strategic substitutes. The economy-wide average compensating variation figure implies that the VER cost the household, on average, \$41, although this figure was as great as \$2366 for some households. Again due to the strategic substitutes, some households were made \$483 better off by the VER.

[21] It should be noted, however, that Japanese profits are actually somewhat lower than what is reported in Table 6. This is because some of the difference between price and cost is kept by the dealer, and these dealers are typically domestically owned.

[22] A further discussion of this method and other applications are found in Pakes, Berry, and Levinsohn (1993).

The next three pairs of lines in Table 7 decompose the economy-wide averages. We estimate that the imposition of the VER would, on average, leave those households who (under the VER) purchased a car \$317 worse off. This figure reflects the twin facts that auto purchasers were adversely affected by a significant amount and that most households in a given year are *not* auto purchasers. The \$317 figure is aggregated over households who purchased a Japanese car (when the VER was imposed) and those that purchased a domestic car. These two groups fared quite differently under the VER. On average the VER cost households that bought a Japanese car \$1242. On the other hand, the VER cost households that purchased a domestic car only about \$30. Consumers of domestic cars themselves were not that adversely affected by the VER.

Table 8 gives the bottom line on our evaluation of the VERs as a strategic trade policy. There, we compute the components of aggregate welfare for each of the years in which the VER was estimated to be binding in our base case. The first column gives the change in domestic profits. The second column gives the compensating variation and is negative since the protection cost domestic consumers. The third column gives the sum of the first two columns and represents the net change in welfare for the VER *as it was actually implemented.* The fourth column presents the foregone tariff revenue (had an import tariff been used instead of the implicit export tax we model.) The fifth column then lists the welfare gain that would have resulted if the VER was instead implemented as a tariff, *and* no other change occured in the nature of the equilibrium. The bottom row of the table gives the cumulative totals over the multiple years, and that is the row on which we focus. Standard errors of all figures are given in parentheses. All figures are in 1983 dollars. In current (1996) dollars, the amounts would be inflated by around 50 percent.

The first effect of the VER was to increase the pure profits of U.S. firms by about 10.2 billion dollars. It is hard to evaluate the magnitude of this figure. To put it into some perspective, though, our estimates imply that the pure profits (not including fixed costs) from Japanese automobile sales in the U.S. in 1990 were about 7.6 billion (1983) dollars, while the profits of U.S. firms in 1990 were about \$23.1 billion. It seems that the profit shifting effects of the VERs was not negligible.

On the other hand, the burden placed on U.S. consumers was not negligible either as the compensating variation of the VERs was just over \$13.1 billion. The standard error of this figure is \$2.48 billion. The net change in welfare due to the VERs was about -\$2.9 billion. Due to the large standard error on the change in profits, the net change has a relatively large standard error– \$7.56 billion.

When one evaluates the typical trade policy, the welfare components number three: profits, consumer welfare, and tariff revenue. The VER was implemented such that it gave the latter of these back to the Japanese firms or government. Suppose the U.S. had instead opted for the tariff that would have resulted in the same industry equilibrium observed under the VER. We assume that all imports from Japan generate tariff revenue, and this includes captive imports as well as the made-in-Japan portion of production of models which were also produced in the U.S. (i.e. Camrys made in Japan raise tariff revenue while those made in Kentucky do not.) This policy would have generated almost \$11.2 billion dollars in revenue for the U.S. government. The foregone revenue with a VER is sometimes referred to as the bribe paid in order to induce Japan to agree to the policy in the first place. Our (precise) estimates suggest this was a hefty bribe. When this figure is added to the net change computed in the third column of Table 8, the welfare gain from the VERs totals \$8.34 billion. Our point estimates suggest that if the government been able to impose a tariff without changing any of the other conditions in the market, the implied protection of the automobile industry could have enhanced U.S. welfare for exactly the sort of reasons that came out of the early theoretical models of trade policy and imperfect competition. Nonetheless, this net figure has a standard error as large as the net figure itself. In terms of what *was* precisely estimated, we conclude that the decrease in consumer welfare was about equal to the foregone tariff revenue.

Does this suggest that tariffs on Japanese automobiles would be in the U.S. economic interest? There are several reasons why this might not be so. For example, we do not model retaliation (nor, though, do most theoretical models of strategic trade policy.) Surely one reason to implement a VER instead of an outright tariff or quota was that the VER bribed the Japanese government into not retaliating. Furthermore, a tariff directed solely at Japanese products would violate the GATT. Also, we are assuming that the imposition of a tariff would not cause Japanese firms to stop marketing some of their models in the U.S. If models were pulled off the U.S. market then consumers with inelastic, as well as those with elastic, demand for that model would be adversely affected.

Just as there are good reasons, though, to wonder whether the \$8.341 billion figure might be unrealistically high, there are also good reasons to believe it is too low. First, we have estimated the welfare effects of the VERs as actually implemented, and there is no reason to believe that they were set to optimize welfare. Second, our theoretical and empirical work did not account for monopoly rents accruing to U.S. workers in the automobile industry.

6. Sensitivity Analyses

Along the way to the punchlines provided in the last section, we have made several possibly objectionable assumptions. For example, we assumed the firms played a Bertrand game, that firms' underlying cost functions were the same, and that the export limits were either binding or not binding on all firms in any given year. We chose not to ignore dfi or captive imports, but did ignore some key ways in which the macro-economy might affect automobile demand. We also assumed that quality changes were exogenous. In this section, we ask, do changes in these assumptions affect our major conclusions.[23]

Table 9 provides results from seven of the alternative specifications we tried. The base case was estimated under a Bertrand assumption. We investigate how robust our results are to a Cournot as well as to a Mixed Nash assumption. We also investigate the possibility that the VER led to collusion among Japanese firms while the Japanese firms collectively maintained a Bertrand strategy vis a vis non-Japanese firms.

There are many ways to compare results across specifications: demand elasticities, markups, profits (which use information from each of the previous two), and the coefficients on the VER dummies. Since the focus of this study is on trade policy, we opt for the latter.

The first column of Table 9 replicates the VER multipliers from our base case. The second column has the estimates obtained under the assumption of Cournot behavior. These estimates are obtained from a structural model in which the firms' first order conditions and resulting markup have been amended to reflect the Cournot assumption.[24] With the Cournot assumption, we find that the multiplier on the 1990 VER dummy variable is less precisely estimated, and we can no longer reject the hypothesis that the VER did not bind that year. On the other hand, the dummy variable for the VER in 1985 becomes statistically significant. Other than 1985 and 1990, the VER is found to be binding in the same set of years as when price setting was assumed to be Bertrand (though the magnitude of the VER multiplier was quite a bit larger in 1986, and somewhat smaller in the other years than in our base case).

A possibly more realistic alternative to Bertrand is the Mixed Nash case. Here the Japanese

[23] There is also the issue of the shape of our objective function, in particular the presence of local minima, and the ability of our numerical procedures (which includes a choice of starting values and of stopping tolerances) to find its overall minimum. We experimented with alternate starting values and tolerances and sometimes found the minimization algorithm stopping at local minima that were slightly different than the overall minima reported in the text. In particular some of these alternate runs indicated that the VER had a larger effect in 1985 and a smaller effect in 1990 than the results reported in the text suggest (though these dummies were *never* significant in 1981 to 1984, and were always significant between 1987 and 1989). The VER dummy coefficients on 1985 and especially 1990 are least stable. Our selected base case is the most representative of our results, but it may be that the VER had a larger effect in 1985 and a smaller effect in 1990 than the base case results suggest. The results for these years, then, should be interpreted with caution.

[24] All else is as in the base case. i.e. We use the same: i) starting values; ii) model for dfi and captive imports; and iii) the same simulation draws as in the base case.

firms set quantities while other firms set prices.[25] If one believed that there were strict export limits given to the Japanese firms, a model where these firms set quantities seems more plausible. The VER multipliers we obtained when we re-estimated our model under the Mixed Nash assumption are given in the third column of Table 9. They are, in terms of magnitudes of estimates and standard errors, very close to those obtained under the Bertrand assumption. The VERs bind in all the same years and the implied specific tax is about the same across the two specifications. We conclude that while it may be reasonable to estimate the model under alternative static equilibrium concepts, it doesn't really seem to impact the policy conclusions drawn. A caveat is in order, though. While the results are robust to the various specifications of the equilibrium, it remains the case in all results presented that the demand and cost sides of the model have been estimated simultaneously. In principle, one could estimate the demand-side of the model alone and then use the estimated elasticities to investigate the cost side of the model. This would be more flexible and would impose less structure on the utility function parameter estimates. We have tried to do this, and are unable to obtain precise estimates of many of the parameters of interest. We conclude that, absent more data, the equilibrium first-order conditions on the cost side contribute to the precision of the demand-side estimates. We are currently working on developing methods, using consumer-level data, that might allow one to estimate the demand-side independently of any equilibrium assumptions. See, for example, Berry, Levinsohn, and Pakes (1997).

The fourth column of Table 9 presents the VER multipliers from the collusion case. The thought experiment here is that the VER induced Japanese firms to collude. From a modeling perspective, this essentially changes the firms' first-order conditions such that all Japanese firms act like a single multi-product oligopolist. The estimated VER multipliers are quite similar to the base case, although the 1985 coefficient becomes statistically significant while the 1986 coefficient becomes statistically insignificant. All point estimates, though, are within one standard deviation of the base case estimates.

Since dfi production was not subject to any restraints, one would expect the presence of dfi to diminish the trade restraining aspect of the VERs. On the other hand, we would not necessarily expect dfi to render the VERs ineffectual for three reasons. First, it takes time to build an automobile plant and bring it up to capacity. Second, the amount of capacity built in the U.S. is determined by perceptions of the future implications of that capacity, including its potential political ramifications, and there is good reason to believe that the U.S. capacity of Japanese models was not built

[25] Once again, we are simply assuming that such an equilibrium exists and then showing that it does exist at the estimated parameter values.

up as fast as otherwise would have been expected. For example, although production costs in 1994 were widely believed to be lower in the U.S. than in Japan for the same vehicle, there were no major new plants on the drawing boards, and this is due in part to political concerns. (Restrictions on Japanese capacity in the U.S. were reported to be discussed during President Bush's "auto" trip to Japan.) Third, if production costs were lower in Japan than in the U.S., the VER might still bind even with the presence of dfi. To investigate how treating dfi differently (and effectively ignoring it) might alter our results, the model is re-estimated ignoring the effects of dfi on the underlying structural model. These results are presented in the fifth column of Table 9.

The general pattern is one in which the VER dummies are similar to the base case, with a few exceptions. When we ignore dfi, the VER appears to be binding in 1985 and not binding in 1986 or 1990. More importantly, ignoring dfi does not affect our finding that the VER contributed to higher prices for Japanese cars in the later 1980's, but not in the first four years of the policy. Although the coefficient estimates of the VER dummies are not that different from the base case, the welfare implications are. This is because the implicit tariff revenue foregone is much higher when dfi is ignored, since no-dfi assumption would attribute foregone tariff revenue to all the cars actually produced by Japanese firms in the U.S.

The next column of Table 9 gives the results when we ignore the role of captive imports. This specification is estimated in order to determined whether ignoring captive imports (as previous studies have) matters to our main results. We find that the results of the no captive imports specification are quite similar to the base case. The main difference is that by ignoring captive imports, it appears that the VER significantly raised prices in 1985, and possibly also in 1984, while our base case indicates the contrary. Although the coefficients are not that different for the no captive imports case, the welfare consequences of ignoring the captive imports are large. Like the story with dfi, this occurs because with captive imports, the consequences for foregone tariff revenue are large.

The next-to-last column of Table 9 presents the VER dummies when an attempt is made to account for macroeconomic influences on the demand system. These runs included GNP and the prime interest rate as linear terms in the utility function. These terms do not have random coefficients. The GNP variable had a positive coefficient on it (with a t-statistic of around 2) while the prime interest rate had a negative coefficient on it (with a t-statistic of around -10). Including these variables is quite *ad hoc.*. In principle, one can argue that shifts in income are already captured by the inclusion of household income in the utility function. Also, while the interest rate certainly matters, it just as certainly would not enter a structural dynamic model of automobile demand in the simple manner with which we experiment. We include these variables, though, to investigate, albeit loosely, whether including some macroeconomic demand shifters substantively alters our conclusions about the VERs. As VER dummies in the last column indicate, our results are not that different. We find that the 1985 coefficient becomes significant, while the 1986 and 1990 coefficients become insignificant, and the other coefficients are slightly smaller in magnitude. This suggests that ignoring macroeconomic influences may make the VER look slightly more binding than in fact it was.

Finally, we investigate the robustness of our results to the implicit assumption that all firms have the same underlying cost function. There are of course many ways in which cost functions might differ across firms. As a first pass at this issue, we allow firm-specific fixed effects in the cost function and re-estimate the model with these 26 fixed effects. The estimated VER multipliers from this experiment are given in the last column off Table 9. The main difference between this case and the base case is that the 1986 coefficient becomes statistically insignificant.

We also conducted some sensitivity analyses in which more than just yearly VER dummies were estimated. Recall that the base case imposed that the export limits were either binding or not binding on all Japanese firms in a given year. Anecdotal evidence suggests that perhaps the smaller Japanese firms were more constrained by the VER (at least in the early years). An approach which would be robust to this and other contingencies would be to estimate separate VER dummies for each firm in each year. This, though, is computationally infeasible and would, in any case, generate imprecise estimates. A middle ground between the infeasible ideal and the base case is to estimate one multiplier for the Big Two in Japan (Toyota and Nissan) and another for the other Japanese firms. The results suggested that the smaller firms might have been more constrained in the first few years of the VER, although the effect is imprecisely measured. The anecdotal evidence may have a grain of truth to it.

The VER, as modeled, enters costs as a year-specific dummy variable for Japanese firms beginning in 1981. There are myriad stories that might lead to an observationally equivalent estimating equation. The VER effects show up as deviations from costs, conditional on trends and cost-shifters, in the very particular way implied by the firms' first order conditions. We estimated the model with quadratic region-specific trends instead of the linear ones. The VER coefficients for 1986 and 1990 cease to be significantly different from zero.

As a "common sense" test of our results, we re-estimated the model with two other sets of country-year dummy variables. Each enter the cost function just as the VER did. In one specification, the model was re-estimated using "VER" dummies for every year, even those prior to the VER. If we were to consistently find significant effects of the "VERs" in the years prior to 1981, one might wonder whether the results for the years after 1980 were really picking up the trade restraints or something altogether different. The coefficients on the VER dummy variables were insignificantly different from zero throughout the 1970s. During the years that the VER was actually in place, the only changes relative to the base case are that the coefficients on the VER in some years were slightly smaller and usually less precisely estimated.

The model was also estimated with year-specific dummy variables for domestic firms during the 1980's. Again, had these dummy variables matched the pattern of the Japanese VER multipliers, one might wonder whether something other than the VER might be motivating the base case results. We found that only one of the 10 year-specific dummy variables for domestic firms was significantly different from zero– about what we would expect if all were zero at the 90 percent level of statistical significance. The point estimates were all quite small.[26]

The model was estimated allowing tastes to differ in the 1970s. This was done by allowing the means of the tastes distributions to differ in the 1970s while constraining the variances of the taste distributions to remain constant over the sample. This was required in order to keep the estimation computationally feasible. The results suggest that the marginal utility of size and air conditioning was lower in the 1970's, a period during which gas prices were high. We can reject that tastes were constant over the sample. The estimated VER coefficients, though, remain substantively the same as the base case.

Finally, we have assumed that quality changes are exogenous. That is, while upgrading occurred, we do not model this as a policy-induced response. Our results, then, are conditional on the exogeneity of the existing product attributes.

From Table 9 and our other sensitivity analyses, we conclude that our base case results are reasonably robust to several plausible alternative specifications. Because the results seem so robust, it is natural to question why they do not replicate the messages of the existing literature on the effect of the VER. Our results are not very much at odds with Feenstra's and the differences are explainable. Feenstra (1988) found substantial quality upgrading, and we also find this in our data. Feenstra found that the VER was initially binding. His methods and data, though, were quite different. He did not use data for the decade prior to the VERs, and he estimated separate sets of coefficients for Japanese cars. Finally, his methods are much more in the spirit of a reduced form,

[26] We do not report the full results here, because this was the one specification for which we had troubles in reliably minimizing the objective function. This problem appeared to arise because of the large number of non-linear cost-side parameters being estimated.

and the underlying framework is not nearly as structural as our equilibrium oligopoly model. (His work also predates ours by about a decade, and many of the econometric tools at our disposal were not available then.) When we use the same years of data as Feenstra and employ simple hedonic regressions as he did, we find that we replicate the gist of his results. The VERs appear binding in the early years, but their magnitude is small and not always precisely estimated. When we add our oligopoly structure, but continue to allow Japanese cars to have different cost functions, we no longer find that the VERs were initially binding. We conclude that what differences there are between our results and Feenstra's emanate from different interpretations to the hedonic regression; we have a model which allows us to impute changes in that regression to changes in underlying costs, in markups, and in the implications of trade policy (the VER dummies).

Though Goldberg's (1995) methods are a lot more similar to ours than Feenstra's, her results, unlike those of Feenstra, are, in some respects, quite different from ours. In particular, as noted earlier, Goldberg finds that the VER was binding in the early years. We investigated several possible sources of this difference but could not account for it. Goldberg did not use data from years prior to the VER, had fewer years of data for the later 1980's, and made use of consumer-level data using the Consumer Expenditure Survey. When we estimate our model using only the same years of data as Goldberg, we continue to find that the VER did not initially bind. We allowed for trends in the data that Goldberg does not account for. We again re-estimate our model excluding all trends. Again, our results remain at odds with Goldberg's. As noted above, ignoring or including macroeconomic variables, direct foreign investment, and/or captive imports do not substantively change our results, and hence could not reconcile them with those reported by Goldberg. We speculate on two possible reasons for the difference. We account for the econometric endogeneity of price, while Goldberg does not. Using consumer-level automobile purchase data (not used in the analysis of this paper), we find that ignoring this endogeneity substantially biases the estimates and that the resulting elasticities are affected. Since these elasticities are key to the analysis, this may account for the difference. Secondly, the demand structures in this study and in Goldberg's are quite different and this too may matter.

7. Conclusions and Caveats

Our estimates indicate that the VERs affected prices, although not necessarily in the years most expected. They raised Japanese prices and domestic sales. The profits of domestic firms increased substantially while those of Japanese firms were less affected. Domestic consumer welfare fell, also quite significantly, and this burden fell disproportionately on consumers with relatively inelastic demands for Japanese products. The "give-away" to Japan in terms of foregone tariff revenue was very large. In sum, our point estimates imply that if tariffs could have been instituted without setting off other changes in the market (in particular with no changes in the cars marketed in the U.S. and no retaliatory responses by the Japanese), strategic trade policy could have enhanced U.S. economic welfare.

When the first economic models of strategic trade policy were being introduced, most of the founders of that literature went to some length to make clear that their models did not mean the traditional arguments for free trade had become inapplicable. This paper may be the first detailed econometric study of a strategic trade policy and similar caveats are in order.

We have computed the standard errors around each of these policy implications. These suggest researchers ought to be circumspect about making policy conclusions even when the individual parameters of the structural model are themselves precisely estimated. We are unable to precisely estimate the impact of the VER on profits. The foregone tariff revenue and the compensating variation, though, are precisely estimated and our estimates suggest that these two components of welfare about cancel each other out.

Standard errors around policy conclusions are only one reason to view the results in this paper with care. The underlying structural model is not a dynamic model and this has multiple implications. First, automobiles are a durable good and expectations about how long the VER was expected to last surely impacted production and consumption decisions. Second, as noted earlier, we take as exogenous both the set of products firms bring to the market and the attributes of those products. A more involved dynamic model would allow one to model these endogenously. Third, we do not model myriad other aspects of the dynamics of automobile purchases such as financing, expected depreciation and resale value. Fourth, on the demand-side, we have assumed that the underlying distributions of tastes are constant. If tastes changed over time due for example to learning, these changes might impact our results. In sum, we realize these dynamic issues are important, and this too adds to our caution in interpreting the results.

References

Berry, Steven (1994) "Estimating Discrete Choice Models of Product Differentiation," *RAND Journal of Economics*, **25**, 242-262.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995) "Voluntary Export Restraints on Automobiles: Evaluating a Strategic Trade Policy," NBER Working Paper 5235.

Berry, Steven, James Levinsohn, and Ariel Pakes (1993) "Applications and Limitations of Some Recent Advances in Empirical Industrial Organization: Price Indexes and the Analysis of Environmental Change," *American Economic Review*, **83**, 247-252.

.

Berry, Steven and Ariel Pakes (1994) "Alternative Empirical Models of Product Differentiation," In process. Yale University.

Bresnahan, T. (1981) "Departures from Marginal-Cost Pricing in the American Automobile Industry," *Journal of Econometrics*, **17**, 201-227.

————- (1987) "Competition and Collusion in the American Automobile Oligopoly: The 1955 Price War," *Journal of Industrial Economics*, **35**, 457-482.

Cardell, N.S. (1991) "Variance Components Structures for the Extreme Value and Logistic Distributions," Mimeo, Washington State University.

Chamberlin, G. (1986) "Asymptotic Efficiency in Estimation with Conditional Moment Restrictions," *Journal of Econometrics*, , 305-334.

Crain Communications (1984-1988 issues) *Automotive News Market Data Book*. Detroit: Crain.

Dinopoulos, Elias and Mordechai Kreinin (1988) "Effects of the U.S.-Japan Auto VER on European Prices and U.S. Welfare," *Review of Economics and Statistics*, , 484-91.

Dixit, Avinash (1988) "Optimal Trade and Industrial Policies for the U.S. Automobile Industry," in *Empirical Methods for International Trade*, ed. Robert Feenstra. Cambridge: MIT Press.

Feenstra, Robert (1988) "Quality Change under Trade Restraints: Theory and Evidence from Japanese Autos," *Quarterly Journal of Economics*, , 131-146.

————— (1984) "Voluntary Export Restraint in U.S. Autos, 1980-81: Quality, Employment, and Welfare Effects," in *Structure and Evolution of Recent U.S. Trade Policy*, ed. Robert Baldwin and Anne Krueger. Chicago: University of Chicago Press.

————— (1995) "Survey of Trade Policy Studies," in *Handbook of International Economics, Volume III*, ed. Gene Grossman and Ken Rogoff . Amsterdam: North Holland.

Feenstra, Robert and James Levinsohn (1995) "Estimating Markups and Market Conduct with Multidimensional Product Attributes," *Review of Economic Studies*, **62**, 19-52.

Feenstra, Robert, Joseph Gagnon, and Michael Knetter (1993) "Market Share and Exchange Rate Pass-Through in the World Automobile Industry," NBER Working Paper 4399.

Fuss, Melvyn, Steven Murphy, and Leonard Waverman (1992) "The State of the North American and Japanese Motor Vehicle Industries: A Partially Calibrated Model to Examine the Impacts of Trade Policy Changes," NBER Working Paper 4225.

Goldberg, Pinelopi (1994) "Trade Policies in the U.S. Automobile Industry," *Japan and the World Economy*, **6**, 175-208.

————- (1995) "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, **63**, 891-952.

Hansen, L. (1982) "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, **50**, 1029-1054.

Hausman, J.A. and D. Wise (1978) "A Conditional Probit Model for Qualitative Choice: Discrete Decisions Recognizing Interdependence and Heterogeneous Preferences," *Econometrica*, **46**, 403-426.

Krishna, Kala (1989) "Trade Restrictions as Facilitating Practices," *Journal of International Economics*, **26**, 251-270.

————- (1987) "Tariffs versus Quotas with Endogenous Quality," *Journal of International Economics*, **23**, 97-112.

Krugman, Paul (1994) *Peddling Prosperity: Economic Sense and Nonsense in the Age of Diminished Expectations.* New York: Norton.

Levinsohn, James (1994) "International Trade and the U.S. Automobile Industry: Current Research, Issues, and Questions," *Japan and the World Economy*, **6**, 335-357.

———— (1989) "Strategic Trade Policy When Firms Can Invest Abroad: When are tariffs and quotas equivalent," *Journal of International Economics*, **27**, 129-146.

Newey, W (1990) "Efficient Instrumental Variables Estimation of Nonlinear Models," *Econometrica*, **58**, 809-839.

Pakes, Ariel (1986) "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, **54**, 755-784.

Pakes, Ariel, Steven Berry and James Levinsohn (1995) "Automobile Prices in Market Equilibrium," *Econometrica*, **63**, 841-890.

Appendix: An Approximation to "Optimal" Instruments

Following Chamberlin (1986), the efficient set of instruments when we have only conditional moment restrictions is:

$$H_j(z) = E\left[\frac{\partial \xi_j(\theta_0)}{\partial \theta}, \frac{\partial \omega_j(\theta_0)}{\partial \theta} \middle| z\right] T(z_j) \equiv D_j(z) T(z_j), \tag{1}$$

where $T(z_j)$ is the matrix that normalizes the error matrix, i.e.

$$T(z)'T(z) = \Omega(z)^{-1} \equiv E((\xi,\omega)(\xi,\omega)'|z)^{-1}.$$

This formula is very intuitive: larger weights should be given to the observations that generate disturbances whose computed values are very sensitive to the choice of θ (at $\theta = \theta_0$). Unfortunately $D_j(z)$ is typically difficult to compute. Since the required derivatives are a function of prices, to calculate $D_j(z)$ we would have to calculate the pricing equilibrium for different $\{\xi_j, \omega_j\}$ sequences, take derivatives at the equilibrium prices, and then integrate out over the distribution of such sequences.

We propose to replace the expectation $D_j(z)$ with the appropriate derivatives evaluated at the expectation of the unobservables. To construct such derivatives, we take the following steps:

(i) Obtain an initial estimate $\hat{\theta}$ from an initial run using cruder instruments.

(ii) Use $\hat{\theta}$ to construct exogenous estimates of δ and mc: $\hat{\delta} = x\hat{\beta}$ and $\hat{mc} = w\hat{\gamma}$.

(iii) Solve the first order conditions of the model for equilibrium prices, $\hat{p}$, and shares, $\hat{s}$ as a function of $\hat{\theta}$, $\hat{\delta}$, $\hat{mc}$ and x.

(iv) Construct the functions defining the unobservables of the model evaluated at the exogenous predictions: $\hat{\xi}(\theta) = \xi(\hat{p}, \hat{s}, \hat{\delta}, x, \theta)$ and $\hat{\omega}(\theta) = \omega(\hat{p}, \hat{s}, \hat{\delta}, \hat{mc}, x, \theta)$. Then use as our (admittedly biased) estimate of the optimal instrument vector

$$\hat{D}_j(z) = \left(\frac{\partial \hat{\xi}_j(\hat{\theta})}{\partial \theta}, \frac{\partial \hat{\omega}_j(\hat{\theta})}{\partial \theta}\right).$$

Further detail and some intuition for a simpler model can be found in the 1995 NBER version of this paper.

TABLE 1

Some Descriptive Statistics

Year	No. of Models	Quantity (1000's)	Price ($'000)	HP/Wt	Size	Air	MP$
1971	92	86.892	7.868	0.490	1.496	0.000	1.850
1972	89	91.763	7.979	0.391	1.510	0.014	1.875
1973	86	92.785	7.535	0.364	1.529	0.022	1.819
1974	72	105.119	7.506	0.347	1.510	0.026	1.453
1975	93	84.775	7.821	0.337	1.479	0.054	1.503
1976	99	93.382	7.787	0.338	1.508	0.059	1.696
1977	95	97.727	7.651	0.340	1.467	0.032	1.835
1978	95	99.444	7.645	0.346	1.405	0.034	1.929
1979	102	82.742	7.599	0.348	1.343	0.047	1.657
1980	103	71.567	7.718	0.350	1.296	0.078	1.466
1981	116	62.030	8.349	0.349	1.286	0.094	1.559
1982	110	61.893	8.831	0.347	1.277	0.134	1.817
1983	115	67.878	8.821	0.351	1.276	0.126	2.087
1984	113	85.933	8.870	0.361	1.293	0.129	2.117
1985	136	78.143	8.938	0.372	1.265	0.140	2.024
1986	130	83.756	9.382	0.379	1.249	0.176	2.856
1987	143	67.667	9.965	0.395	1.246	0.229	2.789
1988	150	67.078	10.069	0.396	1.251	0.237	2.919
1989	147	62.914	10.321	0.406	1.259	0.289	2.806
1990	131	66.377	10.337	0.419	1.270	0.308	2.852
all	2217	78.804	8.604	0.372	1.357	0.116	2.086

Notes: The entry in each cell is the sales weighted mean. Prices are in constant 1983 dollars.

Quantity is the average sales (in thousands) per model.

HP/WT is in 100's of HP divided by 1000's of lbs (i.e. # HP divided by 10's of lbs.)

Size is vehicle width (in inches) times vehicle length (in inches) divided by 1000.

Air is one if air conditioning is standard equipment and zero otherwise.

MP$ is the 10's of miles one can drive on a 1983 dollar's worth of gasoline.

TABLE 2

Prices and Quantities in the U.S. Automobile Industry:
The changing balance of U.S. and Japanese Firms

year	Average Domestic Price ($'000)	Average Japanese Price ($'000)	Domestic Sales (1000's)	Japanese Sales (1000's)	Domestic Market Share	Japanese Market Share
1971	8.204	5.147	6925.510	454.722	86.633	5.688
1972	8.188	5.506	7830.860	365.186	89.216	4.161
1973	7.540	6.248	7438.593	320.709	93.221	4.019
1974	7.586	6.238	6709.888	375.712	88.655	4.964
1975	7.900	6.136	6728.847	653.643	85.348	8.291
1976	7.856	6.039	8099.279	744.676	87.609	8.055
1977	7.687	6.106	7770.924	1041.266	83.702	11.216
1978	7.597	6.788	8076.884	1006.493	85.495	10.654
1979	7.494	6.965	6779.265	1335.962	80.326	15.829
1980	7.758	6.585	5699.259	1409.649	77.316	19.123
1981	8.263	7.096	5331.731	1533.095	74.098	21.306
1982	8.722	7.414	4861.743	1597.300	71.410	23.461
1983	8.735	7.270	5731.447	1674.540	73.424	21.452
1984	8.816	7.624	7604.399	1735.902	78.311	17.877
1985	8.648	7.882	8086.050	2033.145	76.086	19.131
1986	9.223	8.229	7982.851	2357.163	73.316	21.649
1987	9.821	8.765	6794.617	2374.362	70.218	24.538
1988	9.968	8.754	7214.957	2389.055	71.707	23.744
1989	10.147	8.808	6382.100	2412.200	69.008	26.083
1990	10.295	9.205	5927.647	2395.638	68.170	27.551

TABLE 3

A First Pass at Examining the Effect of the VER on Automobile Prices

An OLS Hedonic Regression

Dependent Variable is ln(Price)		
Variable	Parameter Estimater	Standard Error
constant	2.248	0.044
ln(hp/wt)	0.593	0.027
ln(space)	1.038	0.056
ln(MP$)	-0.312	0.035
air	0.479	0.015
trend	0.021	0.004
japan	2.358	2.945
euro	2.357	0.436
jtrend	-0.006	0.018
etrend	-0.018	0.005
ln(e-rate)	-0.272	0.091
lag(ln(e-rate))	0.258	0.089
ln(e-rate)*japan	0.295	0.300
ln(e-rate)*euro	0.374	0.070
VER80	-0.199	0.078
VER81	-0.155	0.083
VER82	-0.156	0.114
VER83	-0.099	0.121
VER84	-0.148	0.135
VER85	-0.149	0.151
VER86	-0.120	0.115
VER87	-0.122	0.118
VER88	-0.191	0.129
VER89	-0.257	0.137
VER90	-0.280	0.150
dom80	-0.056	0.037
dom81	0.018	0.039
dom82	0.112	0.041
dom83	0.130	0.043
dom84	0.109	0.048
dom85	0.076	0.050
dom86	0.216	0.057
dom87	0.171	0.060
dom88	0.164	0.065
dom89	0.111	0.069
dom90	0.063	0.073

TABLE 4
Estimated Parameters of the Demand and Pricing Equations:
Base Case Specification
1971-1990 Data, 2217 observations

	Variable	Parameter Estimate	Standard Error
Demand Side Parameters			
Means ($\bar{\beta}$'s)	Constant	-5.901	0.712
	HP/Weight	2.946	0.486
	Size	3.430	0.342
	Air	0.934	0.199
	MP$	0.202	0.084
Std. Deviations (σ_β's)	Constant	1.112	1.171
	HP/Weight	0.167	4.652
	Size	1.392	0.707
	Air	0.377	0.886
	MP$	0.416	0.132
Term on Price (α)	$(-p/y)$	44.794	4.541
Cost Side Parameters			
	Constant	0.035	0.310
	ln(HP/Weight)	0.604	0.063
	ln(Size)	1.291	0.106
	Air	0.484	0.043
	Trend	0.018	0.004
	Japan	3.255	0.667
	Japan*trend	-0.036	0.008
	Euro	3.205	0.525
	Euro*trend	-0.032	0.006
	lagln(e-rate)	0.026	0.024
	ln(wage)	0.356	0.079
VER Dummies			
	ver81	-0.085	0.187
	ver82	-0.022	0.228
	ver83	0.001	0.248
	ver84	0.403	0.245
	ver85	0.361	0.303
	ver86	0.675	0.307
	ver87	1.558	0.353
	ver88	1.490	0.379
	ver89	1.277	0.458
	ver90	1.063	0.469

TABLE 5

A Sample from 1990 of

Estimated Price–Marginal Cost Markups

Based on Table 4 Estimates

	Price (in 1983 $)	Markup over MC $(p - MC)$	Std. Error of Markup	Markup as Fraction of Price
Mazda 323	$ 5,049	$ 1,219	$164	0.241
Nissan Sentra	$ 5,661	$ 1,451	$171	0.256
Ford Escort	$ 5,663	$ 1,653	$203	0.292
Chevy Cavalier	$ 5,797	$ 2,127	$209	0.367
Honda Accord	$ 9,292	$ 2,880	$198	0.310
Ford Taurus	$ 9,671	$ 3,352	$216	0.347
Buick Century	$ 10,138	$ 4,057	$231	0.400
Nissan Maxima	$ 13,695	$ 4,343	$255	0.317
Acura Legend	$ 18,944	$ 6,487	$383	0.342
Lincoln TownCar	$ 21,412	$ 8,206	$457	0.383
Cadillac Seville	$ 24,353	$ 10,231	$486	0.420
Lexus LS400	$ 27,544	$ 9,973	$646	0.362
BMW 735i	$ 37,490	$ 13,521	$692	0.361

TABLE 6
The Effect of the VER on Prices and Profits:

		Average Price in $1000's				Total Profits in $ millions			
		With VER	No VER	Diff.	Std.Err. of diff.	With VER	No VER	Diff.	Std.Err. of diff.
1986	Japan	8.253	7.506	0.747	0.017	6334	6222	111	351
	U.S.	9.107	9.074	0.034	0.009	27551	25927	1623	1662
	Europe	17.079	17.170	-0.091	0.013	3040	2974	66	171
1987	Japan	8.849	7.162	1.687	0.035	7908	7999	-90	426
	U.S.	9.496	9.304	0.192	0.034	24900	21814	3085	1467
	Europe	18.823	19.050	-0.227	0.020	3012	2863	148	162
1988	Japan	8.955	7.470	1.485	0.033	7544	7654	-110	424
	U.S.	9.625	9.424	0.201	0.028	26923	24159	2764	1568
	Europe	19.874	20.064	-0.189	0.018	2863	2752	111	154
1989	Japan	9.053	7.989	1.064	0.033	7353	7368	-14	453
	U.S.	9.888	9.805	0.083	0.017	24648	23064	1583	1410
	Europe	21.435	21.551	-0.116	0.020	3251	3167	84	173
1990	Japan	9.307	8.510	0.797	0.027	7612	7550	61	469
	U.S.	10.053	9.975	0.078	0.016	23123	21972	1151	1317
	Europe	18.639	18.722	-0.083	0.023	2302	2242	59	122

Average prices are sales-weighted averages. (Average prices do not match those on Table 2 due to treatment of direct foreign investment and captive imports.)

TABLE 7
Decomposing the Compensating Variation
Results from 1987

	Mean	Std.Dev	Min	Max	n
All Households:					
Average change in price of originally purchased good	0.018	0.277	-0.499	2.369	10000
Compensating Variation	-0.041	0.300	-2.366	0.483	10000
Only HH's who purchased a car:					
Average change in price of originally purchased good	0.161	0.814	-0.499	2.369	1120
Compensating Variation	-0.317	0.817	- 2.366	0.483	1120
Only HH's who purchased Japanese car:					
Average change in price of originally purchased good	1.208	1.149	-0.432	2.369	266
Compensating Variation	-1.242	1.012	- 2.366	0.426	266
Only HH's who purchased non-Japanese car:					
Average change in price of originally purchased good	-0.165	0.098	-0.499	-0.013	854
Compensating Variation	-0.030	0.457	- 2.063	0.483	854

Notes: The "originally purchased good" refers to the good purchased when the VER was in place.

TABLE 8
Aggregate Welfare and the VER
(Accounting for Direct Foreign Investment by Japanese Firms)
(in $ billion (1983))

year	Change in Domestic Profits	Compensating Variation	Net Change	Foregone Tariff Equivalent	Welfare Gain from Equivalent Tariff
1986	1.623	-1.636	-.013	1.337	1.323
	(1.662)	(.316)	(1.654)	(.566)	(1.792)
1987	3.085	-4.019	-.934	3.266	2.332
	(1.467)	(.797)	(1.617)	(.677)	(1.770)
1988	2.764	-3.338	-.574	3.012	2.437
	(1.568)	(.710)	(1.664)	(.692)	(1.838)
1989	1.583	-2.505	-.921	2.131	1.209
	(1.410)	(470)	(1.464)	(.708)	(1.641)
1990	1.151	-1.635	-.484	1.521	1.037
	(1.317)	(.360)	(1.371)	(.611)	(1.556)
Total	10.207	-13.135	-2.928	11.269	8.341
	(7.350)	(2.480)	(7.556)	(3.096)	(8.311)

Standard errors are in parentheses.

TABLE 9
Sensitivity Analyses

	Base Case	Cournot	Mixed Nash	Collusion	No DFI	No CI	Macro	Fixed Effects
VER81	-0.085	-0.255	-0.001	-0.075	-0.098	0.111	-0.080	0.014
	(0.187)	(0.201)	(0.205)	(0.203)	(0.227)	(0.208)	(0.144)	(0.167)
VER82	-0.022	-0.347	0.000	-0.094	0.033	0.083	-0.144	-0.197
	(0.228)	(0.251)	(0.248)	(0.246)	(0.281)	(0.225)	(0.178)	(0.204)
VER83	0.001	-0.423	0.117	-0.152	0.434	0.193	-0.183	-0.232
	(0.248)	(0.256)	(0.261)	(0.233)	(0.381)	(0.274)	(0.179)	(0.220)
VER84	0.403	0.069	0.542	0.323	0.374	0.577	0.177	0.204
	(0.245)	(0.279)	(0.255)	(0.223)	(0.309)	(0.294)	(0.200)	(0.217)
VER85	0.361	1.378	0.515	0.603	0.677	0.845	0.443	0.438
	(0.303)	(0.359)	(0.309)	(0.228)	(0.361)	(0.293)	(0.222)	(0.241)
VER86	0.675	1.301	0.883	0.490	0.555	0.769	0.304	0.212
	(0.307)	(0.369)	(0.318)	(0.253)	(0.412)	(0.328)	(0.228)	(0.268)
VER87	1.558	1.152	1.433	1.302	1.129	1.361	1.004	0.659
	(0.353)	(0.411)	(0.351)	(0.296)	(0.431)	(0.394)	(0.288)	(0.336)
VER88	1.490	1.184	1.579	1.494	1.184	1.635	0.906	1.378
	(0.379)	(0.443)	(0.391)	(0.343)	(0.518)	(0.459)	(0.313)	(0.382)
VER89	1.277	0.891	1.462	1.232	1.041	1.554	0.828	1.170
	(0.458)	(0.479)	(0.513)	(0.377)	(0.533)	(0.499)	(0.373)	(0.441)
VER90	1.063	0.570	1.231	1.248	0.837	1.156	0.403	1.259
	(0.469)	(0.517)	(0.502)	(0.387)	(0.564)	(0.517)	(0.399)	(0.430)

Standard errors are in parentheses.